

Chapter 1

Bit wizardry

We give low-level functions for binary words, such as isolation of the lowest set bit or counting all set bits. Sometimes the term ‘one’ is used for a set bit and ‘zero’ for an unset bit. Where it cannot cause confusion, the term ‘bit’ is used for a set bit (as in “counting the bits of a word”).

The C-type `unsigned long` is abbreviated as `ulong` as defined in [FXT: [fxttypes.h](#)]. It is assumed that `BITS_PER_LONG` reflects the size of an `unsigned long`. It is defined in [FXT: [bits/bitsperlong.h](#)] and usually equals the machine word size: 32 on 32-bit architectures, and 64 on 64-bit machines. Further, the quantity `BYTES_PER_LONG` reflects the number of bytes in a machine word: it equals `BITS_PER_LONG` divided by eight. For some functions it is assumed that `long` and `ulong` have the same number of bits.

Many functions will only work on machines that use two’s complement, which is used by all of the current general purpose computers (the only machines using one’s complement appear to be some successors of the UNIVAC system, see [358, entry “UNIVAC 1100/2200 series”]).

The examples of assembler code are for the x86 and the AMD64 architecture. They should be simple enough to be understood by readers who know assembler for any CPU.

1.1 Trivia

1.1.1 Little endian versus big endian

The order in which the bytes of an integer are stored in memory can start with the least significant byte (*little endian* machine) or with the most significant byte (*big endian* machine). The hexadecimal number `0x0D0C0B0A` will be stored in the following manner if memory addresses grow from left to right:

```
adr:  z   z+1   z+2   z+3
mem:  0D   0C   0B   0A   // big endian
mem:  0A   0B   0C   0D   // little endian
```

The difference becomes visible when you cast pointers. Let `V` be the 32-bit integer with the value above. Then the result of `char c = *(char *)&V;` will be `0x0A` (value modulo 256) on a little endian machine but `0x0D` (value divided by 2^{24}) on a big endian machine. Though friends of big endian sometimes refer to little endian as ‘wrong endian’, the desired result of the shown pointer cast is much more often the modulo operation.

Whenever words are serialized into bytes, as with transfer over a network or to a disk, one will need two code versions, one for big endian and one for little endian machines. The C-type `union` (with words and bytes) may also require separate treatment for big and little endian architectures.

1.1.2 Size of pointer is not size of int

If programming for a 32-bit architecture (where the size of `int` and `long` coincide), casting pointers to integers (and back) will usually work. The same code *will* fail on 64-bit machines. If you have to cast pointers to an integer type, cast them to a sufficiently big type. For portable code it is better to avoid casting pointers to integer types.

1.1.3 Shifts and division

With two's complement arithmetic division and multiplication by a power of 2 is a right and left shift, respectively. This is true for unsigned types and for multiplication (left shift) with signed types. Division with signed types rounds toward zero, as one would expect, but right shift is a division (by a power of 2) that rounds to $-\infty$:

```
int a = -1;
int c = a >> 1; // c == -1
int d = a / 2; // d == 0
```

The compiler still uses a shift instruction for the division, but with a 'fix' for negative values:

```
9:test.cc @ int foo(int a)
10:test.cc @ {
285 0003 8B442410      movl 16(%esp),%eax // move argument to %eax
11:test.cc @      int s = a >> 1;
289 0007 89C1         movl %eax,%ecx
290 0009 D1F9         sarl $1,%ecx
12:test.cc @      int d = a / 2;
293 000b 89C2         movl %eax,%edx
294 000d C1EA1F       shr1 $31,%edx // fix: %edx=(%edx<0?1:0)
295 0010 01D0         addl %edx,%eax // fix: add one if a<0
296 0012 D1F8         sarl $1,%eax
```

For unsigned types the shift would suffice. One more reason to use unsigned types whenever possible.

The assembler listing was generated from C code via the following commands:

```
# create assembler code:
c++ -S -fverbose-asm -g -O2 test.cc -o test.s
# create asm interlaced with source lines:
as -alhnd test.s > test.lst
```

There are two types of *right* shifts: a *logical* and an *arithmetical* shift. The logical version (`shr1` in the above fragment) always fills the higher bits with zeros, corresponding to division of unsigned types. The arithmetical shift (`sarl` in the above fragment) fills in ones or zeros, according to the most significant bit of the original word.

Computing remainders modulo a power of 2 with unsigned types is equivalent to a bit-and:

```
ulong a = b % 32; // == b & (32-1)
```

All of the above is done by the compiler's optimization wherever possible.

Division by (compile time) constants can be replaced by multiplications and shifts. The compiler does it for you. A division by the constant 10 is compiled to:

```
5:test.cc @ ulong foo(ulong a)
6:test.cc @ {
7:test.cc @      ulong b = a / 10;
290 0000 8B442404      movl 4(%esp),%eax
291 0004 F7250000      mull .LC33 // value == 0xcccccccc
292 000a 89D0         movl %edx,%eax
293 000c C1E803         shr1 $3,%eax
```

Therefore it is sometimes reasonable to have separate code branches with explicit special values. Similar optimizations can be used for the modulo operation if the modulus is a compile time constant. For example, using modulus 10,000:

```
8:test.cc @ ulong foo(ulong a)
9:test.cc @ {
53 0000 8B4C2404      movl 4(%esp),%ecx
10:test.cc @      ulong b = a % 10000;
57 0004 89C8         movl %ecx,%eax
58 0006 F7250000      mull .LC0 // value == 0xd1b71759
59 000c 89D0         movl %edx,%eax
60 000e C1E80D         shr1 $13,%eax
61 0011 69C01027      imull $10000,%eax,%eax
62 0017 29C1         subl %eax,%ecx
63 0019 89C8         movl %ecx,%eax
```

Algorithms to replace divisions by a constant with multiplications and shifts are given in [168], see also [346].


```
if ( (!a) ^ (!b) )
```

1.1.7 Average without overflow

A routine for the computation of the average $(x+y)/2$ of two arguments x and y is [FXT: [bits/average.h](#)]

```
1 static inline ulong average(ulong x, ulong y)
2 // Return floor( (x+y)/2 )
3 // Use:      x+y == ((x&y)<<1) + (x^y)
4 // that is:  sum == carries + sum_without_carries
5 {
6     return (x & y) + ((x ^ y) >> 1);
7 }
```

The function gives the correct value even if $(x+y)$ does not fit into a machine word. If it is known that $x \geq y$, then we can use the simpler statement `return y+(x-y)/2`. The following version rounds to infinity:

```
1 static inline ulong ceil_average(ulong x, ulong y)
2 // Use:      x+y == ((x|y)<<1) - (x^y)
3 // ceil_average(x,y) == average(x,y) + ((x^y)&1)
4 {
5     return (x | y) - ((x ^ y) >> 1);
6 }
```

1.1.8 Toggling between values

To toggle an integer x between two values a and b , use:

```
pre-calculate: t = a ^ b;
toggle:       x ^= t; // a <--> b
```

The equivalent trick for floating-point types is

```
pre-calculate: t = a + b;
toggle:       x = t - x;
```

Here an overflow could occur with a and b in the allowed range if both are close to overflow.

1.1.9 Next or previous even or odd value

Compute the next or previous even or odd value via [FXT: [bits/evenodd.h](#)]:

```
1 static inline ulong next_even(ulong x) { return x+2-(x&1); }
2 static inline ulong prev_even(ulong x) { return x-2+(x&1); }
3
4 static inline ulong next_odd(ulong x)  { return x+1+(x&1); }
5 static inline ulong prev_odd(ulong x)  { return x-1-(x&1); }
```

The following functions return the unmodified argument if it has the required property, else the nearest such value:

```
1 static inline ulong next0_even(ulong x) { return x+(x&1); }
2 static inline ulong prev0_even(ulong x) { return x-(x&1); }
3
4 static inline ulong next0_odd(ulong x)  { return x+1-(x&1); }
5 static inline ulong prev0_odd(ulong x)  { return x-1+(x&1); }
```

Pedro Gimeno gives [priv. comm.] the following optimized versions:

```
1 static inline ulong next_even(ulong x) { return (x|1)+1; }
2 static inline ulong prev_even(ulong x) { return (x-1)&~1; }
3
4 static inline ulong next_odd(ulong x)  { return (x+1)|1; }
5 static inline ulong prev_odd(ulong x)  { return (x&~1)-1; }

1 static inline ulong next0_even(ulong x) { return (x+1)&~1; }
2 static inline ulong prev0_even(ulong x) { return x&~1; }
3
4 static inline ulong next0_odd(ulong x)  { return x|1; }
5 static inline ulong prev0_odd(ulong x)  { return (x-1)|1; }
```

1.1.10 Integer versus float multiplication

The floating-point multiplier gives the highest bits of the product. Integer multiplication gives the result modulo 2^b where b is the number of bits of the integer type used. As an example we square the number 11111111 using a 32-bit integer type and floating-point types with 24-bit and 53-bit mantissa (significant):

```

a = 11111111 // assignment
a*a == 12345678987654321 // true result

a*a == 1653732529 // result with 32-bit integer multiplication
(a*a)%(2**32) == 1653732529 // ... which is modulo (2**bits_per_int)

a*a == 1.2345679481405440e+16 // result with float multiplication (24 bit mantissa)
a*a == 1.2345678987654320e+16 // result with float multiplication (53 bit mantissa)

```

1.1.11 Double precision float to signed integer conversion

Conversion of double precision floats that have a 53-bit mantissa to signed integers via [11, p.52-53]

```

1 #define DOUBLE2INT(i, d) { double t = ((d) + 6755399441055744.0); i = *((int *)&t); }
2 double x = 123.0;
3 int i;
4 DOUBLE2INT(i, x);

```

can be a faster alternative to

```

1 double x = 123.0;
2 int i = x;

```

The constant used is $6755399441055744 = 2^{52} + 2^{51}$. The method is machine dependent as it relies on the binary representation of the floating-point mantissa. Here it is assumed that, the floating-point number has a 53-bit mantissa with the most significant bit (that is always one with normalized numbers) omitted, and that the address of the number points to the mantissa.

1.1.12 Optimization considerations

Never assume that some code is the ‘fastest possible’. There is always another trick that can still improve performance. Many factors can have an influence on performance, like the number of CPU registers or cost of branches. Code that performs well on one machine might perform badly on another. The old trick to swap variables without using a temporary is pretty much out of fashion today:

```

// a=0, b=0   a=0, b=1   a=1, b=0   a=1, b=1
a ^= b; //      0   0       1   1       1   0       0   1
b ^= a; //      0   0       1   0       1   1       0   1
a ^= b; //      0   0       1   0       0   1       1   1

// equivalent to: tmp = a; a = b; b = tmp;

```

However, under some conditions (like extreme register pressure) it may be the way to go. Note that if both operands are identical (memory locations) then the result is zero.

The only way to find out which version of a function is faster is to actually do benchmarking (timing). The performance does depend on the sequence of instructions surrounding the machine code, assuming that all of these low-level functions get inlined. Studying the generated CPU instructions helps to understand what happens, but can never replace benchmarking. This means that benchmarks for just the isolated routine can at best give a rough indication. Test your application using different versions of the routine in question.

Never ever delete the unoptimized version of some code fragment when introducing a streamlined one. Keep the original in the source. If something nasty happens (think of low level software failures when porting to a different platform), you will be *very* grateful for the chance to temporarily resort to the slow but correct version.

Study the optimization recommendations for your CPU (like [11] and [12] for the AMD64, see also [144]). You can also learn a lot from the documentation for other architectures.

Proper documentation is an absolute must for optimized code. Always assume that nobody will understand the code without comments. You may not be able to understand uncommented code written by yourself after enough time has passed.

1.2 Operations on individual bits

1.2.1 Testing, setting, and deleting bits

The following functions should be self-explanatory. Following the spirit of the C language there is no check whether the indices used are out of bounds. That is, if any index is greater than or equal to `BITS_PER_LONG`, the result is undefined [FXT: [bits/bittest.h](#)]:

```
1 static inline ulong test_bit(ulong a, ulong i)
2 // Return zero if bit[i] is zero,
3 // else return one-bit word with bit[i] set.
4 {
5     return (a & (1UL << i));
6 }
```

The following version returns either zero or one:

```
1 static inline bool test_bit01(ulong a, ulong i)
2 // Return whether bit[i] is set.
3 {
4     return (0 != test_bit(a, i));
5 }
```

Functions for setting, clearing, and changing a bit are:

```
1 static inline ulong set_bit(ulong a, ulong i)
2 // Return a with bit[i] set.
3 {
4     return (a | (1UL << i));
5 }
1 static inline ulong clear_bit(ulong a, ulong i)
2 // Return a with bit[i] cleared.
3 {
4     return (a & ~(1UL << i));
5 }
1 static inline ulong change_bit(ulong a, ulong i)
2 // Return a with bit[i] changed.
3 {
4     return (a ^ (1UL << i));
5 }
```

1.2.2 Copying a bit

To copy a bit from one position to another, we generate a one if the bits at the two positions differ. Then an XOR changes the target bit if needed [FXT: [bits/bitcopy.h](#)]:

```
1 static inline ulong copy_bit(ulong a, ulong isrc, ulong idst)
2 // Copy bit at [isrc] to position [idst].
3 // Return the modified word.
4 {
5     ulong x = ((a>>isrc) ^ (a>>idst)) & 1; // one if bits differ
6     a ^= (x<<idst); // change if bits differ
7     return a;
8 }
```

The situation is more tricky if the bit positions are given as (one bit) masks:

```
1 static inline ulong mask_copy_bit(ulong a, ulong msrc, ulong mdst)
2 // Copy bit according at src-mask (msrc)
3 // to the bit according to the dest-mask (mdst).
4 // Both msrc and mdst must have exactly one bit set.
5 {
6     ulong x = mdst;
7     if (msrc & a) x = 0; // zero if source bit set
8     x ^= mdst; // ==mdst if source bit set, else zero
9     a &= ~mdst; // clear dest bit
```

```

10     a |= x;
11     return a;
12 }

```

The compiler generates branch-free code as the conditional assignment is compiled to a `cmov` (conditional move) assembler instruction. If one or both masks have several bits set, the routine will set all bits of `mdst` if any of the bits in `msrc` is one, or else clear all bits of `mdst`.

1.2.3 Swapping two bits

A function to swap two bits of a word is [FXT: [bits/bitswap.h](#)]:

```

1  static inline ulong bit_swap(ulong a, ulong k1, ulong k2)
2  // Return a with bits at positions [k1] and [k2] swapped.
3  // k1==k2 is allowed (a is unchanged then)
4  {
5      ulong x = ((a>>k1) ^ (a>>k2)) & 1; // one if bits differ
6      a ^= (x<<k2); // change if bits differ
7      a ^= (x<<k1); // change if bits differ
8      return a;
9  }

```

If it is known that the bits do have different values, the following routine should be used:

```

1  static inline ulong bit_swap_01(ulong a, ulong k1, ulong k2)
2  // Return a with bits at positions [k1] and [k2] swapped.
3  // Bits must have different values (!)
4  // (i.e. one is zero, the other one)
5  // k1==k2 is allowed (a is unchanged then)
6  {
7      return a ^ ( (1UL<<k1) ^ (1UL<<k2) );
8  }

```

1.3 Operations on low bits or blocks of a word

The underlying idea of functions operating on the lowest set bit is that addition and subtraction of 1 always changes a burst of bits at the lower end of the word. The functions are given in [FXT: [bits/bitlow.h](#)].

1.3.1 Isolating, setting, and deleting the lowest one

The lowest one (set bit) is isolated via

```

1  static inline ulong lowest_one(ulong x)
2  // Return word where only the lowest set bit in x is set.
3  // Return 0 if no bit is set.
4  {
5      return x & -x; // use: -x == ~x + 1
6  }

```

The lowest zero (unset bit) is isolated using the equivalent of `lowest_one(~x)`:

```

1  static inline ulong lowest_zero(ulong x)
2  // Return word where only the lowest unset bit in x is set.
3  // Return 0 if all bits are set.
4  {
5      x = ~x;
6      return x & -x;
7  }

```

Alternatively, we can use either of

```

return (x ^ (x+1)) & ~x;
return ((x ^ (x+1)) >> 1) + 1;

```

The sequence of returned values for $x = 0, 1, \dots$ is the highest power of 2 that divides $x + 1$, entry [A006519](#) in [312] (see also entry [A001511](#)):

```

x: == x      lowest_zero(x)
0: == .....1  .....1
1: == .....11 .....1
2: == .....111 .....1
3: == .....1111 .....1
4: == .....11111 .....1
5: == .....111111 .....1
6: == .....1111111 .....1
7: == .....11111111 .....1
8: == .....111111111 .....1
9: == .....1111111111 .....1
10: == .....11111111111 .....1

```

The lowest set bit in a word can be cleared by

```

1  static inline ulong clear_lowest_one(ulong x)
2  // Return word where the lowest bit set in x is cleared.
3  // Return 0 for input == 0.
4  {
5      return x & (x-1);
6  }

```

The lowest unset bit can be set by

```

1  static inline ulong set_lowest_zero(ulong x)
2  // Return word where the lowest unset bit in x is set.
3  // Return ~0 for input == ~0.
4  {
5      return x | (x+1);
6  }

```

1.3.2 Computing the index of the lowest one

We compute the index (position) of the lowest bit with an assembler instruction if available [FXT: [bits/bitasm-amd64.h](#)]:

```

1  static inline ulong asm_bsf(ulong x)
2  // Bit Scan Forward
3  {
4      asm ("bsfq %0, %0" : "=r" (x) : "0" (x));
5      return x;
6  }

```

Without the assembler instruction an algorithm that involves $O(\log_2 \text{BITS_PER_LONG})$ operations can be used. The function can be implemented as follows (suggested by Nathan Bullock [priv. comm.], 64-bit version) [FXT: [bits/bitlow.h](#)]:

```

1  static inline ulong lowest_one_idx(ulong x)
2  // Return index of lowest bit set.
3  // Examples:
4  //   ***1 --> 0
5  //   **10 --> 1
6  //   *100 --> 2
7  // Return 0 (also) if no bit is set.
8  {
9      ulong r = 0;
10     x &= -x; // isolate lowest bit
11     if ( x & 0xffffffff00000000UL ) r += 32;
12     if ( x & 0xffff0000ffff0000UL ) r += 16;
13     if ( x & 0xff00ff00ff00ff00UL ) r += 8;
14     if ( x & 0xf0f0f0f0f0f0f0f0UL ) r += 4;
15     if ( x & 0xccccccccccccccUL ) r += 2;
16     if ( x & 0xaaaaaaaaaaaaaaaaUL ) r += 1;
17     return r;
18 }

```

The function returns zero for two inputs, one and zero. If a special value for the input zero is needed, a statement as the following should be added as the first line of the function:

```
if ( 1>=x ) return x-1; // 0 if 1, ~0 if 0
```

The following function returns the parity of the index of the lowest set bit in a binary word

```

1  static inline ulong lowest_one_idx_parity(ulong x)
2  {
3      x &= -x; // isolate lowest bit

```



```

4     return 0 != (x & 0xaaaaaaaaaaaaaaaaUL);
5 }

```

The sequence of values for $x = 0, 1, 2, \dots$ is

```
0010001010100010001000101010001010100010101000100010001010100010...
```

This is the complement of the *period-doubling sequence*, entry [A035263](#) in [312]. See section [38.5.1](#) on page [735](#) for the connection to the towers of Hanoi puzzle.

1.3.3 Isolating blocks of zeros or ones at the low end

Isolate the burst of low ones as follows [FXT: [bits/bitlow.h](#)]:

```

1  static inline ulong low_ones(ulong x)
2  // Return word where all the (low end) ones are set.
3  // Example: 01011011 --> 00000011
4  // Return 0 if lowest bit is zero:
5  //      10110110 --> 0
6  {
7      x = ~x;
8      x &= -x;
9      --x;
10     return x;
11 }

```

The isolation of the low zeros is slightly cheaper:

```

1  static inline ulong low_zeros(ulong x)
2  // Return word where all the (low end) zeros are set.
3  // Example: 01011000 --> 00000111
4  // Return 0 if all bits are set.
5  {
6      x &= -x;
7      --x;
8      return x;
9  }

```

The lowest block of ones (which may have zeros to the right of it) can be isolated by

```

1  static inline ulong lowest_block(ulong x)
2  // Isolate lowest block of ones.
3  // e.g.:
4  // x   = *****011100
5  // l   = 00000000100
6  // y   = *****100000
7  // x^y = 00000111100
8  // ret = 00000011100
9  {
10     ulong l = x & -x; // lowest bit
11     ulong y = x + 1;
12     x ^= y;
13     return x & (x>>1);
14 }

```

1.3.4 Creating a transition at the lowest one

Use the following routines to set a rising or falling edge at the position of the lowest set bit [FXT: [bits/bitlow-edge.h](#)]:

```

1  static inline ulong lowest_one_10edge(ulong x)
2  // Return word where all bits from (including) the
3  // lowest set bit to most significant bit are set.
4  // Return 0 if no bit is set.
5  // Example: 00110100 --> 11111100
6  {
7      return ( x | -x );
8  }

1  static inline ulong lowest_one_01edge(ulong x)
2  // Return word where all bits from (including) the
3  // lowest set bit to the least significant are set.
4  // Return 0 if no bit is set.
5  // Example: 00110100 --> 00000111

```

```

6  {
7      if ( 0==x ) return 0;
8      return  x^(x-1);
9  }

```

1.3.5 Isolating the lowest run of matching bits

Let $x = *0W$ and $y = *1W$, the following function computes W :

```

1  static inline long low_match(ulong x, ulong y)
2  {
3      x ^= y;    // bit-wise difference
4      x &= -x;   // lowest bit that differs in both words
5      x -= 1;   // mask that covers equal bits at low end
6      x &= y;   // isolate matching bits
7      return x;
8  }

```

1.4 Extraction of ones, zeros, or blocks near transitions

We give functions for the creation or extraction of bit-blocks and the isolation of values near transitions. A transition is a place where adjacent bits have different values. A block is a group of adjacent bits of the same value.

1.4.1 Creating blocks of ones

The following functions are given in [FXT: [bits/bitblock.h](#)].

```

1  static inline ulong bit_block(ulong p, ulong n)
2  // Return word with length-n bit block starting at bit p set.
3  // Both p and n are effectively taken modulo BITS_PER_LONG.
4  {
5      ulong x = (1UL<<n) - 1;
6      return  x << p;
7  }

```

A version with indices wrapping around is

```

1  static inline ulong cyclic_bit_block(ulong p, ulong n)
2  // Return word with length-n bit block starting at bit p set.
3  // The result is possibly wrapped around the word boundary.
4  // Both p and n are effectively taken modulo BITS_PER_LONG.
5  {
6      ulong x = (1UL<<n) - 1;
7      return  (x<<p) | (x>>(BITS_PER_LONG-p));
8  }

```

1.4.2 Finding isolated ones or zeros

The following functions are given in [FXT: [bits/bit-isolate.h](#)]:

```

1  static inline ulong single_ones(ulong x)
2  // Return word with only the isolated ones of x set.
3  {
4      return  x & ~( (x<<1) | (x>>1) );
5  }

```

We can assume a word is embedded in zeros or ignore the bits outside the word:

```

1  static inline ulong single_zeros_xi(ulong x)
2  // Return word with only the isolated zeros of x set.
3  {
4      return  single_ones( ~x ); // ignore outside values
5  }

1  static inline ulong single_zeros(ulong x)
2  // Return word with only the isolated zeros of x set.
3  {
4      return  ~x & ( (x<<1) & (x>>1) ); // assume outside values == 0
5  }

```

```

1 static inline ulong single_values(ulong x)
2 // Return word where only the isolated ones and zeros of x are set.
3 {
4     return (x ^ (x<<1)) & (x ^ (x>>1)); // assume outside values == 0
5 }

1 static inline ulong single_values_xi(ulong x)
2 // Return word where only the isolated ones and zeros of x are set.
3 {
4     return single_ones(x) | single_zeros_xi(x); // ignore outside values
5 }

```

1.4.3 Isolating single ones or zeros at the word boundary

```

1 static inline ulong border_ones(ulong x)
2 // Return word where only those ones of x are set that lie next to a zero.
3 {
4     return x & ~( (x<<1) & (x>>1) );
5 }

1 static inline ulong border_values(ulong x)
2 // Return word where those bits of x are set that lie on a transition.
3 {
4     return (x ^ (x<<1)) | (x ^ (x>>1));
5 }

```

1.4.4 Isolating transitions

```

1 static inline ulong high_border_ones(ulong x)
2 // Return word where only those ones of x are set
3 // that lie right to (i.e. in the next lower bin of) a zero.
4 {
5     return x & ( x ^ (x>>1) );
6 }

1 static inline ulong low_border_ones(ulong x)
2 // Return word where only those ones of x are set
3 // that lie left to (i.e. in the next higher bin of) a zero.
4 {
5     return x & ( x ^ (x<<1) );
6 }

```

1.4.5 Isolating ones or zeros at block boundaries

```

1 static inline ulong block_border_ones(ulong x)
2 // Return word where only those ones of x are set
3 // that are at the border of a block of at least 2 bits.
4 {
5     return x & ( (x<<1) ^ (x>>1) );
6 }

1 static inline ulong low_block_border_ones(ulong x)
2 // Return word where only those bits of x are set
3 // that are at left of a border of a block of at least 2 bits.
4 {
5     ulong t = x & ( (x<<1) ^ (x>>1) ); // block_border_ones()
6     return t & (x>>1);
7 }

1 static inline ulong high_block_border_ones(ulong x)
2 // Return word where only those bits of x are set
3 // that are at right of a border of a block of at least 2 bits.
4 {
5     ulong t = x & ( (x<<1) ^ (x>>1) ); // block_border_ones()
6     return t & (x<<1);
7 }

1 static inline ulong block_ones(ulong x)
2 // Return word where only those bits of x are set
3 // that are part of a block of at least 2 bits.
4 {
5     return x & ( (x<<1) | (x>>1) );
6 }

```

1.5 Computing the index of a single set bit

In the function `lowest_one_idx()` given in section 1.3.2 on page 9 we first isolated the lowest one of a word `x` by first setting `x&=-x`. At this point, `x` contains just one set bit (or `x==0`). The following lines in the routine compute the index of the only bit set. This section gives some alternative techniques to compute the index of the one in a single-bit word.

1.5.1 Cohen's trick

modulus	m=11								
k=	0	1	2	3	4	5	6	7	
mt[k]=	0	0	1	8	2	4	9	7	
Lowest bit == 0:	x=1	=	1	x % m=	1	==>	lookup =	0
Lowest bit == 1:	x=1.	=	2	x % m=	2	==>	lookup =	1
Lowest bit == 2:	x=1..	=	4	x % m=	4	==>	lookup =	2
Lowest bit == 3:	x=1..	=	8	x % m=	8	==>	lookup =	3
Lowest bit == 4:	x=	...1....	=	16	x % m=	5	==>	lookup =	4
Lowest bit == 5:	x=	..1.....	=	32	x % m=	10	==>	lookup =	5
Lowest bit == 6:	x=	.1.....	=	64	x % m=	9	==>	lookup =	6
Lowest bit == 7:	x=	1.....	=	128	x % m=	7	==>	lookup =	7

Figure 1.5-A: Determination of the position of a single bit with 8-bit words.

A nice trick is presented in [110]: for N -bit words find a number m such that all powers of 2 are different modulo m . That is, the (multiplicative) order of 2 modulo m must be greater than or equal to N . We use a table `mt[]` of size m that contains the power of 2: `mt[(2**j) mod m] = j` for $j > 0$. To look up the index of a one-bit-word `x` it is reduced modulo m and `mt[x]` is returned.

We demonstrate the method for $N = 8$ where $m = 11$ is the smallest number with the required property. The setup routine for the table is

```

1  const ulong m = 11; // the modulus
2  ulong mt[m+1];
3  static void mt_setup()
4  {
5      mt[0] = 0; // special value for the zero word
6      ulong t = 1;
7      for (ulong i=1; i<m; ++i)
8      {
9          mt[t] = i-1;
10         t *= 2;
11         if ( t>=m ) t -= m; // modular reduction
12     }
13 }
```

The entry in `mt[0]` will be accessed when the input is the zero word. We can use any value to be returned for input zero. Here we simply use zero to always have the same return value as with `lowest_one_idx()`. The index can be computed by

```

1  static inline ulong m_lowest_one_idx(ulong x)
2  {
3      x &= -x; // isolate lowest bit
4      x %= m; // power of 2 modulo m
5      return mt[x]; // lookup
6  }
```

The code is given in the program [FXT: [bits/modular-lookup-demo.cc](#)], the output with $N = 8$ (edited for size) is shown in figure 1.5-A. The following moduli $m(N)$ can be used for N -bit words:

N:	4	8	16	32	64	128	256	512	1024
m:	5	11	19	37	67	131	269	523	1061

The modulus $m(N)$ is the smallest prime greater than N such that 2 is a primitive root modulo $m(N)$.

```

db=...1.111 (De Bruijn sequence)
  k = 0 1 2 3 4 5 6 7
dbt[k] = 0 1 2 4 7 3 6 5

Lowest bit == 0:  x = .....1  db * x = ...1.111  shifted = ..... == 0 ==> lookup = 0
Lowest bit == 1:  x = .....1.  db * x = ..1.111.  shifted = .....1 == 1 ==> lookup = 1
Lowest bit == 2:  x = ....1..  db * x = .1.111..  shifted = .....1. == 2 ==> lookup = 2
Lowest bit == 3:  x = ...1...  db * x = 1.111...  shifted = ....1.1 == 5 ==> lookup = 3
Lowest bit == 4:  x = ..1....  db * x = .111....  shifted = .....11 == 3 ==> lookup = 4
Lowest bit == 5:  x = .1.....  db * x = 111.....  shifted = ....111 == 7 ==> lookup = 5
Lowest bit == 6:  x = .1.....  db * x = 11.....  shifted = .....11. == 6 ==> lookup = 6
Lowest bit == 7:  x = 1.....  db * x = 1.....  shifted = .....1.. == 4 ==> lookup = 7

```

Figure 1.5-B: Computing the position of the single set bit in 8-bit words with a De Bruijn sequence.

1.5.2 Using De Bruijn sequences

The following method (given in [228]) is even more elegant. It uses binary De Bruijn sequences of size N . A binary De Bruijn sequence of length 2^N contains all binary words of length N , see section 4.1.1 on page 864. These are the sequences for 32 and 64 bit, as binary words:

```

#if BITS_PER_LONG == 32
const ulong db = 0x4653ADFUL;
// == 00000100011001010011101011011111
const ulong s = 32-5;
#else
const ulong db = 0x218A392CD3D5DBFUL;
// == 0000010000110001010001110010010110011101001110101011101101111111
const ulong s = 64-6;
#endif

```

Let w_i be the i -th sub-word from the left (high end). We create a table such that the entry with index w_i points to i :

```

1  ulong dbt[BITS_PER_LONG];
2  static void dbt_setup()
3  {
4      for (ulong i=0; i<BITS_PER_LONG; ++i)  dbt[ (db<<i)>>s ] = i;
5  }

```

The computation of the index involves a multiplication and a table lookup:

```

1  static inline ulong db_lowest_one_idx(ulong x)
2  {
3      x &= -x; // isolate lowest bit
4      x *= db; // multiplication by a power of 2 is a shift
5      x >>= s; // use log_2(BITS_PER_LONG) highest bits
6      return dbt[x]; // lookup
7  }

```

The used sequences must start with at least $\log_2(N) - 1$ zeros because in the line `x *= db` the word `x` is shifted (not rotated). The code is given in the demo [FXT: bits/debruijn-lookup-demo.cc], the output with $N = 8$ (edited for size, dots denote zeros) is shown in figure 1.5-B.

1.5.3 Using floating-point numbers

Floating-point numbers are normalized so that the highest bit in the mantissa is set. Therefore if we convert an integer into a float, the position of the *highest* set bit can be read off the exponent. By isolating the lowest bit before that operation, the index can be found with the same trick. However, the conversion between integers and floats is usually slow. Further, the technique is highly machine dependent.

1.6 Operations on high bits or blocks of a word

For functions operating on the highest bit there is no method as trivial as shown for the lower end of the word. With a bit-reverse CPU-instruction available life would be significantly easier. However, almost no CPU seems to have it.


```

1  static inline ulong highest_one(ulong x)
2  // Return word where only the highest bit in x is set.
3  // Return 0 if no bit is set.
4  {
5  #if defined BITS_USE_ASM
6      if ( 0==x ) return 0;
7      x = asm_bsr(x);
8      return 1UL<<x;
9  #else
10     x = highest_one_01edge(x);
11     return x ^ (x>>1);
12 #endif // BITS_USE_ASM
13 }

```

To determine the index of the highest set bit, use

```

1  static inline ulong highest_one_idx(ulong x)
2  // Return index of highest bit set.
3  // Return 0 if no bit is set.
4  {
5  #if defined BITS_USE_ASM
6      return asm_bsr(x);
7  #else // BITS_USE_ASM
8
9      if ( 0==x ) return 0;
10
11     ulong r = 0;
12 #if BITS_PER_LONG >= 64
13     if ( x & 0xfffffff00000000UL ) { x >>= 32; r += 32; }
14 #endif
15     if ( x & 0xffff0000UL ) { x >>= 16; r += 16; }
16     if ( x & 0x0000fff00UL ) { x >>= 8; r += 8; }
17     if ( x & 0x000000f0UL ) { x >>= 4; r += 4; }
18     if ( x & 0x0000000cUL ) { x >>= 2; r += 2; }
19     if ( x & 0x00000002UL ) { r += 1; }
20     return r;
21 #endif // BITS_USE_ASM
22 }

```

The branches in the non-assembler part of the routine can be avoided by a technique given in [215, rel.96, sect.7.1.3] (version for 64-bit words):

```

1  static inline ulong highest_one_idx(ulong x)
2  {
3  #define MU0 0x5555555555555555UL // MU0 == ((-1UL)/3UL) == ...01010101_2
4  #define MU1 0x3333333333333333UL // MU1 == ((-1UL)/5UL) == ...00110011_2
5  #define MU2 0x0f0f0f0f0f0f0f0fUL // MU2 == ((-1UL)/17UL) == ...00001111_2
6  #define MU3 0x00ff00ff00ff00ffUL // MU3 == ((-1UL)/257UL) == (8 ones)
7  #define MU4 0x0000ffff0000ffffUL // MU4 == ((-1UL)/65537UL) == (16 ones)
8  #define MU5 0x00000000ffffffffffUL // MU5 == ((-1UL)/4294967297UL) == (32 ones)
9      ulong r = ld_neq(x, x & MU0)
10         + (ld_neq(x, x & MU1) << 1)
11         + (ld_neq(x, x & MU2) << 2)
12         + (ld_neq(x, x & MU3) << 3)
13         + (ld_neq(x, x & MU4) << 4)
14         + (ld_neq(x, x & MU5) << 5);
15     return r;
16 }

```

The auxiliary function `ld_neq()` is given in [FXT: [bits/bitldeq.h](#)]:

```

1  static inline bool ld_neq(ulong x, ulong y)
2  // Return whether floor(log2(x))!=floor(log2(y))
3  { return ( (x^y) > (x&y) ); }

```

The following version for 64-bit words provided by Sebastiano Vigna [priv. comm.] is an implementation of Brodal's algorithm [215, alg.B, sect.7.1.3]:

```

1  static inline ulong highest_one_idx(ulong x)
2  {
3      if ( x == 0 ) return 0;
4      ulong r = 0;
5      if ( x & 0xfffffff00000000UL ) { x >>= 32; r += 32; }
6      if ( x & 0xffff0000UL ) { x >>= 16; r += 16; }
7      x |= (x << 16);
8      x |= (x << 32);
9      const ulong y = x & 0xff00f0f0ccccaaaaUL;

```

```

10     const ulong z = 0x8000800080008000UL;
11     ulong t = z & ( y | (( y | z ) - ( x ^ y )));
12     t |= (t << 15);
13     t |= (t << 30);
14     t |= (t << 60);
15     return r + ( t >> 60 );
16 }

```

1.6.2 Isolating the highest block of ones or zeros

Isolate the left block of zeros with the function

```

1  static inline ulong high_zeros(ulong x)
2  // Return word where all the (high end) zeros are set.
3  // e.g.: 00011001 --> 11100000
4  // Returns 0 if highest bit is set:
5  //      11011001 --> 00000000
6  {
7      x |= x>>1;
8      x |= x>>2;
9      x |= x>>4;
10     x |= x>>8;
11     x |= x>>16;
12     #if BITS_PER_LONG >= 64
13     x |= x>>32;
14     #endif
15     return ~x;
16 }

```

The left block of ones can be isolated using arithmetical right shifts:

```

1  static inline ulong high_ones(ulong x)
2  // Return word where all the (high end) ones are set.
3  // e.g. 11001011 --> 11000000
4  // Returns 0 if highest bit is zero:
5  //      01110110 --> 00000000
6  {
7      long y = (long)x;
8      y &= y>>1;
9      y &= y>>2;
10     y &= y>>4;
11     y &= y>>8;
12     y &= y>>16;
13     #if BITS_PER_LONG >= 64
14     y &= y>>32;
15     #endif
16     return (ulong)y;
17 }

```

If arithmetical shifts are more expensive than unsigned shifts, use

```
1  static inline ulong high_ones(ulong x) { return high_zeros( ~x ); }
```

A demonstration of selected functions operating on the highest or lowest bit (or block) of binary words is given in [FXT: bits/bithilo-demo.cc]. Part of its output is shown in figure 1.6-A.

1.7 Functions related to the base-2 logarithm

The following functions are given in [FXT: bits/bit2pow.h]. A function that returns $\lfloor \log_2(x) \rfloor$ can be implemented using the obvious algorithm:

```

1  static inline ulong ld(ulong x)
2  // Return floor(log2(x)),
3  // i.e. return k so that 2^k <= x < 2^(k+1)
4  // If x==0, then 0 is returned (!)
5  {
6      ulong k = 0;
7      while ( x>>=1 ) { ++k; }
8      return k;
9  }

```

The result is the same as returned by `highest_one_idx()`:


```
1 static inline ulong ld(ulong x) { return highest_one_idx(x); }
```

The bit-wise algorithm can be faster if the average result is known to be small.

Use the function `one_bit_q()` to determine whether its argument is a power of 2:

```
1 static inline bool one_bit_q(ulong x)
2 // Return whether x \in {1,2,4,8,16,...}
3 {
4     ulong m = x-1;
5     return ((x^m)>>1) == m;
6 }
```

The following function does the same except that it returns `true` also for the zero argument:

```
1 static inline bool is_pow_of_2(ulong x)
2 // Return whether x == 0(!) or x == 2**k
3 { return !(x & (x-1)); }
```

With FFTs where the length of the transform is often restricted to power of 2 the following functions are useful:

```
1 static inline ulong next_pow_of_2(ulong x)
2 // Return x if x=2**k
3 // else return 2**ceil(log_2(x))
4 // Exception: returns 0 for x==0
5 {
6     if ( is_pow_of_2(x) ) return x;
7     x |= x >> 1;
8     x |= x >> 2;
9     x |= x >> 4;
10    x |= x >> 8;
11    x |= x >> 16;
12 #if BITS_PER_LONG == 64
13    x |= x >> 32;
14 #endif
15    return x + 1;
16 }

1 static inline ulong next_exp_of_2(ulong x)
2 // Return k if x=2**k else return k+1.
3 // Exception: returns 0 for x==0.
4 {
5     if ( x <= 1 ) return 0;
6     return ld(x-1) + 1;
7 }
```

The following version should be faster if inline assembler is used for `ld()`:

```
1 static inline ulong next_pow_of_2(ulong x)
2 {
3     if ( is_pow_of_2(x) ) return x;
4     ulong n = 1UL<<ld(x); // n<x
5     return n<<1;
6 }
```

The following routine for comparison of base-2 logarithms without actually computing them is suggested by [215, rel.58, sect.7.1.3] [FXT: [bits/bitldeq.h](#)]:

```
1 static inline bool ld_eq(ulong x, ulong y)
2 // Return whether floor(log2(x))==floor(log2(y))
3 { return ( (x^y) <= (x&y) ); }
```

1.8 Counting the bits and blocks of a word

The following functions count the ones in a binary word. They need $O(\log_2(\text{BITS_PER_LONG}))$ operations. We give mostly the 64-bit versions [FXT: [bits/bitcount.h](#)]:

```
1 static inline ulong bit_count(ulong x)
2 // Return number of bits set
3 {
4     x = (0x5555555555555555UL & x) + (0x5555555555555555UL & (x>> 1)); // 0-2 in 2 bits
5     x = (0x3333333333333333UL & x) + (0x3333333333333333UL & (x>> 2)); // 0-4 in 4 bits
6     x = (0x0f0f0f0f0f0f0f0fUL & x) + (0x0f0f0f0f0f0f0f0fUL & (x>> 4)); // 0-8 in 8 bits
7     x = (0x00ff00ff00ff00ffUL & x) + (0x00ff00ff00ff00ffUL & (x>> 8)); // 0-16 in 16 bits
```

```

8     x = (0x0000ffff0000ffffUL & x) + (0x0000ffff0000ffffUL & (x>>16)); // 0-32 in 32 bits
9     x = (0x00000000ffffffffffUL & x) + (0x00000000ffffffffffUL & (x>>32)); // 0-64 in 64 bits
10    return x;
11 }

```

The underlying idea is to do a search via bit masks. The code can be improved to either

```

1     x = ((x>>1) & 0x5555555555555555UL) + (x & 0x5555555555555555UL); // 0-2 in 2 bits
2     x = ((x>>2) & 0x3333333333333333UL) + (x & 0x3333333333333333UL); // 0-4 in 4 bits
3     x = ((x>>4) + x) & 0x0f0f0f0f0f0f0f0fUL; // 0-8 in 8 bits
4     x += x>> 8; // 0-16 in 8 bits
5     x += x>>16; // 0-32 in 8 bits
6     x += x>>32; // 0-64 in 8 bits
7     return x & 0xff;

```

or (taken from [10])

```

1     x -= (x>>1) & 0x5555555555555555UL; // 0-2 in 2 bits
2     x = ((x>>2) & 0x3333333333333333UL) + (x & 0x3333333333333333UL); // 0-4 in 4 bits
3     x = ((x>>4) + x) & 0x0f0f0f0f0f0f0f0fUL; // 0-8 in 8 bits
4     x *= 0x0101010101010101UL;
5     return x>>56;

```

Which of the latter two versions is faster mainly depends on the speed of integer multiplication.

The following code for 32-bit words (given by Johan Rönblom [priv. comm.]) may be advantageous if loading constants is expensive. Note some constants are in octal notation:

```

1 static inline uint CountBits32(uint a)
2 {
3     uint mask = 011111111111UL;
4     a = (a - ((a&~mask)>>1)) - ((a>>2)&mask);
5     a += a>>3;
6     a = (a & 070707) + ((a>>18) & 070707);
7     a *= 010101;
8     return ((a>>12) & 0x3f);
9 }

```

If the table holds the bit-counts of the numbers 0 .. 255, then the bits can be counted as follows:

```

1 unsigned long bit_count(unsigned long x)
2 {
3     unsigned char ct = 0;
4     ct += tab[ x & 0xff ]; x >>= 8;
5     ct += tab[ x & 0xff ]; x >>= 8;
6     [--snip--] /* BYTES_PER_LONG times */
7     ct += tab[ x & 0xff ];
8     return ct;
9 }

```

However, while table driven methods tend to excel in synthetic benchmarks, they can be very slow if they cause cache misses.

We give a method to count the bits of a word of a special form:

```

1 static inline unsigned long bit_count_01(unsigned long x)
2 // Return number of bits in a word
3 // for words of the special form 00...0001...11
4 {
5     unsigned long ct = 0;
6     unsigned long a;
7 #if BITS_PER_LONG == 64
8     a = (x & (1UL<<32)) >> (32-5); // test bit 32
9     x >>= a; ct += a;
10 #endif
11    a = (x & (1UL<<16)) >> (16-4); // test bit 16
12    x >>= a; ct += a;
13
14    a = (x & (1UL<<8)) >> (8-3); // test bit 8
15    x >>= a; ct += a;
16
17    a = (x & (1UL<<4)) >> (4-2); // test bit 4
18    x >>= a; ct += a;
19
20    a = (x & (1UL<<2)) >> (2-1); // test bit 2
21    x >>= a; ct += a;
22
23    a = (x & (1UL<<1)) >> (1-0); // test bit 1

```

```

24     x >>= a;  ct += a;
25
26     ct += x & 1; // test bit 0
27
28     return ct;
29 }

```

All branches are avoided, thereby the code may be useful on a planet with pink air, for further details see [301].

1.8.1 Sparse counting

If the (average input) word is known to have only a few bits set, the following sparse count variant can be advantageous:

```

1  static inline ulong bit_count_sparse(ulong x)
2  // Return number of bits set.
3  {
4      ulong n = 0;
5      while ( x ) { ++n; x &= (x-1); }
6      return n;
7  }

```

The loop will execute once for each set bit. Partial unrolling of the loop should be an improvement for most cases:

```

1      ulong n = 0;
2      do
3      {
4          n += (x!=0); x &= (x-1);
5          n += (x!=0); x &= (x-1);
6          n += (x!=0); x &= (x-1);
7          n += (x!=0); x &= (x-1);
8      }
9      while ( x );
10     return n;

```

If the number of bits is close to the maximum, use the given routine with the complement:

```

1  static inline ulong bit_count_dense(ulong x)
2  // Return number of bits set.
3  // The loop (of bit_count_sparse()) will execute once for
4  // each unset bit (i.e. zero) of x.
5  {
6      return BITS_PER_LONG - bit_count_sparse( ~x );
7  }

```

If the number of ones is guaranteed to be less than 16, then the following routine (suggested by Gunther Piez [priv. comm.]) can be used:

```

1  static inline ulong bit_count_15(ulong x)
2  // Return number of set bits, must have at most 15 set bits.
3  {
4      x -= (x>>1) & 0x5555555555555555UL; // 0-2 in 2 bits
5      x = ((x>>2) & 0x3333333333333333UL) + (x & 0x3333333333333333UL); // 0-4 in 4 bits
6      x *= 0x1111111111111111UL;
7      return x>>60;
8  }

```

A routine for words with no more than 3 set bits is

```

1  static inline ulong bit_count_3(ulong x)
2  {
3      x -= (x>>1) & 0x5555555555555555UL; // 0-2 in 2 bits
4      x *= 0x5555555555555555UL;
5      return x>>62;
6  }

```

1.8.2 Counting blocks

Compute the number of bit-blocks in a binary word with the following function:

```

1  static inline ulong bit_block_count(ulong x)
2  // Return number of bit blocks.
3  // E.g.:
4  // ..1..11111..111. -> 3

```

```

5 // ...1..11111...111 -> 3
6 // .....1.....1.1.. -> 3
7 // .....111.1111 -> 2
8 {
9     return (x & 1) + bit_count( (x^(x>>1)) ) / 2;
10 }

```

Similarly, the number of blocks with two or more bits can be counted via:

```

1 static inline ulong bit_block_ge2_count(ulong x)
2 // Return number of bit blocks with at least 2 bits.
3 // E.g.:
4 // ..1..11111...111. -> 2
5 // ...1..11111...111 -> 2
6 // .....1.....1.1.. -> 0
7 // .....111.1111 -> 2
8 {
9     return bit_block_count( x & ( (x<<1) & (x>>1) ) );
10 }

```

1.8.3 GCC built-in functions ‡

Newer versions of the C compiler of the GNU Compiler Collection (GCC [146], starting with version 3.4) include a function `__builtin_popcountl(ulong)` that counts the bits of an unsigned long integer. The following list is taken from [147]:

```

int __builtin_ffs (unsigned int x)
    Returns one plus the index of the least significant 1-bit of x,
    or if x is zero, returns zero.

int __builtin_clz (unsigned int x)
    Returns the number of leading 0-bits in x, starting at the
    most significant bit position. If x is 0, the result is undefined.

int __builtin_ctz (unsigned int x)
    Returns the number of trailing 0-bits in x, starting at the
    least significant bit position. If x is 0, the result is undefined.

int __builtin_popcount (unsigned int x)
    Returns the number of 1-bits in x.

int __builtin_parity (unsigned int x)
    Returns the parity of x, i.e. the number of 1-bits in x modulo 2.

```

The names of the corresponding versions for arguments of type unsigned long are obtained by adding ‘l’ (ell) to the names, for the type unsigned long long append ‘ll’. Two more useful built-ins are:

```

void __builtin_prefetch (const void *addr, ...)
    Prefetch memory location addr

long __builtin_expect (long exp, long c)
    Function to provide the compiler with branch prediction information.

```

1.8.4 Counting the bits of many words ‡

x[0]=11111111	a0=11111111	a1=.....	a2=.....	a3=.....	a4=.....
x[1]=11111111	a0=.....	a1=11111111	a2=.....	a3=.....	a4=.....
x[2]=11111111	a0=11111111	a1=11111111	a2=.....	a3=.....	a4=.....
x[3]=11111111	a0=.....	a1=.....	a2=11111111	a3=.....	a4=.....
x[4]=11111111	a0=11111111	a1=.....	a2=11111111	a3=.....	a4=.....
x[5]=11111111	a0=.....	a1=11111111	a2=11111111	a3=.....	a4=.....
x[6]=11111111	a0=11111111	a1=11111111	a2=11111111	a3=.....	a4=.....
x[7]=11111111	a0=.....	a1=.....	a2=.....	a3=11111111	a4=.....
x[8]=11111111	a0=11111111	a1=.....	a2=.....	a3=11111111	a4=.....
x[9]=11111111	a0=.....	a1=11111111	a2=.....	a3=11111111	a4=.....
x[10]=11111111	a0=11111111	a1=11111111	a2=.....	a3=11111111	a4=.....
x[11]=11111111	a0=.....	a1=.....	a2=11111111	a3=11111111	a4=.....
x[12]=11111111	a0=11111111	a1=.....	a2=11111111	a3=11111111	a4=.....
x[13]=11111111	a0=.....	a1=11111111	a2=11111111	a3=11111111	a4=.....
x[14]=11111111	a0=11111111	a1=11111111	a2=11111111	a3=11111111	a4=.....
x[15]=11111111	a0=.....	a1=.....	a2=.....	a3=.....	a4=11111111
x[16]=11111111	a0=11111111	a1=.....	a2=.....	a3=.....	a4=11111111

Figure 1.8-A: Counting the bits of an array (where all bits are set) via vertical addition.

For counting the bits in a long array the technique of *vertical addition* can be useful. For ordinary addition the following relation holds:

$$a + b == (a \oplus b) + ((a \& b) \ll 1)$$

The carry term ($a \& b$) is propagated to the left. We now replace this ‘horizontal’ propagation by a ‘vertical’ one, that is, propagation into another word. An implementation of this idea is [FXT: [bits/bitcount-v-demo.cc](#)]:

```

1  ulong
2  bit_count_leq31(const ulong *x, ulong n)
3  // Return sum(j=0, n-1, bit_count(x[j]) )
4  // Must have n<=31
5  {
6      ulong a0=0, a1=0, a2=0, a3=0, a4=0;
7      //      1,    3,    7,   15,   31,   <--- max n
8      for (ulong k=0; k<n; ++k)
9          {
10             ulong cy = x[k];
11             { ulong t = a0 & cy; a0 ^= cy; cy = t; }
12             { ulong t = a1 & cy; a1 ^= cy; cy = t; }
13             { ulong t = a2 & cy; a2 ^= cy; cy = t; }
14             { ulong t = a3 & cy; a3 ^= cy; cy = t; }
15             { a4 ^= cy; }
16             // [ PRINT x[k], a0, a1, a2, a3, a4 ]
17         }
18
19         ulong b = bit_count(a0);
20         b += (bit_count(a1)<<1);
21         b += (bit_count(a2)<<2);
22         b += (bit_count(a3)<<3);
23         b += (bit_count(a4)<<4);
24         return b;
25     }

```

Figure 1.8-A shows the intermediate values with the computation of a length-17 array of all-ones words. After the loop the values of the variables a_0, \dots, a_4 are

```

a4=11111111
a3=.....
a2=.....
a1=.....
a0=11111111

```

The columns, read as binary numbers, tell us that in all positions of all words there were a total of $17 = 10001_2$ bits. The remaining instructions compute the total bit-count.

After some simplifications and loop-unrolling a routine for counting the bits of 15 words can be given as [FXT: [bits/bitcount-v.cc](#)]:

```

1  static inline ulong bit_count_v15(const ulong *x)
2  // Return sum(j=0, 14, bit_count(x[j]) )
3  // Technique is "vertical" addition.
4  {
5      #define VV(A) { ulong t = A & cy; A ^= cy; cy = t; }
6      ulong a1, a2, a3;
7      ulong a0=x[0];
8      { ulong cy = x[ 1]; VV(a0); a1 = cy; }
9      { ulong cy = x[ 2]; VV(a0); a1 ^= cy; }
10     { ulong cy = x[ 3]; VV(a0); VV(a1); a2 = cy; }
11     { ulong cy = x[ 4]; VV(a0); VV(a1); a2 ^= cy; }
12     { ulong cy = x[ 5]; VV(a0); VV(a1); a2 ^= cy; }
13     { ulong cy = x[ 6]; VV(a0); VV(a1); a2 ^= cy; }
14     { ulong cy = x[ 7]; VV(a0); VV(a1); VV(a2); a3 = cy; }
15     { ulong cy = x[ 8]; VV(a0); VV(a1); VV(a2); a3 ^= cy; }
16     { ulong cy = x[ 9]; VV(a0); VV(a1); VV(a2); a3 ^= cy; }
17     { ulong cy = x[10]; VV(a0); VV(a1); VV(a2); a3 ^= cy; }
18     { ulong cy = x[11]; VV(a0); VV(a1); VV(a2); a3 ^= cy; }
19     { ulong cy = x[12]; VV(a0); VV(a1); VV(a2); a3 ^= cy; }
20     { ulong cy = x[13]; VV(a0); VV(a1); VV(a2); a3 ^= cy; }
21     { ulong cy = x[14]; VV(a0); VV(a1); VV(a2); a3 ^= cy; }
22     #undef VV
23
24     ulong b = bit_count(a0);
25     b += (bit_count(a1)<<1);

```

```

26     b += (bit_count(a2)<<2);
27     b += (bit_count(a3)<<3);
28     return b;
29 }

```

Each of the macros `VV` gives three machine instructions, one AND, XOR, and MOVE. The routine for the user is

```

1  ulong
2  bit_count_v(const ulong *x, ulong n)
3  // Return sum(j=0, n-1, bit_count(x[j]) )
4  {
5      ulong b = 0;
6      const ulong *xe = x + n + 1;
7      while ( x+15 < xe ) // process blocks of 15 elements
8      {
9          b += bit_count_v15(x);
10         x += 15;
11     }
12
13     // process remaining elements:
14     const ulong r = (ulong)(xe-x-1);
15     for (ulong k=0; k<r; ++k) b+=bit_count(x[k]);
16
17     return b;
18 }

```

Compared to the obvious method of bit-counting

```

1  ulong bit_count_v2(const ulong *x, ulong n)
2  {
3      ulong b = 0;
4      for (ulong k=0; k<n; ++k) b += bit_count(x[k]);
5      return b;
6  }

```

our routine uses roughly 30 percent less time when an array of 100,000,000 words is processed. There are many possible modifications of the method. If the bit-count routine is rather slow, one may want to avoid the four calls to it after the processing of every 15 words. Instead, the variables `a0`, ..., `a3` could be added (vertically!) to an array of more elements. If that array has n elements, then only with each block of $2^n - 1$ words n calls to the bit-count routine are necessary.

1.9 Words as bitsets

1.9.1 Testing whether subset of given bitset

The following function tests whether a word u , as a bitset, is a subset of the bitset given as the word e [FXT: [bits/bitsubsetq.h](#)]:

```

1  static inline bool is_subset(ulong u, ulong e)
2  // Return whether the set bits of u are a subset of the set bits of e.
3  // That is, as bitsets, test whether u is a subset of e.
4  {
5      return ( (u & e)==u );
6      // return ( (u & ~e)==0 );
7      // return ( (~u | e)!=0 );
8  }

```

If u contains any bits not set in e , then these bits are cleared in the AND-operation and the test for equality will fail. The second version tests whether no element of u lies outside of e , the third is obtained by complementing the equality. A proper subset of e is a subset $\neq e$:

```

1  static inline bool is_proper_subset(ulong u, ulong e)
2  // Return whether u (as bitset) is a proper subset of e.
3  {
4      return ( (u<e) && ((u & e)==u) );
5  }

```

The generated machine code contains a branch:

```

101  xorl   %eax, %eax    # prephitmp.71
102  cmpq   %rsi, %rdi    # e, u

```

```

103   jae    .L6           #, /* branch to end of function */
104   andq   %rdi, %rsi    # u, e
106   xorl   %eax, %eax    # prehitmp.71
107   cmpq   %rdi, %rsi    # u, e
108   sete   %al          #, prehitmp.71

```

Replace the Boolean operator ‘&&’ by the bit-wise operator ‘&’ to obtain branch-free machine code:

```

101   cmpq   %rsi, %rdi    # e, u
102   setb   %al          #, tmp63
103   andq   %rdi, %rsi    # u, e
105   cmpq   %rdi, %rsi    # u, e
106   sete   %dl          #, tmp66
107   andl   %edx, %eax    # tmp66, tmp63
108   movzbl %al, %eax     # tmp63, tmp61

```

1.9.2 Testing whether an element is in a given set

We determine whether a given number is an element of a given set (which must be a subset of the set $\{0, 1, 2, \dots, \text{BITS_PER_LONG}-1\}$). For example, to determine whether x is a prime less than 32, use the function

```

1   ulong m = (1UL<<2) | (1UL<<3) | (1UL<<5) | ... | (1UL<<31); // precomputed
2   static inline ulong is_tiny_prime(ulong x)
3   {
4       return m & (1UL << x);
5   }

```

The same idea can be applied to look up tiny factors [FXT: [bits/tinyfactors.h](#)]:

```

1   static inline bool is_tiny_factor(ulong x, ulong d)
2   // For x,d < BITS_PER_LONG (!)
3   // return whether d divides x (1 and x included as divisors)
4   // no need to check whether d==0
5   //
6   {
7       return ( 0 != ( (tiny_factors_tab[x]>>d) & 1 ) );
8   }

```

The function uses the precomputed array [FXT: [bits/tinyfactors.cc](#)]:

```

1   extern const ulong tiny_factors_tab[] =
2   {
3       0x0UL, // x = 0:          ( bits: ..... )
4       0x2UL, // x = 1:  1          ( bits: .....1 )
5       0x6UL, // x = 2:  1 2        ( bits: .....11 )
6       0xaUL, // x = 3:  1 3        ( bits: ....1.1 )
7       0x16UL, // x = 4:  1 2 4      ( bits: ...1.11 )
8       0x22UL, // x = 5:  1 5        ( bits: ..1...1 )
9       0x4eUL, // x = 6:  1 2 3 6    ( bits: .1.111 )
10      0x82UL, // x = 7:  1 7        ( bits: 1.....1 )
11      0x116UL, // x = 8:  1 2 4 8
12      0x20aUL, // x = 9:  1 3 9
13  [--snip--]
14      0x20000002UL, // x = 29:  1 29
15      0x4000846eUL, // x = 30:  1 2 3 5 6 10 15 30
16      0x80000002UL, // x = 31:  1 31
17  #if ( BITS_PER_LONG > 32 )
18      0x100010116UL, // x = 32:  1 2 4 8 16 32
19      0x20000080aUL, // x = 33:  1 3 11 33
20  [--snip--]
21      0x2000000000000002UL, // x = 61:  1 61
22      0x4000000080000006UL, // x = 62:  1 2 31 62
23      0x80000000020028aUL, // x = 63:  1 3 7 9 21 63
24  #endif // ( BITS_PER_LONG > 32 )
25  };

```

Bit-arrays of arbitrary size are discussed in section 4.6 on page 164.

1.10 Index of the i -th set bit

To determine the index of the i -th set bit, we use a technique similar to the method for counting the bits of a word. Only the 64-bit version is shown [FXT: [bits/ith-one-idx.h](#)]:

```

1  static inline ulong ith_one_idx(ulong x, ulong i)
2  // Return index of the i-th set bit of x where 0 <= i < bit_count(x).
3  {
4      ulong x2 = x - ((x>>1) & 0x5555555555555555UL);    // 0-2 in 2 bits
5      ulong x4 = ((x2>>2) & 0x3333333333333333UL) +
6                (x2 & 0x3333333333333333UL);           // 0-4 in 4 bits
7      ulong x8 = ((x4>>4) + x4) & 0x0f0f0f0f0f0f0f0fUL; // 0-8 in 8 bits
8      ulong ct = (x8 * 0x0101010101010101UL) >> 56;    // bit count
9
10     ++i;
11     if ( ct < i ) return ~0UL; // less than i bits set
12
13     ulong x16 = (0x00ff00ff00ff00ffUL & x8) + (0x00ff00ff00ff00ffUL & (x8>>8)); // 0-16
14     ulong x32 = (0x0000ffff0000ffffUL & x16) + (0x0000ffff0000ffffUL & (x16>>16)); // 0-32
15
16     ulong w, s = 0;
17
18     w = x32 & 0xffffffffUL;
19     if ( w < i ) { s += 32; i -= w; }
20
21     x16 >>= s;
22     w = x16 & 0xffff;
23     if ( w < i ) { s += 16; i -= w; }
24
25     x8 >>= s;
26     w = x8 & 0xff;
27     if ( w < i ) { s += 8; i -= w; }
28
29     x4 >>= s;
30     w = x4 & 0xf;
31     if ( w < i ) { s += 4; i -= w; }
32
33     x2 >>= s;
34     w = x2 & 3;
35     if ( w < i ) { s += 2; i -= w; }
36
37     x >>= s;
38     s += ( (x&1) != i );
39
40     return s;
41 }

```

1.11 Avoiding branches

Branches are expensive operations with many CPUs, especially if the CPU pipeline is very long. A useful trick is to replace

```
if ( (x<0) || (x>m) ) { ... }
```

where x might be a signed integer, by

```
if ( (unsigned)x > m ) { ... }
```

The obvious code to test whether a point (x, y) lies outside a square box of size m is

```
if ( (x<0) || (x>m) || (y<0) || (y>m) ) { ... }
```

If m is a power of 2, it is better to use

```
if ( ( (ulong)x | (ulong)y ) > (unsigned)m ) { ... }
```

The following functions are given in [FXT: [bits/branchless.h](#)]. This function returns $\max(0, x)$. That is, zero is returned for negative input, else the unmodified input:

```

1  static inline long max0(long x)
2  {
3      return x & ~(x >> (BITS_PER_LONG-1));
4  }

```

There is no restriction on the input range. The trick used is that with negative x the arithmetic shift will give a word of all ones which is then negated and the AND-operation clears all bits. Note this function

will only work if the compiler emits an arithmetic right shift, see section 1.1.3 on page 3. The following routine computes $\min(0, x)$:

```
1  static inline long min0(long x)
2  // Return min(0, x), i.e. return zero for positive input
3  {
4      return x & (x >> (BITS_PER_LONG-1));
5  }
```

The following `upos_*()` functions only work for a limited range. The highest bit must not be set as it is used to emulate the carry flag. Branchless computation of the absolute difference $|a - b|$:

```
1  static inline ulong upos_abs_diff(ulong a, ulong b)
2  {
3      long d1 = b - a;
4      long d2 = (d1 & (d1 >> (BITS_PER_LONG-1))) << 1;
5      return d1 - d2; // == (b - d) - (a + d);
6  }
```

The following routine sorts two values:

```
1  static inline void upos_sort2(ulong &a, ulong &b)
2  // Set {a, b} := {min(a, b), max(a,b)}
3  // Both a and b must not have the most significant bit set
4  {
5      long d = b - a;
6      d &= (d >> (BITS_PER_LONG-1));
7      a += d;
8      b -= d;
9  }
```

Johan Rönblom gives [priv. comm.] the following versions for signed integer minimum, maximum, and absolute value, that can be advantageous for CPUs where immediates are expensive:

```
1  #define B1 (BITS_PER_LONG-1) // bits of signed int minus one
2  #define MINI(x,y) (((x) & (((int)((x)-(y))) >> B1)) + ((y) & ~(((int)((x)-(y))) >> B1)))
3  #define MAXI(x,y) (((x) & ~(((int)((x)-(y))) >> B1)) + ((y) & (((int)((x)-(y))) >> B1)))
4  #define ABSI(x) ((x) & ~(((int)(x)) >> B1)) - ((x) & (((int)(x)) >> B1))
```

Your compiler may be smarter than you thought

The machine code generated for

```
x = x & ~(x >> (BITS_PER_LONG-1)); // max0()
```

is

```
35:  48 99                cqto
37:  48 83 c4 08          add    $0x8,%rsp // stack adjustment
3b:  48 f7 d2             not    %rdx
3e:  48 21 d0             and    %rdx,%rax
```

The variable `x` resides in the register `rAX` both at start and end of the function. The compiler uses a special (AMD64) instruction `cqto`. Quoting [13]:

Copies the sign bit in the `rAX` register to all bits of the `rDX` register. The effect of this instruction is to convert a signed word, doubleword, or quadword in the `rAX` register into a signed doubleword, quadword, or double-quadword in the `rDX:rAX` registers. This action helps avoid overflow problems in signed number arithmetic.

Now the equivalent

```
x = ( x < 0 ? 0 : x ); // max0() "simple minded"
```

is compiled to:

```
35:  ba 00 00 00 00      mov    $0x0,%edx
3a:  48 85 c0             test   %rax,%rax
3d:  48 0f 48 c2         cmovs %rdx,%rax // note %edx is %rdx
```

A conditional move (`cmovs`) instruction is used here. That is, the optimized version is (on my machine) actually worse than the straightforward equivalent.

A second example is a function to adjust a given value when it lies outside a given range [FXT: [bits/branchless.h](#)]:

```

1  static inline long clip_range(long x, long mi, long ma)
2  // Code equivalent to (for mi<=ma):
3  //   if ( x<mi ) x = mi;
4  //   else if ( x>ma ) x = ma;
5  {
6      x -= mi;
7      x = clip_range0(x, ma-mi);
8      x += mi;
9      return x;
10 }
```

The auxiliary function used involves one branch:

```

1  static inline long clip_range0(long x, long m)
2  // Code equivalent (for m>0) to:
3  //   if ( x<0 ) x = 0;
4  //   else if ( x>m ) x = m;
5  //   return x;
6  {
7      if ( (ulong)x > (ulong)m ) x = m & ~(x >> (BITS_PER_LONG-1));
8      return x;
9  }
```

The generated machine code is

```

0: 48 89 f8          mov    %rdi,%rax
3: 48 29 f2          sub    %rsi,%rdx
6: 31 c9             xor    %ecx,%ecx
8: 48 29 f0          sub    %rsi,%rax
b: 78 0a            js     17 <_Z2CL11l+0x17> // the branch
d: 48 39 d0          cmp    %rdx,%rax
10: 48 89 d1          mov    %rdx,%rcx
13: 48 0f 4e c8       cmovle %rax,%rcx
17: 48 8d 04 0e       lea   (%rsi,%rcx,1),%rax
```

Now we replace the code by

```

1  static inline long clip_range(long x, long mi, long ma)
2  {
3      x -= mi;
4      if ( x<0 ) x = 0;
5      // else // commented out to make (compiled) function really branchless
6      {
7          ma -= mi;
8          if ( x>ma ) x = ma;
9      }
10     x += mi;
11 }
```

Then the compiler generates branchless code:

```

0: 48 89 f8          mov    %rdi,%rax
3: b9 00 00 00 00    mov    $0x0,%ecx
8: 48 29 f0          sub    %rsi,%rax
b: 48 0f 48 c1       cmovs %rcx,%rax
f: 48 29 f2          sub    %rsi,%rdx
12: 48 39 d0          cmp    %rdx,%rax
15: 48 0f 4f c2       cmovg %rdx,%rax
19: 48 01 f0          add    %rsi,%rax
```

Still, with CPUs that do not have a conditional move instruction (or some branchless equivalent of it) the techniques shown in this section can be useful.

1.12 Bit-wise rotation of a word

Neither C nor C++ have a statement for bit-wise rotation of a binary word (which may be considered a missing feature). The operation can be emulated via [FXT: [bits/bitrotate.h](#)]:

```

1  static inline ulong bit_rotate_left(ulong x, ulong r)
2  // Return word rotated r bits to the left
3  // (i.e. toward the most significant bit)
```

```

4  {
5      return (x<<r) | (x>>(BITS_PER_LONG-r));
6  }

```

As already mentioned, GCC emits exactly the CPU instruction that is *meant* here, even with non-constant argument *r*. Explicit use of the corresponding assembler instruction should not do any harm:

```

1  static inline ulong bit_rotate_right(ulong x, ulong r)
2  // Return word rotated r bits to the right
3  // (i.e. toward the least significant bit)
4  {
5  #if defined BITS_USE_ASM    // use x86 asm code
6      return asm_ror(x, r);
7  #else
8      return (x>>r) | (x<<(BITS_PER_LONG-r));
9  #endif
10 }

```

Here we use an assembler instruction when available [FXT: [bits/bitasm-amd64.h](#)]:

```

1  static inline ulong asm_ror(ulong x, ulong r)
2  {
3      asm ("rorq   %%cl, %0" : "=r" (x) : "0" (x), "c" (r));
4      return x;
5  }

```

Rotation using only a part of the word length can be implemented as

```

1  static inline ulong bit_rotate_left(ulong x, ulong r, ulong ldn)
2  // Return ldn-bit word rotated r bits to the left
3  // (i.e. toward the most significant bit)
4  // Must have 0 <= r <= ldn
5  {
6      ulong m = ~0UL >> ( BITS_PER_LONG - ldn );
7      x &= m;
8      x = (x<<r) | (x>>(ldn-r));
9      x &= m;
10     return x;
11 }

```

and

```

1  static inline ulong bit_rotate_right(ulong x, ulong r, ulong ldn)
2  // Return ldn-bit word rotated r bits to the right
3  // (i.e. toward the least significant bit)
4  // Must have 0 <= r <= ldn
5  {
6      ulong m = ~0UL >> ( BITS_PER_LONG - ldn );
7      x &= m;
8      x = (x>>r) | (x<<(ldn-r));
9      x &= m;
10     return x;
11 }

```

Finally, the functions

```

1  static inline ulong bit_rotate_sgn(ulong x, long r, ulong ldn)
2  // Positive r --> shift away from element zero
3  {
4      if ( r > 0 ) return bit_rotate_left(x, (ulong)r, ldn);
5      else        return bit_rotate_right(x, (ulong)-r, ldn);
6  }

```

and (full-word version)

```

1  static inline ulong bit_rotate_sgn(ulong x, long r)
2  // Positive r --> shift away from element zero
3  {
4      if ( r > 0 ) return bit_rotate_left(x, (ulong)r);
5      else        return bit_rotate_right(x, (ulong)-r);
6  }

```

are sometimes convenient.

1.13 Binary necklaces ‡

We give several functions related to cyclic rotations of binary words and a class to generate binary necklaces.

1.13.1 Cyclic matching, minimum, and maximum

The following function determines whether there is a cyclic right shift of its second argument so that it matches the first argument. It is given in [FXT: [bits/bitcyclic-match.h](#)]:

```

1  static inline ulong bit_cyclic_match(ulong x, ulong y)
2  // Return r if x==rotate_right(y, r) else return ~0UL.
3  // In other words: return
4  //   how often the right arg must be rotated right (to match the left)
5  // or, equivalently:
6  //   how often the left arg must be rotated left (to match the right)
7  {
8      ulong r = 0;
9      do
10     {
11         if ( x==y ) return r;
12         y = bit_rotate_right(y, 1);
13     }
14     while ( ++r < BITS_PER_LONG );
15     return ~0UL;
16 }

```

The functions shown work on the full length of the words, equivalents for the sub-word of the lowest `ldn` bits are given in the respective files. Just one example:

```

1  static inline ulong bit_cyclic_match(ulong x, ulong y, ulong ldn)
2  // Return r if x==rotate_right(y, r, ldn) else return ~0UL
3  // (using ldn-bit words)
4  {
5      ulong r = 0;
6      do
7      {
8          if ( x==y ) return r;
9          y = bit_rotate_right(y, 1, ldn);
10     }
11     while ( ++r < ldn );
12     return ~0UL;
13 }

```

The minimum among all cyclic shifts of a word can be computed via the following function given in [FXT: [bits/bitcyclic-minmax.h](#)]:

```

1  static inline ulong bit_cyclic_min(ulong x)
2  // Return minimum of all rotations of x
3  {
4      ulong r = 1;
5      ulong m = x;
6      do
7      {
8          x = bit_rotate_right(x, 1);
9          if ( x<m ) m = x;
10     }
11     while ( ++r < BITS_PER_LONG );
12     return m;
13 }

```


in [bits/bit-necklace.h](#)]:

```

1  class bit_necklace
2  {
3  public:
4      ulong a_;    // necklace
5      ulong j_;    // period of the necklace
6      ulong n2_;   // bit representing n: n2==2**(n-1)
7      ulong j2_;   // bit representing j: j2==2**(j-1)
8      ulong n_;    // number of bits in words
9      ulong mm_;   // mask of n ones
10     ulong tfb_;  // for fast factor lookup
11
12 public:
13     bit_necklace(ulong n) { init(n); }
14     ~bit_necklace() { ; }
15
16     void init(ulong n)
17     {
18         if ( 0==n ) n = 1; // avoid hang
19         if ( n>=BITS_PER_LONG ) n = BITS_PER_LONG;
20         n_ = n;
21
22         n2_ = 1UL<<(n-1);
23         mm_ = (~0UL) >> (BITS_PER_LONG-n);
24         tfb_ = tiny_factors_tab[n] >> 1;
25         tfb_ |= n2_; // needed for n==BITS_PER_LONG
26         first();
27     }
28
29     void first()
30     {
31         a_ = 0;
32         j_ = 1;
33         j2_ = 1;
34     }
35
36     ulong data() const { return a_; }
37     ulong period() const { return j_; }

```

The method for computing the successor is

```

1  ulong next()
2  // Create next necklace.
3  // Return the period, zero when current necklace is last.
4  {
5      if ( a_==mm_ ) { first(); return 0; }
6
7      do
8      {
9          // next lines compute index of highest zero, same result as
10         // j_ = highest_zero_idx( a_ ^ (~mm_ ) );
11         // but the direct computation is faster:
12         j_ = n_ - 1;
13         ulong jb = 1UL << j_;
14         while ( 0!=(a_ & jb) ) { --j_; jb>>=1; }
15
16         j2_ = 1UL << j_;
17         ++j_;
18         a_ |= j2_;
19         a_ = bit_copy_periodic(a_, j_, n_);
20     }
21     while ( 0==(tfb_ & j2_) ); // necklaces only
22
23     return j_;
24 }

```

It uses the following function for periodic copying [FXT: [bits/bitperiodic.h](#)]:

```

1  static inline ulong bit_copy_periodic(ulong a, ulong p, ulong ldn)
2  // Return word that consists of the lowest p bits of a repeated
3  // in the lowest ldn bits (higher bits are zero).
4  // E.g.: if p==3, ldn=7 and a=*****xyz (8-bit), the return 0zxyzxyz.
5  // Must have p>0 and ldn>0.
6  {
7      a &= ( ~0UL >> (BITS_PER_LONG-p) );

```

```

8     for (ulong s=p; s<ldn; s<<=1) { a |= (a<<s); }
9     a &= ( ~0UL >> (BITS_PER_LONG-ldn) );
10    return a;
11 }

```

Finally, we can easily detect whether a necklace is a Lyndon word:

```

1     ulong is_lyndon_word() const { return (j2_ & n2_); }
2
3     ulong next_lyn()
4     // Create next Lyndon word.
5     // Return the period (==n), zero when current necklace is last.
6     {
7         if ( a_==mm_ ) { first(); return 0; }
8         do { next(); } while ( !is_lyndon_word() );
9         return n_;
10    }
11 };

```

About 54 million necklaces per second are generated (with $n = 32$), corresponding to a rate of 112 M/s for pre-necklaces [FXT: bits/bit-necklace-demo.cc].

1.13.4 Computing the cyclic distance

A function to compute the cyclic distance between two words [FXT: bits/bitcyclic-dist.h] is:

```

1     static inline ulong bit_cyclic_dist(ulong a, ulong b)
2     // Return minimal bitcount of (t ^ b)
3     // where t runs through the cyclic rotations of a.
4     {
5         ulong d = ~0UL;
6         ulong t = a;
7         do
8         {
9             ulong z = t ^ b;
10            ulong e = bit_count( z );
11            if ( e < d ) d = e;
12            t = bit_rotate_right(t, 1);
13        }
14        while ( t!=a );
15        return d;
16    }

```

If the arguments are cyclic shifts of each other, then zero is returned. A version for partial words is

```

1     static inline ulong bit_cyclic_dist(ulong a, ulong b, ulong ldn)
2     {
3         ulong d = ~0UL;
4         const ulong m = (~0UL>>(BITS_PER_LONG-ldn));
5         b &= m;
6         a &= m;
7         ulong t = a;
8         do
9         {
10            ulong z = t ^ b;
11            ulong e = bit_count( z );
12            if ( e < d ) d = e;
13            t = bit_rotate_right(t, 1, ldn);
14        }
15        while ( t!=a );
16        return d;
17    }

```

1.13.5 Cyclic XOR and its inverse

The functions [FXT: bits/bitcyclic-xor.h]

```

1     static inline ulong bit_cyclic_rxor(ulong x)
2     {
3         return x ^ bit_rotate_right(x, 1);
4     }

```

and

```

1  static inline ulong bit_cyclic_lxor(ulong x)
2  {
3      return x ^ bit_rotate_left(x, 1);
4  }

```

return a word whose number of set bits is even. A word and its complement produce the same result.

The inverse functions need no rotation at all, the inverse of `bit_cyclic_rxor()` is the inverse Gray code (see section 1.16 on page 41):

```

1  static inline ulong bit_cyclic_inv_rxor(ulong x)
2  // Return v so that bit_cyclic_rxor(v) == x.
3  {
4      return inverse_gray_code(x);
5  }

```

The argument `x` must have an even number of bits. If this is the case, the lowest bit of the result is zero. The complement of the returned value is also an inverse of `bit_cyclic_rxor()`.

The inverse of `bit_cyclic_lxor()` is the inverse reversed code (see section 1.16.6 on page 45):

```

1  static inline ulong bit_cyclic_inv_lxor(ulong x)
2  // Return v so that bit_cyclic_lxor(v) == x.
3  {
4      return inverse_rev_gray_code(x);
5  }

```

We do not need to mask out the lowest bit because for valid arguments (that have an even number of bits) the high bits of the result are zero. This function can be used to solve the quadratic equation $v^2 + v = x$ in the finite field $\text{GF}(2^n)$ when normal bases are used, see section 42.6.2 on page 903.

1.14 Reversing the bits of a word

The bits of a binary word can efficiently be reversed by a sequence of steps that reverse the order of certain blocks. For 16-bit words, we need $4 = \log_2(16)$ such steps [FXT: bits/revbin-steps-demo.cc]:

```

[ 0 1 2 3 4 5 6 7 8 9 a b c d e f ]
[ 1 0 3 2 5 4 7 6 9 8 b a d c f e ] <--- pairs swapped
[ 3 2 1 0 7 6 5 4 b a 9 8 f e d c ] <--- groups of 2 swapped
[ 7 6 5 4 3 2 1 0 f e d c b a 9 8 ] <--- groups of 4 swapped
[ f e d c b a 9 8 7 6 5 4 3 2 1 0 ] <--- groups of 8 swapped

```

1.14.1 Swapping adjacent bit blocks

We need a couple of auxiliary functions given in [FXT: bits/bitswap.h]. Pairs of adjacent bits can be swapped via

```

1  static inline ulong bit_swap_1(ulong x)
2  // Return x with neighbor bits swapped.
3  {
4      #if BITS_PER_LONG == 32
5          ulong m = 0x55555555UL;
6      #else
7          #if BITS_PER_LONG == 64
8              ulong m = 0x5555555555555555UL;
9          #endif
10         #endif
11         return ((x & m) << 1) | ((x & (~m)) >> 1);
12     }

```

The 64-bit branch is omitted in the following examples. Adjacent groups of 2 bits are swapped by

```

1  static inline ulong bit_swap_2(ulong x)
2  // Return x with groups of 2 bits swapped.
3  {
4      ulong m = 0x33333333UL;
5      return ((x & m) << 2) | ((x & (~m)) >> 2);
6  }

```

Equivalently,


```

1 static inline ulong bit_swap_4(ulong x)
2 // Return x with groups of 4 bits swapped.
3 {
4     ulong m = 0x0f0f0f0fUL;
5     return ((x & m) << 4) | ((x & (~m)) >> 4);
6 }

```

and

```

1 static inline ulong bit_swap_8(ulong x)
2 // Return x with groups of 8 bits swapped.
3 {
4     ulong m = 0x00ff00ffUL;
5     return ((x & m) << 8) | ((x & (~m)) >> 8);
6 }

```

When swapping half-words (here for 32-bit architectures)

```

1 static inline ulong bit_swap_16(ulong x)
2 // Return x with groups of 16 bits swapped.
3 {
4     ulong m = 0x0000ffffUL;
5     return ((x & m) << 16) | ((x & (m<<16)) >> 16);
6 }

```

we could also use the bit-rotate function from section 1.12 on page 27, or

```
return (x << 16) | (x >> 16);
```

The GCC compiler recognizes that the whole operation is equivalent to a (left or right) word rotation and indeed emits just a single rotate instruction.

1.14.2 Bit-reversing binary words

The following is a function to reverse the bits of a binary word [FXT: [bits/revbin.h](#)]:

```

1 static inline ulong revbin(ulong x)
2 // Return x with reversed bit order.
3 {
4     x = bit_swap_1(x);
5     x = bit_swap_2(x);
6     x = bit_swap_4(x);
7     x = bit_swap_8(x);
8     x = bit_swap_16(x);
9     #if BITS_PER_LONG >= 64
10    x = bit_swap_32(x);
11 #endif
12    return x;
13 }

```

The steps after `bit_swap_4()` correspond to a byte-reverse operation. This operation is just one assembler instruction for many CPUs. The inline assembler with GCC for AMD64 CPUs is given in [FXT: [bits/bitasm-amd64.h](#)]:

```

1 static inline ulong asm_bswap(ulong x)
2 {
3     asm ("bswap %0" : "=r" (x) : "0" (x));
4     return x;
5 }

```

We use it for byte reversal if available:

```

1 static inline ulong bswap(ulong x)
2 // Return word with reversed byte order.
3 {
4     #ifdef BITS_USE_ASM
5         x = asm_bswap(x);
6     #else
7         x = bit_swap_8(x);
8         x = bit_swap_16(x);
9     #if BITS_PER_LONG >= 64
10    x = bit_swap_32(x);
11 #endif
12 #endif // def BITS_USE_ASM
13    return x;

```

```
14 }
```

The function actually used for bit reversal is good for both 32 and 64 bit words:

```
1 static inline ulong revbin(ulong x)
2 {
3     x = bit_swap_1(x);
4     x = bit_swap_2(x);
5     x = bit_swap_4(x);
6     x = bswap(x);
7     return x;
8 }
```

The masks can be generated in the process:

```
1 static inline ulong revbin(ulong x)
2 {
3     ulong s = BITS_PER_LONG >> 1;
4     ulong m = ~0UL >> s;
5     while ( s )
6     {
7         x = ( (x & m) << s ) ^ ( (x & (~m)) >> s );
8         s >>= 1;
9         m ^= (m<<s);
10    }
11    return x;
12 }
```

The above function will not always beat the obvious, bit-wise algorithm:

```
1 static inline ulong revbin(ulong x)
2 {
3     ulong r = 0, ldn = BITS_PER_LONG;
4     while ( ldn-- != 0 )
5     {
6         r <<= 1;
7         r += (x&1);
8         x >>= 1;
9     }
10    return r;
11 }
```

Therefore the function

```
1 static inline ulong revbin(ulong x, ulong ldn)
2 // Return word with the ldn least significant bits
3 // (i.e. bit_0 ... bit_{ldn-1}) of x reversed,
4 // the other bits are set to zero.
5 {
6     return revbin(x) >> (BITS_PER_LONG-ldn);
7 }
```

should only be used if `ldn` is not too small, else be replaced by the trivial algorithm.

We can use table lookups so that, for example, eight bits are reversed at a time using a 256-byte table.

The routine for full words is

```
1 unsigned char revbin_tab[256]; // reversed 8-bit words
2 ulong revbin_t(ulong x)
3 {
4     ulong r = 0;
5     for (ulong k=0; k<BYTES_PER_LONG; ++k)
6     {
7         r <<= 8;
8         r |= revbin_tab[ x & 255 ];
9         x >>= 8;
10    }
11    return r;
12 }
```

The routine can be optimized by unrolling to avoid all branches:

```
1 static inline ulong revbin_t(ulong x)
2 {
3     ulong r      = revbin_tab[ x & 255 ]; x >>= 8;
4     r <<= 8; r |= revbin_tab[ x & 255 ]; x >>= 8;
5     r <<= 8; r |= revbin_tab[ x & 255 ]; x >>= 8;
6     #if BYTES_PER_LONG > 4
```

```

7     r <<= 8; r |= revbin_tab[ x & 255 ]; x >>= 8;
8     r <<= 8; r |= revbin_tab[ x & 255 ]; x >>= 8;
9     r <<= 8; r |= revbin_tab[ x & 255 ]; x >>= 8;
10    r <<= 8; r |= revbin_tab[ x & 255 ]; x >>= 8;
11    #endif
12    r <<= 8; r |= revbin_tab[ x ];
13    return r;
14 }

```

However, reversing the first 2^{30} binary words with this routine takes (on a 64-bit machine) longer than with the routine using the `bit_swap_NN()` calls, see [FXT: [bits/revbin-tab-demo.cc](#)].

1.14.3 Generating the bit-reversed words in order

If the bit-reversed words have to be generated in the (reversed) counting order, there is a significantly cheaper way to do the update [FXT: [bits/revbin-upd.h](#)]:

```

1  static inline ulong revbin_upd(ulong r, ulong h)
2  // Let n=2**ldn and h=n/2.
3  // Then, with r == revbin(x, ldn) at entry, return revbin(x+1, ldn)
4  // Note: routine will hang if called with r the all-ones word
5  {
6      while ( !((r^=h)&h) ) h >>= 1;
7      return r;
8  }

```

Now assume we want to generate the bit-reversed words of all $N = 2^n - 1$ words less than 2^n . The total number of branches with the `while`-loop can be estimated by observing that for half of the updates just one bit changes, two bits change for a quarter, three bits change for one eighth of all updates, and so on. So the loop executes less than $2N$ times:

$$N \left(\frac{1}{2} + \frac{2}{4} + \frac{3}{8} + \frac{4}{16} + \cdots + \frac{\log_2(N)}{N} \right) = N \sum_{j=1}^{\log_2(N)} \frac{j}{2^j} < 2N \quad (1.14-1)$$

For large values of N the following method can be significantly faster if a fast routine is available for the computation of the least significant bit in a word. The underlying observation is that for a fixed word of size n there are just n different patterns of bit-changes with incrementing. We generate a lookup table of the bit-reversed patterns, `utab[]`, an array of `BITS_PER_LONG` elements:

```

1  static inline void make_revbin_upd_tab(ulong ldn)
2  // Initialize lookup table used by revbin_tupd()
3  {
4      utab[0] = 1UL<<(ldn-1);
5      for (ulong k=1; k<ldn; ++k) utab[k] = utab[k-1] | (utab[k-1]>>1);
6  }

```

The change patterns for $n = 5$ start as

pattern	reversed pattern
...11	11...
...111	111..
...1111	1111.
...11111	11111.
...111111	111111.
...1111111	1111111.
...11111111	11111111.
...111111111	111111111.

The pattern with x set bits is used for the update of k to $k + 1$ when the lowest zero of k is at position $x - 1$:

	reversed	used when the lowest zero of k is at index:
utab[0]=	1...	0
utab[1]=	11...	1
utab[2]=	111..	2
utab[3]=	1111.	3
utab[4]=	11111.	4

The update routine can now be implemented as

```

1  static inline ulong revbin_tupd(ulong r, ulong k)
2  // Let r==revbin(k, ldn) then
3  // return revbin(k+1, ldn).
4  // NOTE 1: need to call make_revbin_upd_tab(ldn) before usage
5  //       where ldn=log_2(n)
6  // NOTE 2: different argument structure than revbin_upd()
7  {
8      k = lowest_one_idx(~k); // lowest zero idx
9      r ^= utab[k];
10     return r;
11 }

```

The revbin-update routines are used for the revbin permutation described in section 2.6.

	30 bits	16 bits	8 bits	
Update, bit-wise	1.00	1.00	1.00	revbin_upd()
Update, table	0.99	1.08	1.15	revbin_tupd()
Full, masks	0.74	0.81	0.86	revbin()
Full, 8-bit table	1.77	1.94	2.06	revbin_t()
Full32, 8-bit table	0.83	0.90	0.96	revbin_t_le32()
Full16, 8-bit table	—	0.54	0.58	revbin_t_le16()
Full, generated masks	2.97	3.25	3.45	[page 35]
Full, bit-wise	8.76	5.77	2.50	[page 35]

Figure 1.14-A: Relative performance of the revbin-update and (full) revbin routines. The timing of the bit-wise update routine is normalized to 1. Values in each column should be compared, smaller values correspond to faster routines. A column labeled “ N bits” gives the timing for reversing the N least significant bits of a word.

The relative performance of the different revbin routines is shown in figure 1.14-A. As a surprise, the full-word revbin function is consistently faster than both of the update routines, mainly because the machine used (see appendix B on page 922) has a byte swap instruction. As the performance of table lookups is highly machine dependent your results can be very different.

1.14.4 Alternative techniques for in-order generation

The following loop, due to Brent Lehmann [priv. comm.], also generates the bit-reversed words in succession:

```

1      ulong n = 32; // a power of 2
2      ulong p = 0, s = 0, n2 = 2*n;
3      do
4      {
5          // here: s is the bit-reversed word
6          p += 2;
7          s ^= n - (n / (p&-p));
8      }
9      while ( p<n2 );

```

The revbin-increment is branchless but involves a division which usually is an expensive operation. With a fast bit-scan function the loop should be replaced by

```

1      do
2      {
3          p += 1;
4          s ^= n - (n >> (lowest_one_idx(p)+1));
5      }
6      while ( p<n );

```

A recursive algorithm for the generation of the bit-reversed words in order is given in [FXT: [bits/revbin-rec-demo.cc](#)]:

```

1  ulong N;
2  void revbin_rec(ulong f, ulong n)
3  {
4      // visit( f )
5      for (ulong m=N>>1; m>n; m>>=1) revbin_rec(f+m, m);

```

```
6 }
```

Call `revbin_rec(0, 0)` to generate all N-bit bit-reversed words.

A technique to generate all revbin pairs in a pseudo random order is given in section [41.4](#) on page [873](#).

1.15 Bit-wise zip

The bit-wise zip (bit-zip) operation moves the bits in the lower half to even indices and the bits in the upper half to odd indices. For example, with 8-bit words the permutation of bits is

```
[ a b c d A B C D ] |--> [ a A b B c C d D ]
```

A straightforward implementation is

```
1  ulong bit_zip(ulong a, ulong b)
2  {
3      ulong x = 0;
4      ulong m = 1, s = 0;
5      for (ulong k=0; k<(BITS_PER_LONG/2); ++k)
6      {
7          x |= (a & m) << s;
8          ++s;
9          x |= (b & m) << s;
10         m <<= 1;
11     }
12     return x;
13 }
```

Its inverse (bit-unzip) moves even indexed bits to the lower half-word and odd indexed bits to the upper half-word:

```
1  void bit_unzip(ulong x, ulong &a, ulong &b)
2  {
3      a = 0; b = 0;
4      ulong m = 1, s = 0;
5      for (ulong k=0; k<(BITS_PER_LONG/2); ++k)
6      {
7          a |= (x & m) >> s;
8          ++s;
9          m <<= 1;
10         b |= (x & m) >> s;
11         m <<= 1;
12     }
13 }
```

For a faster implementation we will use the `butterfly_*`(-)-functions which are defined in [[FXT: bits/bitbutterfly.h](#)] (64-bit version):

```
1  static inline ulong butterfly_4(ulong x)
2  // Swap in each block of 16 bits the two central blocks of 4 bits.
3  {
4      const ulong ml = 0x0f000f000f000f00UL;
5      const ulong s = 4;
6      const ulong mr = ml >> s;
7      const ulong t = ((x & ml) >> s) | ((x & mr) << s);
8      x = (x & ~(ml | mr)) | t;
9      return x;
10 }
```

The following version of the function may look more elegant but is actually slower:

```
1  static inline ulong butterfly_4(ulong x)
2  {
3      const ulong m = 0x0ff00ff00ff00ff0UL;
4      ulong c = x & m;
5      c ^= (c<<4) ^ (c>>4);
6      c &= m;
7      return x ^ c;
8 }
```

The optimized versions of the bit-zip and bit-unzip routines are [[FXT: bits/bitzip.h](#)]:

```
1  static inline ulong bit_zip(ulong x)
2  {
```

```

3  #if BITS_PER_LONG == 64
4      x = butterfly_16(x);
5  #endif
6      x = butterfly_8(x);
7      x = butterfly_4(x);
8      x = butterfly_2(x);
9      x = butterfly_1(x);
10     return x;
11 }

```

and

```

1  static inline ulong bit_unzip(ulong x)
2  {
3      x = butterfly_1(x);
4      x = butterfly_2(x);
5      x = butterfly_4(x);
6      x = butterfly_8(x);
7  #if BITS_PER_LONG == 64
8      x = butterfly_16(x);
9  #endif
10     return x;
11 }

```

Laszlo Hars suggests [priv. comm.] the following routine (version for 32-bit words), which can be obtained by making the compile-time constants explicit:

```

1  static inline uint32 bit_zip(uint32 x)
2  {
3      x = ((x & 0x0000ff00) << 8) | ((x >> 8) & 0x0000ff00) | (x & 0xff0000ff);
4      x = ((x & 0x00f000f0) << 4) | ((x >> 4) & 0x00f000f0) | (x & 0xf00ff00f);
5      x = ((x & 0x0c0c0c0c) << 2) | ((x >> 2) & 0x0c0c0c0c) | (x & 0xc3c3c3c3);
6      x = ((x & 0x22222222) << 1) | ((x >> 1) & 0x22222222) | (x & 0x99999999);
7      return x;
8  }

```

A bit-zip version for words whose upper half is zero is (64-bit version)

```

1  static inline ulong bit_zip0(ulong x)
2  // Return word with lower half bits in even indices.
3  {
4      x = (x | (x<<16)) & 0x0000ffff0000ffffUL;
5      x = (x | (x<<8)) & 0x00ff00ff00ff00ffUL;
6      x = (x | (x<<4)) & 0x0f0f0f0f0f0f0fUL;
7      x = (x | (x<<2)) & 0x3333333333333333UL;
8      x = (x | (x<<1)) & 0x5555555555555555UL;
9      return x;
10 }

```

Its inverse is

```

1  static inline ulong bit_unzip0(ulong x)
2  // Bits at odd positions must be zero.
3  {
4      x = (x | (x>>1)) & 0x3333333333333333UL;
5      x = (x | (x>>2)) & 0x0f0f0f0f0f0f0fUL;
6      x = (x | (x>>4)) & 0x00ff00ff00ff00ffUL;
7      x = (x | (x>>8)) & 0x0000ffff0000ffffUL;
8      x = (x | (x>>16)) & 0x00000000ffffffUL;
9      return x;
10 }

```

The simple structure of the routines suggests trying the following versions of bit-zip and its inverse:

```

1  static inline ulong bit_zip(ulong x)
2  {
3      ulong y = (x >> 32);
4      x &= 0xffffffffUL;
5      x = (x | (x<<16)) & 0x0000ffff0000ffffUL;
6      y = (y | (y<<16)) & 0x0000ffff0000ffffUL;
7      x = (x | (x<<8)) & 0x00ff00ff00ff00ffUL;
8      y = (y | (y<<8)) & 0x00ff00ff00ff00ffUL;
9      x = (x | (x<<4)) & 0x0f0f0f0f0f0f0fUL;
10     y = (y | (y<<4)) & 0x0f0f0f0f0f0f0fUL;
11     x = (x | (x<<2)) & 0x3333333333333333UL;
12     y = (y | (y<<2)) & 0x3333333333333333UL;
13     x = (x | (x<<1)) & 0x5555555555555555UL;

```

```

14     y = (y | (y<<1)) & 0x5555555555555555UL;
15     x |= (y<<1);
16     return x;
17 }

1  static inline ulong bit_unzip(ulong x)
2  {
3     ulong y = (x >> 1) & 0x5555555555555555UL;
4     x &= 0x5555555555555555UL;
5     x = (x | (x>>1)) & 0x3333333333333333UL;
6     y = (y | (y>>1)) & 0x3333333333333333UL;
7     x = (x | (x>>2)) & 0x0f0f0f0f0f0f0f0fUL;
8     y = (y | (y>>2)) & 0x0f0f0f0f0f0f0f0fUL;
9     x = (x | (x>>4)) & 0x00ff00ff00ff00ffUL;
10    y = (y | (y>>4)) & 0x00ff00ff00ff00ffUL;
11    x = (x | (x>>8)) & 0x0000ffff0000ffffUL;
12    y = (y | (y>>8)) & 0x0000ffff0000ffffUL;
13    x = (x | (x>>16)) & 0x00000000ffffffffUL;
14    y = (y | (y>>16)) & 0x00000000ffffffffUL;
15    x |= (y<<32);
16    return x;
17 }

```

As the statements involving the variables *x* and *y* are independent the CPU-internal parallelism can be used. However, these versions turn out to be slightly slower than those given before.

The following function moves the bits of the lower half-word of *x* into the even positions of *lo* and the bits of the upper half-word into *hi* (two versions given):

```

1  #define BPLH (BITS_PER_LONG/2)
2
3  static inline void bit_zip2(ulong x, ulong &lo, ulong &hi)
4  {
5     #if 1
6         x = bit_unzip(x);
7         lo = x & 0x5555555555555555UL;
8         hi = (x>>1) & 0x5555555555555555UL;
9     #else
10    hi = bit_zip0( x >> BPLH );
11    lo = bit_zip0( x << BPLH ) >> (BPLH) );
12 #endif
13 }

```

The inverse function is

```

1  static inline ulong bit_unzip2(ulong lo, ulong hi)
2  // Inverse of bit_zip2(x, lo, hi).
3  {
4     #if 1
5         return bit_unzip( (hi<<1) | lo );
6     #else
7         return bit_unzip0(lo) | (bit_unzip0(hi) << BPLH);
8     #endif
9  }

```

Functions that zip/unzip the bits of the lower half of two words are

```

1  static inline ulong bit_zip2(ulong x, ulong y)
2  // 2-word version:
3  // only the lower half of x and y are merged
4  {
5     return bit_zip( (y<<BPLH) + x );
6  }

```

and (64-bit version)

```

1  static inline void bit_unzip2(ulong t, ulong &x, ulong &y)
2  // 2-word version:
3  // only the lower half of x and y are filled
4  {
5     t = bit_unzip(t);
6     y = t >> BPLH;
7     x = t & 0x00000000ffffffffUL;
8  }

```

1.16 Gray code and parity

k:	bin(k)	g(k)	g ⁻¹ (k)	g(2*k)	g(2*k+1)
0:1111111
1:1111111
2:111111111111
3:111111111111
4:111111111111
5:11111111111111111
6:11111111111111111
7:11111111111111111
8:11111111111111111
9:11111111111111111
10:1111111111111111111111
11:1111111111111111111111
12:1111111111111111111111
13:1111111111111111111111
14:1111111111111111111111
15:111111111111111111111111111
16:111111111111111111111111111
17:111111111111111111111111111
18:111111111111111111111111111
19:111111111111111111111111111
20:111111111111111111111111111
21:111111111111111111111111111
22:111111111111111111111111111
23:111111111111111111111111111
24:111111111111111111111111111
25:111111111111111111111111111
26:111111111111111111111111111
27:111111111111111111111111111
28:111111111111111111111111111
29:111111111111111111111111111
30:111111111111111111111111111
31:111111111111111111111111111

Figure 1.16-A: Binary words, their Gray code, inverse Gray code, and Gray codes of even and odd values (from left to right).

The *Gray code* of a binary word can easily be computed by [FXT: [bits/graycode.h](#)]

```
1 static inline ulong gray_code(ulong x) { return x ^ (x>>1); }
```

Gray codes of consecutive values differ in one bit. Gray codes of values that differ by a power of 2 differ in two bits. Gray codes of even/odd values have an even/odd number of bits set, respectively. This is demonstrated in [FXT: [bits/gray-demo.cc](#)], whose output is given in figure 1.16-A.

To produce a random value with an even/odd number of bits set, set the lowest bit of a random number to 0/1, respectively, and return its Gray code.

Computing the inverse Gray code is slightly more expensive. As the Gray code is the bit-wise difference modulo 2, we can compute the inverse as bit-wise sums modulo 2:

```
1 static inline ulong inverse_gray_code(ulong x)
2 {
3     // VERSION 1 (integration modulo 2):
4     ulong h=1, r=0;
5     do
6     {
7         if ( x & 1 ) r^=h;
8         x >>= 1;
9         h = (h<<1)+1;
10    }
11    while ( x!=0 );
12    return r;
13 }
```

For n -bit words, n -fold application of the Gray code gives back the original word. Using the symbol G for the Gray code (operator), we have $G^n = \text{id}$, so $G^{n-1} \circ G = \text{id} = G^{-1} \circ G$. That is, applying the Gray code computation $n - 1$ times gives the inverse Gray code. Thus we can simplify to

```
1 // VERSION 2 (apply graycode BITS_PER_LONG-1 times):
2 ulong r = BITS_PER_LONG;
3 while ( --r ) x ^= x>>1;
4 return x;
```


Applying the Gray code twice is identical to $x^{\wedge}=x>>2$; applying it four times is $x^{\wedge}=x>>4$; and the idea holds for all power of 2. This leads to the most efficient way to compute the inverse Gray code:

```

1      // VERSION 3 (use: gray ** BITS_PER_LONG == id):
2      x ^= x>>1; // gray ** 1
3      x ^= x>>2; // gray ** 2
4      x ^= x>>4; // gray ** 4
5      x ^= x>>8; // gray ** 8
6      x ^= x>>16; // gray ** 16
7      // here: x = gray**31(input)
8      // note: the statements can be reordered at will
9      #if BITS_PER_LONG >= 64
10     x ^= x>>32; // for 64bit words
11     #endif
12     return x;

```

1.16.1 The parity of a binary word

The *parity* of a word is its bit-count modulo 2. The lowest bit of the inverse Gray code of a word contains the parity of the word. So we can compute the parity as [FXT: [bits/parity.h](#)]:

```

1  static inline ulong parity(ulong x)
2  // Return 0 if the number of set bits is even, else 1
3  {
4      return inverse_gray_code(x) & 1;
5  }

```

Each bit of the inverse Gray code contains the parity of the partial input left from it (including itself).

Be warned that the parity flag of many CPUs is the complement of the above. With the x86-architecture the parity bit also only takes into account the lowest byte. The following routine computes the parity of a full word [FXT: [bits/bitasm-i386.h](#)]:

```

1  static inline ulong asm_parity(ulong x)
2  {
3      x ^= (x>>16);
4      x ^= (x>>8);
5      asm ("addl $0, %0 \n"
6          "setnp %%al \n"
7          "movzx %%al, %0"
8          : "=r" (x) : "0" (x) : "eax");
9      return x;
10 }

```

The equivalent code for the AMD64 CPU is [FXT: [bits/bitasm-amd64.h](#)]:

```

1  static inline ulong asm_parity(ulong x)
2  {
3      x ^= (x>>32);
4      x ^= (x>>16);
5      x ^= (x>>8);
6      asm ("addq $0, %0 \n"
7          "setnp %%al \n"
8          "movzx %%al, %0"
9          : "=r" (x) : "0" (x) : "rax");
10     return x;
11 }

```

1.16.2 Byte-wise Gray code and parity

A byte-wise Gray code can be computed using (32-bit version)

```

1  static inline ulong byte_gray_code(ulong x)
2  // Return the Gray code of bytes in parallel
3  {
4      return x ^ ((x & 0xfefefefe)>>1);
5  }

```

Its inverse is

```

1  static inline ulong byte_inverse_gray_code(ulong x)
2  // Return the inverse Gray code of bytes in parallel
3  {

```

```

4     x ^= ((x & 0xfefefefeUL)>>1);
5     x ^= ((x & 0xfcfcfcfcUL)>>2);
6     x ^= ((x & 0xf0f0f0f0UL)>>4);
7     return x;
8 }

```

And the parities of all bytes can be computed as

```

1 static inline ulong byte_parity(ulong x)
2 // Return the parities of bytes in parallel
3 {
4     return byte_inverse_gray_code(x) & 0x01010101UL;
5 }

```

1.16.3 Incrementing (counting) in Gray code

k:	g(k)	g(2*k)	g(k) p	diff p	set
0:	{}
1:1111 1+ 1	{0}
2:1111.11+1 .	{0, 1}
3:1.1.11. 1- 1	{1}
4:11.11..11.+1. .	{1, 2}
5:1111111111 111+ 1	{0, 1, 2}
6:1.11.1.1.1-1 .	{0, 2}
7:1..1..11.. 1- 1	{2}
8:11..11...11..+1.. .	{2, 3}
9:11.111.1111.1 111.+ 1	{0, 2, 3}
10:11111111.111111+1 .	{0, 1, 2, 3}
11:111.111.1111. 1111- 1	{1, 2, 3}
12:1.1.1.1..1.1.-1. .	{1, 3}
13:1.111.1111.11 1-1.+ 1	{0, 1, 3}
14:1..11..1.1..1-1. - 1	{0, 3}
15:1...1...11... 1- 1	{3}
16:11...11....11...+1... .	{3, 4}
17:11..111..1111..1 111..+ 1	{0, 3, 4}

Figure 1.16-B: The Gray code equals the Gray code of doubled value shifted to the right once. Equivalently, we can separate the lowest bit which equals the parity of the other bits. The last column shows that the changes with each increment always happen one position left of the rightmost bit.

Let $g(k)$ be the Gray code of a number k . We are interested in efficiently generating $g(k+1)$. We can implement a fast Gray counter if we use a spare bit to keep track of the parity of the Gray code word, see figure 1.16-B The following routine does this [FXT: [bits/nextgray.h](#)]:

```

1 static inline ulong next_gray2(ulong x)
2 // With input x==gray_code(2*k) the return is gray_code(2*k+2).
3 // Let x1 be the word x shifted right once
4 // and i1 its inverse Gray code.
5 // Let r1 be the return r shifted right once.
6 // Then r1 = gray_code(i1+1).
7 // That is, we have a Gray code counter.
8 // The argument must have an even number of bits.
9 {
10     x ^= 1;
11     x ^= (lowest_one(x) << 1);
12     return x;
13 }

```

Start with $x=0$, increment with $x=\text{next_gray2}(pg)$ and use the words $g=x>>1$:

```

1     ulong x = 0;
2     for (ulong k=0; k<n2; ++k)
3     {
4         ulong g = x>>1;
5         x = next_gray2(x);
6         // here: g == gray_code(k);
7     }
8

```

This is shown in [FXT: [bits/bit-nextgray-demo.cc](#)]. To start at an arbitrary (Gray code) value g , compute

```
x = (g<<1) ^ parity(g)
```

Then use the statement `x=next_gray2(x)` for later increments.

If working with a set whose elements are the set bits in the Gray code, the parity is the set size k modulo 2. Compute the increment as follows:

1. If k is even, then goto step 2, else goto step 3.
2. If the first element is zero, then remove it, else prepend the element zero.
3. If the first element equals the second minus one, then remove the second element, else insert at the second position the element equal to the first element plus one.

A method to decrement is obtained by simply swapping the actions for even and odd parity.

When working with an array that contains the elements of the set, it is more convenient to do the described operations at the end of the array. This leads to the (loopless) algorithm for subsets in minimal-change order given in section 8.2.2 on page 206. Properties of the Gray code are discussed in [127].

1.16.4 The Thue-Morse sequence

The sequence of parities of the binary words

```
011010011001011010010110011010011001011001101001...
```

is called the *Thue-Morse sequence* (entry [A010060](#) in [312]). It appears in various seemingly unrelated contexts, see [8] and section 38.1 on page 726. The sequence can be generated with [FXT: `class thue_morse` in `bits/thue-morse.h`]:

```

1  class thue_morse
2  // Thue-Morse sequence
3  {
4  public:
5      ulong k_;
6      ulong tm_;
7
8  public:
9      thue_morse(ulong k=0) { init(k); }
10     ~thue_morse() { ; }
11
12     ulong init(ulong k=0)
13     {
14         k_ = k;
15         tm_ = parity(k_);
16         return tm_;
17     }
18
19     ulong data() { return tm_; }
20
21     ulong next()
22     {
23         ulong x = k_ ^ (k_ + 1);
24         ++k_;
25         x ^= x>>1; // highest bit that changed with increment
26         x &= 0x5555555555555555UL; // 64-bit version
27         tm_ ^= (x!=0); // change if highest changed bit was at even index
28         return tm_;
29     }
30 };
```

The rate of generation is about 366 M/s (6 cycles per update) [FXT: [bits/thue-morse-demo.cc](#)].

1.16.5 The Golay-Rudin-Shapiro sequence ‡

The function [FXT: [bits/grsnegative.h](#)]

```
1  static inline ulong grs_negative_q(ulong x) { return parity( x & (x>>1) ); }
```

returns +1 for indices where the *Golay-Rudin-Shapiro sequence* (or *GRS sequence*, entry [A020985](#) in [312]) has the value -1 . The algorithm is to count the bit-pairs modulo 2. The pairs may overlap: the

```

++
+++--
++++- ++-+
++++- ++++  +++-  --+-
++++- ++++  +++-  --+-  +++-  ++++  ----+  ++++
++++-  ++-+  +++-  --+-  +++-  ++++  ----+  ++++  +++-  ++-+  +++-  --+-  ...
      ^      ^      ^      ^      ^      ^      ^      ^
      3,   6,   11,12,13,15, 19, 22, ...

```

Figure 1.16-C: A construction for the Golay-Rudin-Shapiro (GRS) sequence.

sequence [1111] contains the three bit-pairs [11..], [.11.], and [..11]. The function returns +1 for x in the sequence

3, 6, 11, 12, 13, 15, 19, 22, 24, 25, 26, 30, 35, 38, 43, 44, 45, 47, 48, 49, 50, 52, 53, ...

This is entry [A022155](#) in [312], see also section 38.3 on page 731. The sequence can be computed by starting with two ones, and appending the left half and the negated right half of the values so far in each step, see figure 1.16-C. To compute the successor in the GRS sequence, use

```

1  static inline ulong grs_next(ulong k, ulong g)
2  // With g == grs_negative_q(k), compute grs_negative_q(k+1).
3  {
4      const ulong cm = 0x5555555555555554UL; // 64-bit version
5      ulong h = ~k; h &= -h; // == lowest_zero(k);
6      g ^= ( (h&cm) ^ ((k>>1)&h) !=0 );
7      return g;
8  }

```

With incrementing k , the lowest run of ones of k is replaced by a one at the lowest zero of k . If the length of the lowest run is odd and ≥ 2 then a change of parity happens. This is the case if the lowest zero of k is at one of the positions

bin 0101 0101 0101 0100 == hex 5 5 5 4 == cm

If the position of the lowest zero is adjacent to the next block of ones, another change of parity will occur. The element of the GRS sequence changes if exactly one of the parity changes takes place.

The update function can be used as shown in [FXT: [bits/grs-next-demo.cc](#)]:

```

1      ulong n = 65; // Generate this many values of the sequence.
2      ulong k0 = 0; // Start point of the sequence.
3      ulong g = grs_negative_q(k0);
4      for (ulong k=k0; k<k0+n; ++k)
5      {
6          // Do something with g here.
7          g = grs_next(k, g);
8      }

```

The rate of generation is about 347 M/s, direct computation gives a rate of 313 M/s.

1.16.6 The reversed Gray code

We define the *reversed Gray code* to be the bit-reversed word of the Gray code of the bit-reversed word. That is,

```
rev_gray_code(x) := revbin( gray_code( revbin(x) ) )
```

It turns out that the corresponding functions are identical to the Gray code versions up to the reversed shift operations (C-language operators '>>' replaced by '<<'). So computing the reversed Gray code is as easy as [FXT: [bits/revgraycode.h](#)]:

```
1  static inline ulong rev_gray_code(ulong x) { return x ^ (x<<1); }
```

Its inverse is

```

1  static inline ulong inverse_rev_gray_code(ulong x)
2  {
3      // use: rev_gray ** BITSPERLONG == id:
4      x ^= x<<1; // rev_gray ** 1
5      x ^= x<<2; // rev_gray ** 2
6      x ^= x<<4; // rev_gray ** 4

```

```

-----
111.1111...1111..... = 0xef0f0000 == word
1..11...1...1...1..... = gray_code
..11...1...1...1..... = rev_gray_code
1.11.1.11111.1.1111111111111111 = inverse_gray_code
1.1..1.1.....1.1..... = inverse_rev_gray_code
-----
...1...1111...1111111111111111 = 0x10f0ffff == word
...11...1...1...1..... = gray_code
..11...1...1...1..... = rev_gray_code
...11111.1.11111.1.1.1.1.1.1.1.1 = inverse_gray_code
1111.....1.1.....1.1.1.1.1.1.1.1 = inverse_rev_gray_code
-----
.....1..... = 0x2000000 == word
.....11..... = gray_code
.....11..... = rev_gray_code
.....1111111111111111111111111111 = inverse_gray_code
1111111..... = inverse_rev_gray_code
-----
111111.1111111111111111111111111111 = 0xfdfdfdf == word
1.....11..... = gray_code
.....11..... = rev_gray_code
1.1.1..1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1 = inverse_gray_code
1.1.1.11.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1 = inverse_rev_gray_code
-----

```

Figure 1.16-D: Examples of the Gray code, reversed Gray code, and their inverses with 32-bit words.

```

7     x ^= x<<8; // rev_gray ** 8
8     x ^= x<<16; // rev_gray ** 16
9     // here: x = rev_gray**31(input)
10    // note: the statements can be reordered at will
11    #if BITS_PER_LONG >= 64
12        x ^= x<<32; // for 64bit words
13    #endif
14    return x;
15 }

```

Some examples with 32-bit words are shown in figure 1.16-D.

Let G and E denote be the Gray code and reversed Gray code of a word X , respectively. Write G^{-1} and E^{-1} for their inverses. Then E preserves the lowest bit of X , while E preserves the highest. Also E preserves the lowest *set* bit of X , while E preserves the highest. Further, E^{-1} contains at each bit the parity of all bits of X right from it, including the bit itself. Especially, the word parity can be found in the highest bit of E^{-1} .

Let \bar{X} denote the complement of X , p its parity, and let S the right shift by one of G^{-1} . Then we have

$$G^{-1} \text{ XOR } E^{-1} = \begin{cases} X & \text{if } p = 0 \\ \bar{X} & \text{otherwise} \end{cases} \quad (1.16-1a)$$

$$S \text{ XOR } E^{-1} = \begin{cases} 0 & \text{if } p = 0 \\ \bar{0} & \text{otherwise} \end{cases} \quad (1.16-1b)$$

We note that taking the reversed Gray code of a binary word corresponds to multiplication with the binary polynomial $x + 1$ and the inverse reversed Gray code is a method for fast exact division by $x + 1$, see section 40.1.6 on page 826. The inverse reversed Gray code can be used to solve the reduced quadratic equation for binary normal bases, see section 42.6.2 on page 903.

1.17 Bit sequency ‡

The *sequency* of a binary word is the number of zero-one transitions in the word. A function to determine the sequency is [FXT: [bits/bitsequency.h](#)]:

```

1 static inline ulong bit_sequency(ulong x) { return bit_count( gray_code(x) ); }

```

seq=	0	1	2	3	4	5	6
1	...1.	..1.1	..1.1.	.1.1.1	1.1.1.
	11	...11.	..11.1	..11.1.	.11.1.1	1.1.1.1
		...111	...1..	..1..1	..1..1.	.1..1.1	1..1.1.1
		..1111	..111.	..1.11	..1.11.	.1.11.1	1.11.1.1
		.11111	..11..	..11.1	..11.1.	.1.1..1	1.1..1.1
		111111	..1...	..11..1	..11..1.	.1.1.1.1	1.1.1.1.1
			.1111.	..11.11	..11.11.	.11.1.1	1.1.1.1.1
			.111..	..1..1	..11.11.	.11.11.	1.1.1.1.1
			.11...	..1..11	..11.11.	.11.1..	1.1.1.1.1
			.1.....	..1.111	..1..1.	1.1.1.1.	1.1.1.1.1
			11111.	..111.1	..1..11.	1.1.1.1.	1.1.1.1.1
			1111..	..111..1	..1..1..	1.1.1..	1.1.1.1.1
			111...	..111.11	..1.111.	1.111.	1.1.1.1.1
			11....	..11...1	..1.11..	1.11..	1.1.1.1.1
			1.....	..11...11	..1.1...1	1.1...1	1.1.1.1.1
				11.111	1.1.111	1.1.111	1.1.111
				1...1	1.1...1	1.1...1	1.1...1
				1..11	1.1..11	1.1..11	1.1..11
				1..111	1.1.111	1.1.111	1.1.111
				1.1111	1.1.1111	1.1.1111	1.1.1111

Figure 1.17-A: 6-bit words of prescribed sequency as generated by `next_sequency()`.

The function assumes that all bits to the left of the word are zero and all bits to the right are equal to the lowest bit, see figure 1.17-A. For example, the sequency of the 8-bit word [00011111] is one. To take the lowest bit into account, add it to the sequency (then all sequencies are even).

The minimal binary word with given sequency can be computed as follows:

```

1 static inline ulong first_sequency(ulong k)
2 // Return the first (i.e. smallest) word with sequency k,
3 // e.g. 00..00010101010 (seq 8)
4 // e.g. 00..00101010101 (seq 9)
5 // Must have: 0 <= k <= BITS_PER_LONG
6 {
7     return inverse_gray_code( first_comb(k) );
8 }

```

A faster version is (32-bit branch only):

```

1     if ( k==0 ) return 0;
2     const ulong m = 0xaaaaaaaaUL;
3     return m >> (BITS_PER_LONG-k);

```

The maximal binary word with given sequency can be computed via

```

1 static inline ulong last_sequency(ulong k)
2 // Return the last (i.e. biggest) word with sequency k.
3 {
4     return inverse_gray_code( last_comb(k) );
5 }

```

The functions `first_comb(k)` and `last_comb(k)` return a word with k bits set at the low and high end, respectively (see section 1.24 on page 62).

For the generation of all words with a given sequency, starting with the smallest, we use a function that computes the next word with the same sequency:

```

1 static inline ulong next_sequency(ulong x)
2 {
3     x = gray_code(x);
4     x = next_colex_comb(x);
5     x = inverse_gray_code(x);
6     return x;
7 }

```

The inverse function, returning the previous word with the same sequency, is

```

1 static inline ulong prev_sequency(ulong x)
2 {
3     x = gray_code(x);
4     x = prev_colex_comb(x);
5     x = inverse_gray_code(x);
6     return x;

```

```
7 }

```

The list of all 6-bit words ordered by sequency is shown in figure 1.17-A. It was created with the program [FXT: [bits/bitsequency-demo.cc](#)].

The sequency of a word can be complemented as follows (32-bit version):

```
1 static inline ulong complement_sequency(ulong x)
2 // Return word whose sequency is BITS_PER_LONG - s
3 // where s is the sequency of x
4 {
5     return x ^ 0xaaaaaaaaUL;
6 }
```

1.18 Powers of the Gray code †

<pre>1..... .1..... ..1..... ...1.....1.....1 G^0=id</pre>	<pre>11..... .11..... ..11..... ...11.....11.....11 G^1=G</pre>	<pre>1,1..... .1,1..... ..1,1..... ...1,1.....1,1.....1,1 G^2</pre>	<pre>1111..... .1111..... ..1111..... ...1111.....1111.....1111 G^3</pre>	<pre>1..1... .1..1... ..1..1... ...1..1...1..1...1..1 G^4</pre>	<pre>11..11.. .11..11.. ..11..11.. ...11..11..11..11..11..11 G^5</pre>	<pre>1,1,1,1,1,1 .1,1,1,1,1,1 ..1,1,1,1,1,1 ...1,1,1,1,1,11,1,1,1,1,11,1,1,1,1,1 G^6</pre>	<pre>11111111 .11111111 ..11111111 ...111111111111111111111111 G^7=G^(-1)</pre>
<pre>1i..... .i..... ..i..... ...i.....i.....i E^0=id</pre>	<pre>1i..... .i..... ..i..... ...i.....i.....i E^1=E</pre>	<pre>1i..... .i..... ..i..... ...i.....i.....i E^2</pre>	<pre>1i..... .i..... ..i..... ...i.....i.....i E^3</pre>	<pre>1i..... .i..... ..i..... ...i.....i.....i E^4</pre>	<pre>1i..... .i..... ..i..... ...i.....i.....i E^5</pre>	<pre>1i..... .i..... ..i..... ...i.....i.....i E^6</pre>	<pre>1i..... .i..... ..i..... ...i.....i.....i E^7=E^(-1)</pre>

Figure 1.18-A: Powers of the matrices for the Gray code (top) and the reversed Gray code (bottom).

The Gray code is a bit-wise linear transform of a binary word. The 2^k -th power of the Gray code of x can be computed as $x \wedge (x \gg k)$. The e -th power can be computed as the bit-wise sum of the powers corresponding to the bits in the exponent. This motivates [FXT: [bits/graypower.h](#)]:

```
1 static inline ulong gray_pow(ulong x, ulong e)
2 // Return (gray_code**e)(x)
3 // gray_pow(x, 1) == gray_code(x)
4 // gray_pow(x, BITS_PER_LONG-1) == inverse_gray_code(x)
5 {
6     e &= (BITS_PER_LONG-1); // modulo BITS_PER_LONG
7     ulong s = 1;
8     while ( e )
9     {
10         if ( e & 1 ) x ^= x >> s; // gray ** s
11         s <<= 1;
12         e >>= 1;
13     }
14     return x;
15 }
```

The Gray code $g = [g_0, g_1, \dots, g_7]$ of a 8-bit binary word $x = [x_0, x_1, \dots, x_7]$ can be expressed as a matrix multiplication over GF(2) (dots for zeros):

$$\begin{array}{rcl}
 \mathbf{g} & = & \mathbf{G} \quad \mathbf{x} \\
 [g_0] & & [11 \dots \dots] [x_0] \\
 [g_1] & & [. 11 \dots \dots] [x_1] \\
 [g_2] & & [.. 11 \dots \dots] [x_2] \\
 [g_3] & = & [... 11 \dots \dots] [x_3] \\
 [g_4] & & [.... 11 \dots \dots] [x_4] \\
 [g_5] & & [..... 11 \dots \dots] [x_5] \\
 [g_6] & & [..... 11] [x_6] \\
 [g_7] & & [..... 1] [x_7]
 \end{array}$$

The powers of the Gray code correspond to multiplication with powers of the matrix G , shown in figure 1.18-A (bottom). The powers of the inverse Gray code for N -bit words (where N is a power of 2)

can be computed by the relation $G^e G^{N-e} = G^N = \text{id}$.

```

1  static inline ulong inverse_gray_pow(ulong x, ulong e)
2  // Return (inverse_gray_code**(e))(x)
3  // == (gray_code**(-e))(x)
4  // inverse_gray_pow(x, 1) == inverse_gray_code(x)
5  // inverse_gray_pow(x, BITS_PER_LONG-1) == gray_code(x)
6  {
7      return gray_pow(x, -e);
8  }

```

The matrices corresponding to the powers of the reversed Gray code are shown in figure 1.18-A (bottom). We just have to reverse the shift operator in the functions:

```

1  static inline ulong rev_gray_pow(ulong x, ulong e)
2  // Return (rev_gray_code**e)(x)
3  {
4      e &= (BITS_PER_LONG-1); // modulo BITS_PER_LONG
5      ulong s = 1;
6      while ( e )
7      {
8          if ( e & 1 ) x ^= x << s; // rev_gray ** s
9          s <<= 1;
10         e >>= 1;
11     }
12     return x;
13 }

```

The inverse function is

```

1  static inline ulong inverse_rev_gray_pow(ulong x, ulong e)
2  // Return (inverse_rev_gray_code**(e))(x)
3  {
4      return rev_gray_pow(x, -e);
5  }

```

1.19 Invertible transforms on words ‡

The functions presented in this section are invertible transforms on binary words. The names are chosen as ‘some code’, emphasizing the result of the transforms, similar to the convention used with the name ‘Gray code’. The functions are given in [FXT: [bits/bittransforms.h](#)].

In the transform (*blue code*)

```

1  static inline ulong blue_code(ulong a)
2  {
3      ulong s = BITS_PER_LONG >> 1;
4      ulong m = ~0UL << s;
5      do
6      {
7          a ^= ( (a&m) >> s );
8          s >>= 1;
9          m ^= (m>>s);
10     }
11     while ( s );
12     return a;
13 }

```

the masks ‘m’ are (32-bit binary)

```

11111111111111111111111111111111
11111111111111111111111111111111
11111111111111111111111111111111
11111111111111111111111111111111
11111111111111111111111111111111
11111111111111111111111111111111
11111111111111111111111111111111
11111111111111111111111111111111

```

The same masks are used in the *yellow code*

```

1  static inline ulong yellow_code(ulong a)
2  {
3      ulong s = BITS_PER_LONG >> 1;
4      ulong m = ~0UL >> s;
5      do
6      {
7          a ^= ( (a&m) << s );
8          s >>= 1;

```


red			green		
0:	0	0	0
1:	1.....	1	32	11111111111111111111111111111111	32
2:	11.....	2	16	1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1	16
3:	1.1.....	1	16	1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1	16
4:	1.1.1.....	2	16	..11..11..11..11..11..11..11..11..11	16
5:	..1.....	1	16	11..11..11..11..11..11..11..11..11..	16
6:	..11.....	2	16	..11..11..11..11..11..11..11..11..11..	16
7:	111.....	3	16	1..11..11..11..11..11..11..11..11..1	16
8:	1111.....	4	8	...1...1...1...1...1...1...1...1	8
9:	..111.....	3	24	111.111.111.111.111.111.111.111.	24
10:	..11.1.....	2	8	1..1...1...1...1...1...1...1...1	8
11:	1.11.....	3	24	1.111.111.111.111.111.111.111.111.11	24
12:	1.1.1.....	2	8	..1...1...1...1...1...1...1...1	8
13:	11.1.....	3	24	11.111.111.111.111.111.111.111.111.1	24
14:	1.1.1.....	2	24	1.111.111.111.111.111.111.111.111.111	24
15:	..1.....	1	8	1...1...1...1...1...1...1...1	8
16:	1..1.....	2	16	...1111...1111...1111...1111...1111	16
17:	..1.1.....	1	16	1111...1111...1111...1111...1111	16
18:	1.1.1.....	2	16	1.11.1.1.11.1.11.1.11.1.11.1.11.1.1	16
19:	11.1.....	3	16	1.1.1.11.1.11.1.11.1.11.1.11.1.11	16
20:	..1.1.....	2	16	..1111...1111...1111...1111...1111	16
21:	1.1.1.....	3	16	11...1111...1111...1111...1111...11	16
22:	111.1.....	4	16	..11.1.1.11.1.11.1.11.1.11.1.11.1.1	16
23:	..11.1.....	3	16	1..1.11.1..1.11.1..1.11.1..1.11.1..	16
24:	..1111.....	4	16	...1111...1111...1111...1111...1111	16
25:	11111.....	5	16	111...1111...1111...1111...1111...1	16
26:	1.111.....	4	16	1..1.11.1.1.11.1.11.1.11.1.11.1.11	16
27:	..111.....	3	16	1.11.1.1.11.1.11.1.11.1.11.1.11.1..	16
28:	11.11.....	4	16	..1.11.1..1.11.1..1.11.1..1.11.1..	16
29:	..1.11.....	3	16	11.1.1.11.1.11.1.11.1.11.1.11.1.1.	16
30:	..11.....	2	16	..1111...1111...1111...1111...1111	16
31:	1..11.....	3	16	1....1111...1111...1111...1111...111	16

Figure 1.19-B: Red and green transforms of the binary words 0, 1, ..., 31.

```

12         m ^= (m<<s);
13     }
14     while ( s );
15     return a;
16 }

and (green code)

1  static inline ulong green_code(ulong a)
2  {
3      ulong s = BITS_PER_LONG >> 1;
4      ulong m = ~0UL << s;
5      do
6      {
7          ulong u = a & m;
8          ulong v = a ^ u;
9          a = v ^ (u>>s);
10         a ^= (v<<s);
11         s >>= 1;
12         m ^= (m>>s);
13     }
14     while ( s );
15     return a;
16 }

```

use the masks

```

.....11111111111111111111111111111111
.....11111111111111111111111111111111
..11111111111111111111111111111111
.i11111111111111111111111111111111
.i.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1

```

The transforms of the binary words up to 31 are shown in figure 1.19-B, which was created with the program [FXT: bits/bittransforms-red-demo.cc]. The red code can also be computed by the statement

```
revbin( blue_code( x ) );
```

and the green code by

```
blue_code( revbin( x ) );
```

		i	r	B	Y	R	E
i	i	r	B	Y	R	E	
r	r	i	R*	E*	B*	Y*	
B	B	E*	i	R*	Y*	r*	
Y	Y	R*	E*	i	r*	B*	
R	R	Y*	r*	B*	E	i	
E	E	B*	Y*	r*	i	R	

Figure 1.19-C: Multiplication table for the transforms.

1.19.1 Relations between the transforms

We write B for the blue code (transform), Y for the yellow code and r for bit-reversal (the `revbin`-function). We have the following relations between B and Y :

$$B = Y r Y = r Y r \quad (1.19-1a)$$

$$Y = B r B = r B r \quad (1.19-1b)$$

$$r = Y B Y = B Y B \quad (1.19-1c)$$

As said, B and Y are self-inverse:

$$B^{-1} = B, \quad B B = \text{id} \quad (1.19-2a)$$

$$Y^{-1} = Y, \quad Y Y = \text{id} \quad (1.19-2b)$$

We write R for the red code, and E for the green code. The red code and the green code are not involutions (square roots of identity) but third roots of identity:

$$R R R = \text{id}, \quad R^{-1} = R R = E \quad (1.19-3a)$$

$$E E E = \text{id}, \quad E^{-1} = E E = R \quad (1.19-3b)$$

$$R E = E R = \text{id} \quad (1.19-3c)$$

Figure 1.19-C shows the multiplication table. The R in the third column of the second row says that $r B = R$. The letter i is used for identity (id). An asterisk says that $x y \neq y x$.

By construction we have

$$R = r B \quad (1.19-4a)$$

$$E = r Y \quad (1.19-4b)$$

Relations between R and E are:

$$R = E r E = r E r \quad (1.19-5a)$$

$$E = R r R = r R r \quad (1.19-5b)$$

$$R = R E R \quad (1.19-5c)$$

$$E = E R E \quad (1.19-5d)$$

For the bit-reversal we have

$$r = Y R = R B = B E = E Y \quad (1.19-6)$$

Some products for the transforms are

$$B = R Y = Y E = R B R = E B E \quad (1.19-7a)$$

$$Y = E B = B R = R Y R = E Y E \quad (1.19-7b)$$

$$R = B Y = B E B = Y E Y \quad (1.19-7c)$$

$$E = Y B = B R B = Y R Y \quad (1.19-7d)$$

Some triple products that give the identical transform are

$$\text{id} = BYE = RYB \tag{1.19-8a}$$

$$\text{id} = EBY = BR Y \tag{1.19-8b}$$

$$\text{id} = YEB = YBR \tag{1.19-8c}$$

1.19.2 Relations to Gray code and reversed Gray code

Write g for the Gray code, then:

$$gBgB = \text{id} \tag{1.19-9a}$$

$$gBg = B \tag{1.19-9b}$$

$$g^{-1}Bg^{-1} = B \tag{1.19-9c}$$

$$gB = Bg^{-1} \tag{1.19-9d}$$

Let S_k be the operator that rotates a word by k bits (bit 0 is moved to position k), then

$$YS_{+1}Y = g \tag{1.19-10a}$$

$$YS_{-1}Y = g^{-1} \tag{1.19-10b}$$

$$YS_kY = g^k \tag{1.19-10c}$$

Shift in the sequency domain is bit-wise derivative in time domain. Relation 1.19-10c, together with an algorithm to generate the cycle leaders of the Gray permutation (section 2.12.1 on page 128) gives a curious method to generate the binary necklaces whose length is a power of 2, described in section 18.1.6 on page 376. Let e be the operator for the reversed Gray code, then

$$BS_{+1}B = e^{-1} \tag{1.19-11a}$$

$$BS_{-1}B = e \tag{1.19-11b}$$

$$BS_kB = e^{-k} \tag{1.19-11c}$$

1.19.3 Fixed points of the blue code ‡

0 = : =	0	16 = .1..... :	.1...1.... =	272
1 =1 :1 =	1	17 = .1...1 :	.1.11.1... =	360
2 =1 :11 =	6	18 = .1..1. :	.1.....1... =	260
3 =11 :111 =	7	19 = .1..11 :	.1.1111... =	380
4 = ...1. :1.1 =	20	20 = .1.1. :	.1...1.11. =	278
5 = ...1.1 :1..1 =	18	21 = .1.1.1 :	.1.11.111. =	366
6 = ...11 :1.1.1 =	21	22 = .1.11. :	.1.....1. =	258
7 = ...111 :1..11 =	19	23 = .1.111 :	.1.1111.1 =	378
8 = ..1... :	...1111... =	120	24 = .11... :	.1...1...1 =	273
9 = ..1..1 :	...11.11... =	108	25 = .11..1 :	.1.11.1..1 =	361
10 = ..1.1 :	...111111. =	126	26 = .11.1. :	.1...1.1.1 =	261
11 = ..1.11 :	...11.1.1. =	106	27 = .11.11 :	.1.11111.1 =	381
12 = ..11. :	...1111..1 =	121	28 = .111. :	.1...1.111 =	279
13 = ..11.1 :	...11.11.1 =	109	29 = .111.1 :	.1.11.1111 =	367
14 = ..111 :	...1111111 =	127	30 = .1111. :	.1.....11 =	259
15 = ..1111 :	...11.1.11 =	107	31 = .11111 :	.1.1111.11 =	379

Figure 1.19-D: The first fixed points of the blue code. The highest bit of all fixed points lies at an even index. There are $2^{n/2}$ fixed points with highest bit at index n .

The sequence of fixed points of the blue code is (entry A118666 in [312])

0, 1, 6, 7, 18, 19, 20, 21, 106, 107, 108, 109, 120, 121, 126, 127, 258, 259, ...

If f is a fixed point, then f XOR 1 is also a fixed point. Further, $2(f \text{ XOR } (2f))$ is a fixed point. These facts can be cast into a function that returns a unique fixed point for each argument [FXT: bits/blue-fixed-points.h]:

```

1  static inline ulong blue_fixed_point(ulong s)
2  {
3      if ( 0==s ) return 0;
4      ulong f = 1;
5      while ( s>1 )
6      {
7          f ^= (f<<1);
8          f <<= 1;
9          f |= (s&1);
10         s >>= 1;
11     }
12     return f;
13 }

```

The output for the first few arguments is shown in figure 1.19-D. Note that the fixed points are not in ascending order. The list was created by the program [FXT: bits/bittransforms-blue-fp-demo.cc].

Now write $f(x)$ for the binary polynomial corresponding to f (see chapter 40 on page 822), if $f(x)$ is a fixed point (that is, $Bf(x) = f(x+1) = f(x)$), then both $(x^2+x)f(x)$ and $1+(x^2+x)f(x)$ are fixed points. The function `blue_fixed_point()` repeatedly multiplies by x^2+x and adds one if the corresponding bit of the argument is set.

For the inverse function, we exploit that polynomial division by $x+1$ can be done with the inverse reversed Gray code (see section 1.16.6 on page 45) if the polynomial is divisible by $x+1$:

```

1  static inline ulong blue_fixed_point_idx(ulong f)
2  // Inverse of blue_fixed_point()
3  {
4      ulong s = 1;
5      while ( f )
6      {
7          s <<= 1;
8          s ^= (f & 1);
9          f >>= 1;
10         f = inverse_rev_gray_code(f); // == bitpol_div(f, 3);
11     }
12     return s >> 1;
13 }

```

1.19.4 More transforms by symbolic powering

The idea of powering a transform (as with the Gray code, see section 1.18 on page 48) can be applied to the ‘color’-transforms as exemplified for the blue code:

```

1  static inline ulong blue_xcode(ulong a, ulong x)
2  {
3      x &= (BITS_PER_LONG-1); // modulo BITS_PER_LONG
4      ulong s = BITS_PER_LONG >> 1;
5      ulong m = ~0UL << s;
6      while ( s )
7      {
8          if ( x & 1 ) a ^= ( a&m ) >> s ;
9          x >>= 1;
10         s >>= 1;
11         m ^= (m>>s);
12     }
13     return a;
14 }

```

The result is *not* the power of the blue code which would be pretty boring as $BB = \text{id}$. The transforms (and the equivalents for Y , R and E , see [FXT: bits/bitxtransforms.h]) are more interesting: all relations between the transforms are still valid, if the symbolic exponent is identical with all terms in the relation. For example, we had $BB = \text{id}$, now $B^x B^x = \text{id}$ is true for all x . Similarly, $EE = R$ now has to be $E^x E^x = R^x$. That is, we have BITS_PER_LONG different versions of our four transforms that share their properties with the ‘simple’ versions. Among them are BITS_PER_LONG transforms B^x and Y^x that are involutions and E^x and R^x that are third roots of the identity: $E^x E^x E^x = R^x R^x R^x = \text{id}$.

While not powers of the simple versions, we still have $B^0 = Y^0 = R^0 = E^0 = \text{id}$. Further, let e be the ‘exponent’ of all ones and Z be any of the transforms, then $Z^e = Z$. Writing ‘+’ for the XOR operation,

we have $Z^x Z^y = Z^{x+y}$ and so $Z^x Z^y = Z$ whenever $x + y = e$.

1.19.5 The building blocks of the transforms

Consider the following transforms on 2-bit words where addition is bit-wise (that is, XOR):

$$\text{id}_2 v = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} a \\ b \end{bmatrix} \quad (1.19-12a)$$

$$r_2 v = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} b \\ a \end{bmatrix} \quad (1.19-12b)$$

$$B_2 v = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} a+b \\ b \end{bmatrix} \quad (1.19-12c)$$

$$Y_2 v = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} a \\ a+b \end{bmatrix} \quad (1.19-12d)$$

$$R_2 v = \begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} b \\ a+b \end{bmatrix} \quad (1.19-12e)$$

$$E_2 v = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} a+b \\ a \end{bmatrix} \quad (1.19-12f)$$

It can easily be verified that for these the same relations hold as for id , r , B , Y , R , E . In fact the ‘color-transforms’, bit-reversal, and identity are the transforms obtained as repeated Kronecker-products of the matrices (see section 23.3 on page 462). The transforms are linear over GF(2):

$$Z(\alpha a + \beta b) = \alpha Z(a) + \beta Z(b) \quad (1.19-13)$$

The corresponding version of the bit-reversal is [FXT: [bits/revbin.h](#)]:

```

1  static inline ulong xrevbin(ulong a, ulong x)
2  {
3      x &= (BITS_PER_LONG-1); // modulo BITS_PER_LONG
4      ulong s = BITS_PER_LONG >> 1;
5      ulong m = ~0UL >> s;
6      while ( s )
7      {
8          if ( x & 1 ) a = ( (a & m) << s ) ^ ( (a & (~m)) >> s );
9          x >>= 1;
10         s >>= 1;
11         m ^= (m<<s);
12     }
13     return a;
14 }
```

Then, for example, $R^x = r^x B^x$ (see relation 1.19-4a on page 52). The yellow code is the bit-wise Reed-Muller transform (described in section 23.11 on page 486) of a binary word. The symbolic powering is equivalent to selecting individual levels of the transform.

1.20 Scanning for zero bytes

The following function (32-bit version) determines if any sub-byte of the argument is zero from [FXT: [bits/zerobyte.h](#)]:

```

1  static inline ulong contains_zero_byte(ulong x)
2  {
3      return ((x-0x01010101UL)^x) & (~x) & 0x80808080UL;
4  }
```

It returns zero when x contains no zero-byte and nonzero when it does. The idea is to subtract one from each of the bytes and then look for bytes where the borrow propagated all the way to the most significant bit. A simplified version is given in [215, sect.7.1.3, rel.90]:

```

1      return 0x80808080UL & ( x - 0x01010101UL ) & ~x;
```

To scan for other values than zero (e.g. 0xa5), we can use

```
contains_zero_byte( x ^ 0xa5a5a5a5UL )
```

For very long strings and word sizes of 64 or more bits the following function may be a win [FXT: aux1/bytescan.cc]:

```

1  ulong long_strlen(const char *str)
2  // Return length of string starting at str.
3  {
4      ulong x;
5      const char *p = str;
6
7      // Alignment: scan bytes up to word boundary:
8      while ( (ulong)p % BYTES_PER_LONG )
9      {
10         if ( 0 == *p ) return (ulong)(p-str);
11         ++p;
12     }
13
14     x = *(ulong *)p;
15     while ( ! contains_zero_byte(x) )
16     {
17         p += BYTES_PER_LONG;
18         x = *(ulong *)p;
19     }
20
21     // now a zero byte is somewhere in x:
22     while ( 0 != *p ) { ++p; }
23
24     return (ulong)(p-str);
25 }
```

1.21 Inverse and square root modulo 2^n

1.21.1 Computation of the inverse

The inverse modulo 2^n where n is the number of bits in a word can be computed using an iteration (see section 29.1.5 on page 569) with quadratic convergence. The number to be inverted has to be odd [FXT: bits/bit2adic.h]:

```

1  static inline ulong inv2adic(ulong x)
2  // Return inverse modulo 2**BITS_PER_LONG
3  // x must be odd
4  // The number of correct bits is doubled with each step
5  // ==> loop is executed prop. log_2(BITS_PER_LONG) times
6  // precision is 3, 6, 12, 24, 48, 96, ... bits (or better)
7  {
8      if ( 0==(x&1) ) return 0; // not invertible
9      ulong i = x; // correct to three bits at least
10     ulong p;
11     do
12     {
13         p = i * x;
14         i *= (2UL - p);
15     }
16     while ( p!=1 );
17     return i;
18 }
```

Let m be the modulus (a power of 2), then the computed value i is the inverse of x modulo m : $i \equiv x^{-1} \pmod{m}$. It can be used for the *exact division*: to compute the quotient a/x for a number a that is known to be divisible by x , simply multiply by i . This works because $a = bx$ (a is divisible by x), so $ai \equiv bxi \equiv b \pmod{m}$.

1.21.2 Exact division by $C = 2^k \pm 1$

We use the following relation where $Y = 1 - C$:

$$\frac{A}{C} = \frac{A}{1-Y} = A(1+Y)(1+Y^2)(1+Y^4)(1+Y^8)\dots(1+Y^{2^n}) \pmod{Y^{2^{n+1}}} \quad (1.21-1)$$

The relation can be used for efficient exact division over \mathbb{Z} by $C = 2^k \pm 1$. For $C = 2^k + 1$ use

$$\frac{A}{C} = A(1-2^k)(1+2^{k^2})(1+2^{k^4})(1+2^{k^8})\dots(1+2^{k^{2^u}}) \pmod{2^N} \quad (1.21-2)$$

where $k \cdot 2^u \geq N$. For $C = 2^k - 1$ use ($A/C = -A/-C$)

$$\frac{A}{C} = -A(1+2^k)(1+2^{k^2})(1+2^{k^4})(1+2^{k^8})\dots(1+2^{k^{2^u}}) \pmod{2^N} \quad (1.21-3)$$

The equivalent method for exact division by polynomials (over $\text{GF}(2)$) is given in section 40.1.6 on page 826.

1.21.3 Computation of the square root

x =1 = 1	x = 11::11::11::11::11::11::11::11::11 = 5
inv =1	inv = 11::11::11::11::11::11::11::11::11 = 5
$sqrt$ =1	$sqrt$ = 11::11::11::11::11::11::11::11::11 = 5
x = 11111111111111111111111111111111 = -1	x = 11111111111111111111111111111111.11 = -5
inv = 11111111111111111111111111111111	inv = ..11::11::11::11::11::11::11::11::11
x =1. = 2	x =11. = 6
x = 11111111111111111111111111111111. = -2	x = 11111111111111111111111111111111.1. = -6
x =1.11 = 3	x =111 = 7
inv = i:i:i:i:i:i:i:i:i:i:i:i:i:i:i:i:11	inv = i:i:i:i:i:i:i:i:i:i:i:i:i:i:i:i:111
x = 11111111111111111111111111111111.1 = -3	x = 11111111111111111111111111111111..1 = -7
inv = .1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1	inv = .1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1
$sqrt$ =1.1	$sqrt$ = i:i:i:i:i:i:i:i:i:i:i:i:i:i:i:i:11
x =1. = 4	x =1... = 8
$sqrt$ =1.	$sqrt$ =1... = 8
x = 11111111111111111111111111111111.. = -4	x = 11111111111111111111111111111111... = -8
	x =1..11 = 9
	inv =11
	$sqrt$ = 11111111111111111111111111111111.1

Figure 1.21-A: Examples of the inverse and square root modulo 2^n of x where $-9 \leq x \leq +9$. Where no inverse or square root is given, it does not exist.

With the inverse square root we choose the start value to match $\lfloor d/2 \rfloor + 1$ as that guarantees four bits of initial precision. Moreover, we control which of the two possible values of the inverse square root is computed. The argument modulo 8 has to be equal to 1.

```

1  static inline ulong invsqrt2adic(ulong d)
2  // Return inverse square root modulo 2**BITS_PER_LONG
3  // Must have: d==1 mod 8
4  // The number of correct bits is doubled with each step
5  // ==> loop is executed prop. log_2(BITS_PER_LONG) times
6  // precision is 4, 8, 16, 32, 64, ... bits (or better)
7  {
8      if ( 1 != (d&7) ) return 0; // no inverse sqrt
9      // start value: if d == ****10001 ==> x := ****1001
10     ulong x = (d >> 1) | 1;
11     ulong p, y;
12     do
13     {
14         y = x;
15         p = (3 - d * y * y);
16         x = (y * p) >> 1;
17     }
18     while ( x!=y );
19     return x;
20 }

```


The square root is computed as $d \cdot 1/\sqrt{d}$:

```

1  static inline ulong sqrt2adic(ulong d)
2  // Return square root modulo 2**BITS_PER_LONG
3  // Must have: d==1 mod 8 or d==4 mod 32, d==16 mod 128
4  // ... d==4**k mod 4**(k+3)
5  // Result undefined if condition does not hold
6  {
7      if ( 0==d ) return 0;
8      ulong s = 0;
9      while ( 0==(d&1) ) { d >>= 1; ++s; }
10     d *= invsqrt2adic(d);
11     d <<= (s>>1);
12     return d;
13 }

```

Note that the square root modulo 2^n is something completely different from the integer square root in general. If the argument d is a perfect square, then the result is $\pm\sqrt{d}$. The output of the program [FXT: bits/bit2adic-demo.cc] is shown in figure 1.21-A. For further information see [213, ex.31, p.213], [135, chap.6, p.126], and also [208].

1.22 Radix -2 (minus two) representation

The radix -2 representation of a number n is

$$n = \sum_{k=0}^{\infty} t_k (-2)^k \quad (1.22-1)$$

where the t_k are zero or one. For integers n the sum is terminating: the highest nonzero t_k is at most two positions beyond the highest bit of the binary representation of the absolute value of n (with two's complement).

1.22.1 Conversion from binary

k:	bin(k)	m=bin2neg(k)	g=gray(m)	dec(g)
0:	0 <= 0
1:111	1 <= 1
2:11111.1	5
3:1111111..	4
4:1..1..11.	2
5:1.11.1111	3 <= 5
6:11.	..11.1.	..1.111	19
7:111	..11.11	..1.11.	18
8:	...1...	..11...	..1.1..	20
9:	...1..1	..11..1	..1.1.1	21
10:	...1.1.	..1111.	..1...1	17
11:	...1.11	..11111	..1....	16
12:	...11..	..111..	..1...1	14
13:	...11.1	..111.1	..1...11	15
14:	...111.	..1...1.	..11.11	7
15:	...1111	..1...11	..11.1.	6
16:	..1....	..1....	..11...	8
17:	..1...1	..1...1	..11..1	9
18:	..1...1.	..1.11.	..111.1	13
19:	..1...11	..1.111	..111..	12
20:	..1.1..	..1.1..	..1111.	10
21:	..1.1.1	..1.1.1	..11111	11 <= 21
22:	..1.11.	11.1.1.	1.11111	75
23:	..1.111	11.1.11	1.1111.	74
24:	..11...	11.1...	1.111..	76
25:	..11..1	11.1..1	1.111.1	77
26:	..11.1.	11.111.	1.11..1	73
27:	..11.11	11.1111	1.11...	72
28:	..111..	11.11..	1.11.1.	70
29:	..111.1	11.11.1	1.11.11	71
30:	..1111.	11...1.	1.1...11	79
31:	..11111	11...11	1.1...1.	78

Figure 1.22-A: Radix -2 representations and their Gray codes. Lines ending in '<=N' indicate that all values $\leq N$ occur in the last column up to that point.

A surprisingly simple algorithm to compute the coefficients t_k of the radix -2 representation of a binary number is [39, item 128] [FXT: [bits/negbin.h](#)]:

```

1  static inline ulong bin2neg(ulong x)
2  // binary --> radix(-2)
3  {
4      const ulong m = 0xaaaaaaaaUL; // 32 bit version
5      x += m;
6      x ^= m;
7      return x;
8  }

```

An example:

$$14 \text{ --> } \dots 111. = 16 - 2 = (-2)^4 + (-2)^1$$

The inverse routine executes the inverse of the two steps in reversed order:

```

1  static inline ulong neg2bin(ulong x)
2  // radix(-2) --> binary
3  // inverse of bin2neg()
4  {
5      const ulong m = 0xaaaaaaaaUL; // 32-bit version
6      x ^= m;
7      x -= m;
8      return x;
9  }

```

Figure 1.22-A shows the output of the program [FXT: [bits/negbin-demo.cc](#)]. The sequence of Gray codes of the radix -2 representation is a Gray code for the numbers in the range $0, \dots, k$ for the following values of k (entry [A002450](#) in [312]):

$$k = 1, 5, 21, 85, 341, 1365, 5461, 21845, 87381, 349525, 1398101, \dots, (4^n - 1)/3$$

1.22.2 Fixed points of the conversion †

0:1	64:1.....	256: ..1.....	320: ..1.1.....
1:1.1	65:1.1....1	257: ..1.1.....1	321: ..1.1.1....1
4:1.1.1	68:1.1.1..1	260: ..1.1.1....1	324: ..1.1.1.1..1
5:1.1.1.1	69:1.1.1.1.1	261: ..1.1.1.1..1	325: ..1.1.1.1.1.1
16:1.....	80:1.1.....	272: ..1.1.1....	336: ..1.1.1.1....
17:1.1....	81:1.1.1....	273: ..1.1.1.1..1	337: ..1.1.1.1.1..1
20:1.1.1..	84:1.1.1.1..	276: ..1.1.1.1.1..	340: ..1.1.1.1.1.1..
21:1.1.1.1.1	85:1.1.1.1.1.1	277: ..1.1.1.1.1.1.1	341: ..1.1.1.1.1.1.1.1

Figure 1.22-B: The fixed points of the conversion and their binary representations (dots denote zeros).

The sequence of fixed points of the conversion starts as

$$0, 1, 4, 5, 16, 17, 20, 21, 64, 65, 68, 69, 80, 81, 84, 85, 256, \dots$$

The binary representations have ones only at even positions (see figure 1.22-B). This is the *Moser – De Bruijn sequence*, entry [A000695](#) in [312]. The generating function of the sequence is

$$\frac{1}{1-x} \sum_{j=0}^{\infty} \frac{4^j x^{2^j}}{1+x^{2^j}} = x + 4x^2 + 5x^3 + 16x^4 + 17x^5 + 20x^6 + 21x^7 + 64x^8 + 65x^9 + \dots \quad (1.22-2)$$

The sequence also appears as exponents in the power series (see also section 38.10.1 on page 750)

$$\prod_{k=0}^{\infty} (1 + x^{4^k}) = 1 + x + x^4 + x^5 + x^{16} + x^{17} + x^{20} + x^{21} + x^{64} + x^{65} + x^{68} + \dots \quad (1.22-3)$$

The k -th fixed point is computed by moving all bits of the binary representation of k to position $2x$ where $x \geq 0$ is the index of the bit under consideration:

```

1  static inline ulong negbin_fixed_point(ulong k)
2  {
3      return bit_zip0(k);
4  }

```

The bit-zip function is given in section 1.15 on page 39. The sequence of radix -2 representations of 0, 1, 2, ..., interpreted as binary numbers, is entry A005351 in [312]:

0,1,6,7,4,5,26,27,24,25,30,31,28,29,18,19,16,17,22,23,20,21,106,107,104,105,110,111, ...

The corresponding sequence for the negative numbers $-1, -2, -3, \dots$ is entry A005352:

3,2,13,12,15,14,9,8,11,10,53,52,55,54,49,48,51,50,61,60,63,62,57,56,59,58,37,36,39,38, ...

More information about ‘non-standard’ representations of numbers can be found in [213].

1.22.3 Generating negbin words in order

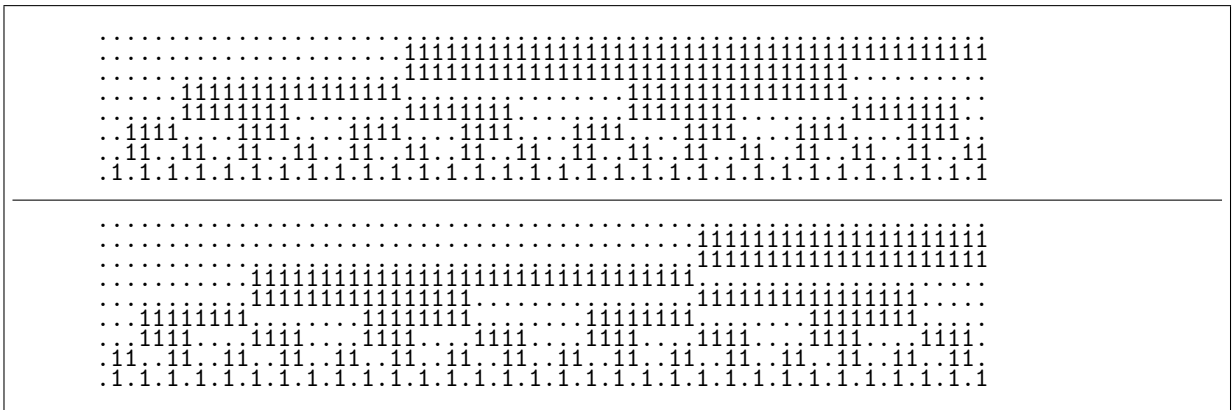


Figure 1.22-C: Radix -2 representations of the numbers $0 \dots +63$ (top) and $0 \dots -63$ (bottom).

A radix -2 representation can be incremented by the function [FXT: [bits/negbin.h](#)] (32-bit versions in what follows):

```

1  static inline ulong next_negbin(ulong x)
2  // With x the radix(-2) representation of n
3  // return radix(-2) representation of n+1.
4  {
5      const ulong m = 0xaaaaaaaaUL;
6      x ^= m;
7      ++x;
8      x ^= m;
9      return x;
10 }

```

A version without constants is

```

1      ulong s = x << 1;
2      ulong y = x ^ s;
3      y += 1;
4      s ^= y;
5      return s;

```

Decrementing can be done via

```

1  static inline ulong prev_negbin(ulong x)
2  // With x the radix(-2) representation of n
3  // return radix(-2) representation of n-1.
4  {
5      const ulong m = 0xaaaaaaaaUL;
6      x ^= m;
7      --x;
8      x ^= m;
9      return x;
10 }

```

or via

```

1  const ulong m = 0x55555555UL;
2  x ^= m;
3  ++x;
4  x ^= m;
5  return x;

```

The functions are quite fast, about 730 million words per second are generated (3 cycles per increment or decrement). Figure 1.22-C shows the generated words in forward (top) and backward (bottom) order. It was created with the program [FXT: [bits/negin2-demo.cc](http://bits.negin2-demo.cc)].

1.23 A sparse signed binary representation

0:P	0 =	
1:1P	1 =	+1
2:1.P.	2 =	+2
3:11P.M	3 =	+4 -1
4:1..P..	4 =	+4
5:1.1P.P	5 =	+4 +1
6:11.P.M.	6 =	+8 -2
7:111P..M	7 =	+8 -1
8:	...1...	...P...	8 =	+8
9:	...1..1	...P..P	9 =	+8 +1
10:	...1.1.	...P.P.	10 =	+8 +2
11:	...1.11	...P.M.M	11 =	+16 -4 -1
12:	...11..	...P.M..	12 =	+16 -4
13:	...11.1	...P.M.P	13 =	+16 -4 +1
14:	...111.	...P..M.	14 =	+16 -2
15:	...1111	...P...M	15 =	+16 -1
16:	..1....	..P....	16 =	+16
17:	..1...1	..P...P	17 =	+16 +1
18:	..1..1.	..P..P.	18 =	+16 +2
19:	..1..11	..P.P.M	19 =	+16 +4 -1
20:	..1.1..	..P.P..	20 =	+16 +4
21:	..1.1.1	..P.P.P	21 =	+16 +4 +1
22:	..1.11.	..P.M.M.	22 =	+32 -8 -2
23:	..1.111	..P.M..M	23 =	+32 -8 -1
24:	..11...	..P.M...	24 =	+32 -8
25:	..11..1	..P.M..P	25 =	+32 -8 +1
26:	..11.1.	..P.M.P.	26 =	+32 -8 +2
27:	..11.11	..P..M.M	27 =	+32 -4 -1
28:	..111..	..P..M..	28 =	+32 -4
29:	..111.1	..P..M.P	29 =	+32 -4 +1
30:	..1111.	..P...M.	30 =	+32 -2
31:	..11111	..P....M	31 =	+32 -1
32:	.1.....	.P.....	32 =	+32

Figure 1.23-A: Sparse signed binary representations (nonadjacent form, NAF). The symbols ‘P’ and ‘M’ are respectively used for +1 and -1, dots denote zeros.

0:P	0 =	
1:1P	1 =	+1
2:1.P.	2 =	+2
4:1..P..	4 =	+4
5:1.1P.P	5 =	+4 +1
8:	...1...	...P...	8 =	+8
9:	...1..1	...P..P	9 =	+8 +1
10:	...1.1.	...P.P.	10 =	+8 +2
16:	..1....	..P....	16 =	+16
17:	..1...1	..P...P	17 =	+16 +1
18:	..1..1.	..P..P.	18 =	+16 +2
20:	..1..11	..P.P.M	20 =	+16 +4
21:	..1.1..	..P.P..	21 =	+16 +4 +1
32:	..1.....	..P.....	32 =	+32
33:	..1...1	..P...P	33 =	+32 +1
34:	..1..1.	..P..P.	34 =	+32 +2
36:	..1..11	..P.P.M	36 =	+32 +4
37:	..1.1..	..P.P..	37 =	+32 +4 +1
40:	..1.1.1	..P.P.P	40 =	+32 +8
41:	..1.1.1	..P.P.P	41 =	+32 +8 +1
42:	..1.1.1	..P.P.P	42 =	+32 +8 +2
64:	.1.....	.P.....	64 =	+64

Figure 1.23-B: The numbers whose negative part in the NAF representation is zero.

An algorithm to compute a representation of a number x as

$$x = \sum_{k=0}^{\infty} s_k \cdot 2^k \quad \text{where } s_k \in \{-1, 0, +1\} \quad (1.23-1)$$

such that two consecutive digits s_k, s_{k+1} are never simultaneously nonzero is given in [275]. Figure 1.23-A gives the representation of several small numbers. It is the output of [FXT: [bits/bin2naf-demo.cc](#)].

We can convert the binary representation of x into a pair of binary numbers that correspond to the positive and negative digits [FXT: [bits/bin2naf.h](#)]:

```

1  static inline void bin2naf(ulong x, ulong &np, ulong &nm)
2  // Compute (nonadjacent form, NAF) signed binary representation of x:
3  // the unique representation of x as
4  // x=\sum_{k}{d_k*2^k} where d_j \in {-1,0,+1}
5  // and no two adjacent digits d_j, d_{j+1} are both nonzero.
6  // np has bits j set where d_j==+1
7  // nm has bits j set where d_j==-1
8  // We have: x = np - nm
9  {
10     ulong xh = x >> 1; // x/2
11     ulong x3 = x + xh; // 3*x/2
12     ulong c = xh ^ x3;
13     np = x3 & c;
14     nm = xh & c;
15 }
```

Converting back to binary is trivial:

```

1  static inline ulong naf2bin(ulong np, ulong nm) { return ( np - nm ); }
```

The representation is one example of a *nonadjacent form* (NAF). A method for the computation of certain nonadjacent forms (*w*-NAF) is given in [255]. A Gray code for the signed binary words is described in section 14.7 on page 315.

If a binary word contains no consecutive ones, then the negative part of the NAF representation is zero. The sequence of values is [0, 1, 2, 4, 5, 8, 9, 10, 16, ...], entry A003714 in [312], see figure 1.23-B. The numbers are called the *Fibbinary numbers*.

1.24 Generating bit combinations

1.24.1 Co-lexicographic (colex) order

Given a binary word with k bits set the following routine computes the binary word that is the next combination of k bits in co-lexicographic order. In the co-lexicographic order the reversed sets are sorted, see figure 1.24-A. The method to determine the successor is to determine the lowest block of ones and move its highest bit one position up. Then the rest of the block is moved to the low end of the word [FXT: [bits/bitcombclex.h](#)]:

```

1  static inline ulong next_colex_comb(ulong x)
2  {
3     ulong r = x & -x; // lowest set bit
4     x += r; // replace lowest block by a one left to it
5
6     if ( 0==x ) return 0; // input was last combination
7
8     ulong z = x & -x; // first zero beyond lowest block
9     z -= r; // lowest block (cf. lowest_block())
10
11     while ( 0==(z&1) ) { z >>= 1; } // move block to low end of word
12     return x | (z>>1); // need one bit less of low block
13 }
```

One could replace the while-loop by a bit scan and shift combination. The combinations $\binom{32}{20}$ are generated at a rate of about 142 million per second. The rate is about 120 M/s for the combinations $\binom{32}{12}$, the rate with $\binom{60}{7}$ is 70 M/s, and with $\binom{60}{53}$ it is 160 M/s.

	word	=	set	=	set (reversed)
1:	...111	=	{ 0, 1, 2 }	=	{ 2, 1, 0 }
2:	..1.11	=	{ 0, 1, 3 }	=	{ 3, 1, 0 }
3:	..11.1	=	{ 0, 2, 3 }	=	{ 3, 2, 0 }
4:	..111.	=	{ 1, 2, 3 }	=	{ 3, 2, 1 }
5:	.1..11	=	{ 0, 1, 4 }	=	{ 4, 1, 0 }
6:	.1.1.1	=	{ 0, 2, 4 }	=	{ 4, 2, 0 }
7:	.1.11.	=	{ 1, 2, 4 }	=	{ 4, 2, 1 }
8:	.11..1	=	{ 0, 3, 4 }	=	{ 4, 3, 0 }
9:	.11.1.	=	{ 1, 3, 4 }	=	{ 4, 3, 1 }
10:	.111..	=	{ 2, 3, 4 }	=	{ 4, 3, 2 }
11:	1...11	=	{ 0, 1, 5 }	=	{ 5, 1, 0 }
12:	1..1.1	=	{ 0, 2, 5 }	=	{ 5, 2, 0 }
13:	1..11.	=	{ 1, 2, 5 }	=	{ 5, 2, 1 }
14:	1.1..1	=	{ 0, 3, 5 }	=	{ 5, 3, 0 }
15:	1.1.1.	=	{ 1, 3, 5 }	=	{ 5, 3, 1 }
16:	1.11..	=	{ 2, 3, 5 }	=	{ 5, 3, 2 }
17:	11...1	=	{ 0, 4, 5 }	=	{ 5, 4, 0 }
18:	11..1.	=	{ 1, 4, 5 }	=	{ 5, 4, 1 }
19:	11.1..	=	{ 2, 4, 5 }	=	{ 5, 4, 2 }
20:	111...	=	{ 3, 4, 5 }	=	{ 5, 4, 3 }

Figure 1.24-A: Combinations $\binom{6}{3}$ in co-lexicographic order. The reversed sets are sorted.

A variant of the method which involves a division appears in [39, item 175]. The routine given here is due to Doug Moore and Glenn Rhoads.

The following routine computes the predecessor of a combination:

```

1 static inline ulong prev_colex_comb(ulong x)
2 // Inverse of next_colex_comb()
3 {
4     x = next_colex_comb( ~x );
5     if ( 0!=x ) x = ~x;
6     return x;
7 }
```

The first and last combination can be computed via

```

1 static inline ulong first_comb(ulong k)
2 // Return the first combination of (i.e. smallest word with) k bits,
3 // i.e. 00..001111..1 (k low bits set)
4 // Must have: 0 <= k <= BITS_PER_LONG
5 {
6     ulong t = ~0UL >> ( BITS_PER_LONG - k );
7     if ( k==0 ) t = 0; // shift with BITS_PER_LONG is undefined
8     return t;
9 }
```

and

```

1 static inline ulong last_comb(ulong k, ulong n=BITS_PER_LONG)
2 // return the last combination of (biggest n-bit word with) k bits
3 // i.e. 1111..100..00 (k high bits set)
4 // Must have: 0 <= k <= n <= BITS_PER_LONG
5 {
6     return first_comb(k) << ( n - k );
7 }
```

The if-statement in `first_comb()` is needed because a shift by more than `BITS_PER_LONG-1` is undefined by the C-standard, see section 1.1.5 on page 4.

The listing in figure 1.24-A can be created with the program [FXT: bits/bitcombcplex-demo.cc]:

```

1     ulong n = 6, k = 3;
2     ulong last = last_comb(k, n);
3     ulong g = first_comb(k);
4     ulong gg = 0;
5     do
6     {
7         // visit combination given as word g
8         gg = g;
```

```

9      g = next_colex_comb(g);
10     }
11     while ( gg!=last );

```

1.24.2 Lexicographic (lex) order

lex (5, 3)		colex (5, 2)	
word	= set	word	= set
1: ..111	= { 0, 1, 2 }	...11	= { 0, 1 }
2: .1.11	= { 0, 1, 3 }	..1.1	= { 0, 2 }
3: 1..11	= { 0, 1, 4 }	..11.	= { 1, 2 }
4: .11.1	= { 0, 2, 3 }	.1..1	= { 0, 3 }
5: 1.1.1	= { 0, 2, 4 }	.1.1.	= { 1, 3 }
6: 11..1	= { 0, 3, 4 }	.11..	= { 2, 3 }
7: .111.	= { 1, 2, 3 }	1...1	= { 0, 4 }
8: 1.11.	= { 1, 2, 4 }	1..1.	= { 1, 4 }
9: 11.1.	= { 1, 3, 4 }	1.1..	= { 2, 4 }
10: 111..	= { 2, 3, 4 }	11...	= { 3, 4 }

Figure 1.24-B: Combinations $\binom{5}{3}$ in lexicographic order (left). The sets are sorted. The binary words with lex order are the bit-reversed complements of the words with colex order (right).

The binary words corresponding to combinations $\binom{n}{k}$ in lexicographic order are the bit-reversed complements of the words for the combinations $\binom{n}{n-k}$ in co-lexicographic order, see figure 1.24-B. A more precise term for the order is subset-lex (for sets written with elements in increasing order). The sequence is identical to the delta-set-colex order backwards.

The program [FXT: [bits/bitcomplex-demo.cc](#)] shows how to compute the subset-lex sequence efficiently:

```

1      ulong n = 5, k = 3;
2      ulong x = first_comb(n-k);      // first colex (n-k choose n)
3      const ulong m = first_comb(n);  // aux mask
4      const ulong l = last_comb(k, n); // last colex
5      ulong ct = 0;
6      ulong y;
7      do
8      {
9          y = revbin(~x, n) & m; // lex order
10         // visit combination given as word y
11         x = next_colex_comb(x);
12     }
13     while ( y != l );

```

The bit-reversal routine `revbin()` is shown in section 1.14 on page 33. Sections 6.2.1 on page 177 and section 6.2.2 give iterative algorithms for combinations (represented by arrays) in lex and colex order, respectively.

1.24.3 Shifts-order

1: 1....	1: 11...	1: 111..	1: 1111.
2: .1...	2: .11..	2: .111.	2: .1111
3: ..1..	3: ..11.	3: ..111	3: .111.1
4: ...1.	4: ...11	4: .11.1	4: .11.11
5:1	5: 1.1..	5: .11.1	5: 1.111
	6: .1.1.	6: .11..1	
	7: ..1.1	7: 1.11.	
	8: 1..1.	8: .1.11	
	9: .1..1	9: 1.1.1	
	10: 1...1	10: 1..11	

Figure 1.24-C: Combinations $\binom{5}{k}$, for $k = 1, 2, 3, 4$ in shifts-order.

Figure 1.24-C shows combinations in *shifts-order*. The order for combinations $\binom{n}{k}$ is obtained from the shifts-order for subsets (section 8.4 on page 208) by discarding all subsets whose number of elements are $\neq k$ and reversing the list order. The first combination is $[1^k 0^{n-k}]$ and the successor is computed as follows (see figure 1.24-D):

1:	1111...		18:	.11..11	
2:	.1111..		19:	11..1.1	< S
3:	..1111.		20:	11...11	< S-2
4:	...1111		21:	1.111..	< S-2
5:	111.1..	< S	22:	.1.111.	
6:	.111.1.		23:	..1.111	
7:	..111.1		24:	1.11.1.	< S
8:	111..1.	< S	25:	.1.11.1	
9:	.111..1		26:	1.11..1	< S
10:	111...1	< S	27:	1.1.11.	< S-2
11:	11.11..	< S-2	28:	.1.1.11	
12:	.11.11.		29:	1.1.1.1	< S
13:	..11.11		30:	1.1..11	< S-2
14:	11.1.1.	< S	31:	1..111.	< S-2
15:	.11.1.1		32:	.1..111	
16:	11.1..1	< S	33:	1..11.1	< S
17:	11..1.1	< S-2	34:	1..1.11	< S-2
18:	.11..11		35:	1...111	< S-2

Figure 1.24-D: Updates with combinations $\binom{7}{4}$: simple split ‘S’, split second ‘S-2’, easy case unmarked.

1. Easy case: if the rightmost one is not in position zero (least significant bit), then shift the word to the right and return the combination.
2. Finished?: if the combination is the last one ($[0^n]$, $[0^{n-1}1]$, $[10^{n-k}1^{k-1}]$), then return zero.
3. Shift back: shift the word to the left such that the leftmost one is in the leftmost position (this can be a no-op).
4. Simple split: if the rightmost one is not the least significant bit, then move it one position to the right and return the combination.
5. Split second block: move the rightmost bit of the second block (from the right) of ones one position to the right and attach the lowest block of ones and return the combination.

An implementation is given in [FXT: [bits/bitcombshifts.h](#)]:

```

1  class bit_comb_shifts
2  {
3  public:
4      ulong x_; // the combination
5      ulong s_; // how far shifted to the right
6      ulong n_, k_; // combinations (n choose k)
7      ulong last_; // last combination
8
9  public:
10     bit_comb_shifts(ulong n, ulong k)
11     {
12         n_ = n; k_ = k;
13         first();
14     }
15
16     ulong first(ulong n, ulong k)
17     {
18         s_ = 0;
19         x_ = last_comb(k, n);
20
21         if ( k>1 ) last_ = first_comb(k-1) | (1UL<<(n-1)); // [10000111]
22         else      last_ = k; // [000001] or [000000]
23
24         return x_;
25     }
26
27     ulong first() { return first(n_, k_); }
28
29     ulong next()
30     {
31         if ( 0==(x_&1) ) // easy case:
32         {
33             ++s_;
34             x_ >>= 1;
35             return x_;
36         }
37         else // splitting cases:

```



```

38     {
39         if ( x_ == last_ ) return 0; // combination was last
40
41         x_ <<= s_; s_ = 0; // shift back to the left
42         ulong b = x_ & -x_; // lowest bit
43
44
45         if ( b!=1UL ) // simple split
46         {
47             x_ -= (b>>1); // move rightmost bit to the right
48             return x_;
49         }
50         else // split second block and attach first
51         {
52             ulong t = low_ones(x_); // block of ones at lower end
53             x_ ^= t; // remove block
54             ulong b2 = x_ & -x_; // (second) lowest bit
55
56             b2 >>= 1;
57             x_ -= b2; // move bit to the right
58
59             // attach block:
60             do { t<<=1; } while ( 0==(t&x_) );
61             x_ |= (t>>1);
62             return x_;
63         }
64     }
65 }
66 };

```

The combinations $\binom{32}{20}$ are generated at a rate of about 150 M/s, for the combinations $\binom{32}{12}$ the rate is about 220 M/s [FXT: bits/bitcombshifts-demo.cc]. The rate with the combinations $\binom{60}{7}$ is 415 M/s and with $\binom{60}{53}$ it is 110 M/s. The generation is very fast for the sparse case.

1.24.4 Minimal-change order †

The following routine is due to Doug Moore [FXT: bits/bitcombminchange.h]:

```

1  static inline ulong igc_next_minchange_comb(ulong x)
2  // Return the inverse Gray code of the next combination in minimal-change order.
3  // Input must be the inverse Gray code of the current combination.
4  {
5      ulong g = rev_gray_code(x);
6      ulong i = 2;
7      ulong cb; // ==candidate bits;
8      do
9      {
10         ulong y = (x & ~(i-1)) + i;
11         ulong j = lowest_one(y) << 1;
12         ulong h = !(y & j);
13         cb = ((j-h) ^ g) & (j-i);
14         i = j;
15     }
16     while ( 0==cb );
17     return x + lowest_one(cb);
18 }
19 }

```

It can be used as suggested by the routine

```

1  static inline ulong next_minchange_comb(ulong x, ulong last)
2  // Not efficient, just to explain the usage of igc_next_minchange_comb()
3  // Must have: last==igc_last_comb(k, n)
4  {
5      x = inverse_gray_code(x);
6      if ( x==last ) return 0;
7      x = igc_next_minchange_comb(x);
8      return gray_code(x);
9  }

```

The auxiliary function `igc_last_comb()` is (32-bit version only)

```

1  static inline ulong igc_last_comb(ulong k, ulong n)
2  // Return the (inverse Gray code of the) last combination

```

```

3 // as in igc_next_minchange_comb()
4 {
5     if ( 0==k ) return 0;
6
7     const ulong f = 0xaaaaaaaaUL >> (BITS_PER_LONG-k); // == first_sequency(k);
8     const ulong c = ~0UL >> (BITS_PER_LONG-n); // == first_comb(n);
9     return c ^ (f>>1);
10    // ^= (by Doug Moore)
11    // return ((1UL<<n) - 1) ^ (((1UL<<k) - 1) / 3);
12 }

```

Successive combinations differ in exactly two positions. For example, with $n = 5$ and $k = 3$:

x	inverse_gray_code(x)
..111	..1.1 == first_sequency(k)
.11.1	.1.1.1
.111.	.1.111
1.11.	1.11.1
1.1.1	1.111.
111.	1.1111
111.1	1.11111
111.11	111.11
111.111	1111.1
111.1111	11111.
111.11111	111111 == igc_last_comb(k, n)

The same run of bit combinations would be generated by going through the Gray codes and omitting all words where the bit-count is not equal to k . The algorithm shown here is much more efficient.

For greater efficiency one may prefer code which avoids the repeated computation of the inverse Gray code, for example:

```

1     ulong last = igc_last_comb(k, n);
2     ulong c, nc = first_sequency(k);
3     do
4     {
5         c = nc;
6         nc = igc_next_minchange_comb(c);
7         ulong g = gray_code(c);
8         // Here g contains the bit-combination
9     }
10    while ( c!=last );

```

n = 6 k = 2	n = 6 k = 3	n = 6 k = 4
...11	...111	...1111
...11.	...11.1	...11.11
...1.1	...11.11	...1.1.11
...11.	...1.111	...1.11.11
...1.1.	...1.1.111	...1.111.11
...1.11.	...1.11.111	...1.1111.11
...1.111.	...1.111.111	...1.11111.11
...111.	...1.1111.111	...1.111111.11
...111.	...1.11111.111	...1.1111111.11
...1111.	...11.111	...1.1111.111
...1111.	...11.11.11	...1.1111.1111
...1111.	...11.11.111	...1.1111.11111
...1111.	...11.111.11	...1.1111.111111
...1111.	...11.1111.11	...1.1111.1111111
...1111.	...11.11111.11	...1.1111.11111111
...11111.	...111.111	...1.1111.11111.11
...11111.	...111.11.11	...1.1111.11111.111
...11111.	...111.11.111	...1.1111.11111.1111
...11111.	...111.111.11	...1.1111.11111.11111
...11111.	...111.1111.11	...1.1111.11111.111111
...11111.	...111.11111.11	...1.1111.11111.1111111
...111111.	...1111.111	...1.1111.11111.11111.11
...111111.	...1111.11.11	...1.1111.11111.11111.111
...111111.	...1111.11.111	...1.1111.11111.11111.1111
...111111.	...1111.111.11	...1.1111.11111.11111.11111
...111111.	...1111.1111.11	...1.1111.11111.11111.111111
...111111.	...1111.11111.11	...1.1111.11111.11111.1111111
...1111111.	...11111.111	...1.1111.11111.11111.11111.11
...1111111.	...11111.11.11	...1.1111.11111.11111.11111.111
...1111111.	...11111.11.111	...1.1111.11111.11111.11111.1111
...1111111.	...11111.111.11	...1.1111.11111.11111.11111.11111
...1111111.	...11111.1111.11	...1.1111.11111.11111.11111.111111
...1111111.	...11111.11111.11	...1.1111.11111.11111.11111.1111111

Figure 1.24-E: Minimal-change combinations, their inverse Gray codes, and the differences of the inverse Gray codes. The differences are powers of 2.

The difference of the inverse Gray codes of two successive combinations is always a power of 2, see figure 1.24-E (the listings were created with the program [FXT: bits/bitcombminchange-demo.cc]). With this observation we can derive a different version that checks the pattern of the change:

```

1 static inline ulong igc_next_minchange_comb(ulong x)
2 // Alternative version.
3 {
4     ulong gx = gray_code( x );
5     ulong i = 2;
6     do
7     {

```

```

8     ulong y = x + i;
9     i <<= 1;
10    ulong gy = gray_code( y );
11    ulong r = gx ^ gy;
12
13    // Check that change consists of exactly one bit
14    // of the new and one bit of the old pattern:
15    if ( is_pow_of_2( r & gy ) && is_pow_of_2( r & gx ) ) break;
16    // is_pow_of_2(x):=((x & -x) == x) returns 1 also for x==0.
17    // But this cannot happen for both tests at the same time
18    }
19    while ( 1 );
20    return y;
21 }

```

This version is the fastest: the combinations $\binom{32}{12}$ are generated at a rate of about 96 million per second, the combinations $\binom{32}{20}$ at a rate of about 83 million per second.

Here is another version which needs the number of set bits as a second parameter:

```

1  static inline ulong igc_next_minchange_comb(ulong x, ulong k)
2  // Alternative version, uses the fact that the difference
3  // of two successive x is the smallest possible power of 2.
4  {
5      ulong y, i = 2;
6      do
7      {
8          y = x + i;
9          i <<= 1;
10     }
11     while ( bit_count( gray_code(y) ) != k );
12     return y;
13 }

```

The routine will be fast if the CPU has a bit-count instruction. The necessary modification for the generation of the previous combination is trivial:

```

1  static inline ulong igc_prev_minchange_comb(ulong x, ulong k)
2  // Returns the inverse graycode of the previous combination in minimal-change order.
3  // Input must be the inverse graycode of the current combination.
4  // With input==first the output is the last for n=BITS_PER_LONG
5  {
6      ulong y, i = 2;
7      do
8      {
9          y = x - i;
10         i <<= 1;
11     }
12     while ( bit_count( gray_code(y) ) != k );
13     return y;
14 }

```

1.25 Generating bit subsets of a given word

1.25.1 Counting order

To generate all subsets of the set of ones of a binary word we use the sparse counting idea shown in section 1.8.1 on page 20. The implementation is [FXT: class `bit_subset` in `bits/bit_subset.h`]:

```

1  class bit_subset
2  {
3  public:
4      ulong u_; // current subset
5      ulong v_; // the full set
6
7  public:
8      bit_subset(ulong v) : u_(0), v_(v) { ; }
9      ~bit_subset() { ; }
10     ulong current() const { return u_; }
11     ulong next() { u_ = (u_ - v_) & v_; return u_; }
12     ulong prev() { u_ = (u_ - 1) & v_; return u_; }
13 }

```



```

4     ulong u1 = S_.current();
5     if ( u1 == 0 ) return last();
6     ulong u0 = S_.prev();
7     ulong x = ~u0 & u1;
8     g_ ^= x;
9     return g_;
10  }
11
12  ulong last(ulong v) { S_.last(v); h_=highest_one(v); g_=h_; return g_; }
13  ulong last() { S_.last(); g_=h_; return g_; }
14 };

```

About 365 million subsets per second are generated with both `next()` and `prev()`.

1.26 Binary words in lexicographic order for subsets

1.26.1 Next and previous word in lexicographic order

1:	1...	= 8	{0}
2:	11..	= 12	{0, 1}
3:	111.	= 14	{0, 1, 2}
4:	1111	= 15	{0, 1, 2, 3}
5:	11.1	= 13	{0, 1, 3}
6:	1.1.	= 10	{0, 2}
7:	1.11	= 11	{0, 2, 3}
8:	1..1	= 9	{0, 3}
9:	.1..	= 4	{1}
10:	.11.	= 6	{1, 2}
11:	.111	= 7	{1, 2, 3}
12:	.1.1	= 5	{1, 3}
13:	..1.	= 2	{2}
14:	..11	= 3	{2, 3}
15:	...1	= 1	{3}

Figure 1.26-A: Binary words corresponding to nonempty subsets of the 4-element set in lexicographic order with respect to subsets. Note the first element of the subsets corresponds to the highest set bit.

[0:	= 0 *	16:	.1...1	= 17	32:	1....1	= 33	48:	11..11	= 51
1:1	= 1 *	17:	.1..11	= 19	33:	1...11	= 35	49:	11..1.	= 50
2:11	= 3	18:	.1..1.	= 18 *	34:	1...1.	= 34 *	50:	11.1.1	= 53
3:	...1.	= 2	19:	.1.1.1	= 21	35:	1..1.1	= 37	51:	11.111	= 55
4:	...1.1	= 5	20:	.1.111	= 23	36:	1..111	= 39	52:	11.11.	= 54
5:	...111	= 7	21:	.1.11.	= 22	37:	1..11.	= 38	53:	11.1..	= 52
6:	...11.	= 6 *	22:	.1.1..	= 20	38:	1..1..	= 36	54:	111..1	= 57
7:	..1..	= 4	23:	.11..1	= 25	39:	1.1..1	= 41	55:	111.11	= 59
8:	..1.1	= 9	24:	.11.11	= 27	40:	1.1.11	= 43	56:	111.1.	= 58
9:	..1.11	= 11	25:	.11.1.	= 26	41:	1.1.1.	= 42	57:	1111.1	= 61
10:	..1.1.	= 10 *	26:	.111.1	= 29	42:	1.11.1	= 45	58:	111111	= 63
11:	..11.1	= 13	27:	.11111	= 31	43:	1.1111	= 47	59:	11111.	= 62
12:	..1111	= 15	28:	.1111.	= 30	44:	1.111.	= 46	60:	1111..	= 60 *
13:	..111.	= 14	29:	.111..	= 28	45:	1.11..	= 44	61:	111...	= 56
14:	..11..	= 12	30:	.11...	= 24	46:	1.1...	= 40	62:	11....	= 48
15:	..1...	= 8	31:	.1....	= 16	47:	11...1	= 49	63:	1.....	= 32

Figure 1.26-B: Binary words corresponding to the subsets of the 6-element set, as generated by `prev_lexrev()`. Fixed points are marked with asterisk.

The (bit-reversed) binary words in lexicographic order with respect to the subsets shown in figure 1.26-A can be generated by successive calls to the following function [FXT: [bits/bitlex.h](#)]:

```

1  static inline ulong next_lexrev(ulong x)
2  // Return next word in subset-lex order.
3  {
4     ulong x0 = x & -x; // lowest bit
5     if ( 1!=x0 ) // easy case: set bit right of lowest bit
6     {
7         x0 >>= 1;
8         x ^= x0;
9         return x;

```

```

10     }
11     else // lowest bit at word end
12     {
13         x ^= x0; // clear lowest bit
14         x0 = x & -x; // new lowest bit ...
15         x0 >>= 1; x -= x0; // ... is moved one to the right
16         return x;
17     }
18 }

```

The bit-reversed representation was chosen because the isolation of the lowest bit is often cheaper than the same operation on the highest bit. Starting with a one-bit word at position $n - 1$, we generate the 2^n subsets of the word of n ones. The function is used as follows [FXT: bits/bitlex-demo.cc]:

```

ulong n = 4; // n-bit binary words
ulong x = 1UL<<(n-1); // first subset
do
{
    // visit word x
}
while ( (x=next_lexrev(x)) );

```

The following function goes backward:

```

1  static inline ulong prev_lexrev(ulong x)
2  // Return previous word in subset-lex order.
3  {
4      ulong x0 = x & -x; // lowest bit
5      if ( x & (x0<<1) ) // easy case: next higher bit is set
6      {
7          x ^= x0; // clear lowest bit
8          return x;
9      }
10     else
11     {
12         x += x0; // move lowest bit to the left
13         x |= 1; // set rightmost bit
14         return x;
15     }
16 }

```

The sequence of all n -bit words is generated by 2^n calls to `prev_lexrev()`, starting with zero. The words corresponding to subsets of the 6-element set are shown in figure 1.26-B. The sequence [1, 3, 2, 5, 7, 6, 4, 9, ...] in the right column is entry A108918 in [312].

The rate of generation using `next()` is about 274 million per second and about 253 million per second with `prev()`. An equivalent routine for arrays is given in section 8.1.2 on page 203. The routines are useful for a special version of fast Walsh transforms described in section 23.5.3 on page 472.

1.26.2 Conversion between binary and lex-ordered words

A little contemplation on the structure of the binary words in lexicographic order leads to the routine that allows random access to the k -th lex-rev word (unrank algorithm) [FXT: bits/bitlex.h]:

```

1  static inline ulong negidx2lexrev(ulong k)
2  {
3      ulong z = 0;
4      ulong h = highest_one(k);
5      while ( k )
6      {
7          while ( 0==(h&k) ) h >>= 1;
8          z ^= h;
9          ++k;
10         k &= h - 1;
11     }
12     return z;
13 }

```

Let the inverse function be $T(x)$, then we have $T(0) = 0$ and, with $h(x)$ being the highest power of 2 not greater than x ,

$$T(x) = h(x) - 1 + \begin{cases} T(x - h(x)) & \text{if } x - h(x) \neq 0 \\ h(x) & \text{otherwise} \end{cases} \quad (1.26-1)$$

The ranking algorithm starts with the lowest bit:

```

1  static inline ulong lexrev2negidx(ulong x)
2  {
3      if ( 0==x ) return 0;
4      ulong h = x & -x; // lowest bit
5      ulong r = (h-1);
6      while ( x^=h )
7      {
8          r += (h-1);
9          h = x & -x; // next higher bit
10     }
11     r += h; // highest bit
12     return r;
13 }

```

1.26.3 Minimal decompositions into terms $2^k - 1$ ‡

...1 1	...1 = 1	= 1	
...11 2	...1. = 2	= 1 + 1	
...1. 1	...11 = 3	= 3	
..1.1 2	..1.. = 4	= 3 + 1	
..111 3	..1.1 = 5	= 3 + 1 + 1	
..11. 2	..11. = 6	= 3 + 3	
..1.. 1	..111 = 7	= 7	
.1..1 2	.1... = 8	= 7 + 1	
.1.11 3	.1..1 = 9	= 7 + 1 + 1	
.1.1. 2	.1.1. = 10	= 7 + 3	
.11.1 3	.1.11 = 11	= 7 + 3 + 1	
.1111 4	.11.. = 12	= 7 + 3 + 1 + 1	
.111. 3	.11.1 = 13	= 7 + 3 + 3	
.11.. 2	.111. = 14	= 7 + 7	
.1... 1	.1111 = 15	= 15	
1...1 2	1.... = 16	= 15 + 1	
1..11 3	1...1 = 17	= 15 + 1 + 1	
1..1. 2	1..1. = 18	= 15 + 3	
1.1.1 3	1..11 = 19	= 15 + 3 + 1	
1.111 4	1.1.. = 20	= 15 + 3 + 1 + 1	
1.11. 3	1.1.1 = 21	= 15 + 3 + 3	
1.1.. 2	1.11. = 22	= 15 + 7	
11..1 3	1.111 = 23	= 15 + 7 + 1	
11.11 4	11... = 24	= 15 + 7 + 1 + 1	
11.1. 3	11..1 = 25	= 15 + 7 + 3	
111.1 4	11.1. = 26	= 15 + 7 + 3 + 1	
11111 5	11.11 = 27	= 15 + 7 + 3 + 1 + 1	
1111. 4	111.. = 28	= 15 + 7 + 3 + 3	
111.. 3	111.1 = 29	= 15 + 7 + 7	
11... 2	1111. = 30	= 15 + 15	
1.... 1	11111 = 31	= 31	

Figure 1.26-C: Binary words in subset-lex order and their bit counts (left columns). The least number of terms of the form $2^k - 1$ needed in the sum $x = \sum_k (2^k - 1)$ (right columns) equals the bit count.

The least number of terms needed in the sum $x = \sum_k (2^k - 1)$ equals the number of bits of the lex-word as shown in figure 1.26-C. The number can be computed as

```
c = bit_count( negidx2lexrev( x ) );
```

Alternatively, we can subtract the greatest integer of the form $2^k - 1$ until x is zero and count the number of subtractions. The sequence of these numbers is entry [A100661](#) in [312]:

1, 2, 1, 2, 3, 2, 1, 2, 3, 2, 3, 4, 3, 2, 1, 2, 3, 2, 3, 4, 3, 2, 3, 4, 3, 2, 3, 4, 3, 4, 5, 4, 3, 2, 1, 2, 3, 2, 3, ...

The following function can be used to compute the sequence:

```

1  void S(ulong f, ulong n) // A100661
2  {
3      static int s = 0;
4      ++s;
5      cout << s << ", ";
6      for (ulong m=1; m<n; m<<=1) S(f+m, m);
7      --s;
8      cout << s << ", ";
9  }

```

If called with arguments $f = 0$ and $n = 2^k$, it prints the first $2^{k+1} - 1$ numbers of the sequence followed by a zero. A generating function of the sequence is given by

$$Z(x) := \frac{-1 + 2(1-x) \prod_{n=1}^{\infty} (1+x^{2^n-1})}{(1-x)^2} = \quad (1.26-2)$$

$$1 + 2x + x^2 + 2x^3 + 3x^4 + 2x^5 + x^6 + 2x^7 + 3x^8 + 2x^9 + 3x^{10} + 4x^{11} + 3x^{12} + 2x^{13} + \dots$$

1.26.4 The sequence of fixed points †

0:	514:	.1.....1.
1:1	540:	.1....111..
6:11.	556:	.1...1.11..
10:1.1.	[--snip--]	
18:1..1.	1556:	.11....1.1..
34:1...1.	1572:	.11...1..1..
60:1111..	1604:	.11..1...1..
66:1...1.	1668:	.11.1...1..
92:1.111..	1796:	.111....1..
108:11.11..	2040:	.11111111..
116:111.1..	2050:	1.....1.
130:1...1.	2076:	1.....111..
156:1..111..	2092:	1.....1.11..
172:1.1.11..	2100:	1.....11.1..
180:1.11.1..	2124:	1.....1..11..
204:11..11..	2132:	1.....1.1.1..
212:11.1.1..	2148:	1.....11..1..
228:111..1..	[--snip--]	
258:	..1.....1.	4644:	1..1...1..1..
284:	..1...111..	4676:	1..1..1...1..
300:	..1..1.11..	4740:	1..1.1...1..
308:	..1..11.1..	4868:	1..11....1..
332:	..1.1..11..	5112:	1..1111111..
340:	..1.1.1.1..	5132:	1.1.....11..
356:	..1.11..1..	5140:	1.1.....1.1..
396:	..11...11..	5156:	1.1....1..1..
404:	..11..1.1..	5188:	1.1...1...1..
420:	..11.1..1..	5252:	1.1..1....1..
452:	..111...1..	5380:	1.1.i.....1..

Figure 1.26-D: Fixed points of the binary to lex-rev conversion.

The sequence of fixed points of the conversion to and from indices starts as

0, 1, 6, 10, 18, 34, 60, 66, 92, 108, 116, 130, 156, 172, 180, 204, 212, 228, 258, 284, 300, 308, 332, 340, 356, 396, 404, 420, 452, 514, 540, 556, ...

This sequence is entry [A079471](#) in [312]. The values as bit patterns are shown in figure 1.26-D. The crucial observation is that a word is a fixed point if it equals zero or its bit-count equals 2^j where j is the index of the lowest set bit.

Now we can find out whether x is a fixed point of the sequence by the following function:

```

1  static inline bool is_lexrev_fixed_point(ulong x)
2  // Return whether x is a fixed point in the prev_lexrev() - sequence
3  {
4      if ( x & 1 )
5      {
6          if ( 1==x ) return true;
7          else       return false;
8      }
9      else
10     {
11         ulong w = bit_count(x);
12         if ( w != (w & -w) ) return false;
13         if ( 0==x ) return true;
14         return 0 != ( (x & -x) & w );
15     }
16 }
```

Alternatively, use either of the following tests:

```

x == negidx2lexrev(x)
x == lexrev2negidx(x)
```


1.26.5 Recursive generation and relation to a power series ‡

```

Start: 1
Rules:
  0 --> 0
  1 --> 110
-----
0:    (#=2)
1
1:    (#=4)
110
2:    (#=8)
1101100
3:    (#=16)
110110011011000
4:    (#=32)
1101100110110001101100110110000
5:    (#=64)
110110011011000110110011011000011011001101100011011001101100000

```

Figure 1.26-E: String substitution with rules $\{0 \rightarrow 0, 1 \mapsto 110\}$.

The following function generates the bit-reversed binary words in reversed lexicographic order:

```

1 void C(ulong f, ulong n, ulong w)
2 {
3     for (ulong m=1; m<n; m<<=1) C(f+m, m, w^m);
4     print_bin(" ", w, 10); // visit
5 }

```

By calling `C(0, 64, 0)` we generate the list of words shown in figure 1.26-B with the all-zeros word moved to the last position. A slight modification of the function

```

1 void A(ulong f, ulong n)
2 {
3     cout << "1,";
4     for (ulong m=1; m<n; m<<=1) A(f+m, m);
5     cout << "0,";
6 }

```

generates the power series (sequence [A079559](#) in [312])

$$\prod_{n=1}^{\infty} (1 + x^{2^n - 1}) = 1 + x + x^3 + x^4 + x^7 + x^8 + x^{10} + x^{11} + x^{15} + x^{16} + \dots \quad (1.26-3)$$

By calling `A(0, 32)` we generate the sequence

1, 1, 0, 1, 1, 0, 0, 1, 1, 0, 1, 1, 0, 0, 0, 1, 1, 0, 1, 1, 0, 0, 1, 1, 0, 1, 1, 0, 0, 1, 1, 0, 0, 0, 0, ...

Indeed, the lowest bit of the k -th word of the bit-reversed sequence in reversed lexicographic order equals the $(k-1)$ -st coefficient in the power series. The sequence can also be generated by the string substitution shown in figure 1.26-E.

The sequence of sums, prepended by 1,

$$1 + x \frac{\prod_{n=1}^{\infty} (1 + x^{2^n - 1})}{1 - x} = 1 + 1x + 2x^2 + 2x^3 + 3x^4 + 4x^5 + 4x^6 + \dots \quad (1.26-4)$$

has series coefficients

1, 1, 2, 2, 3, 4, 4, 4, 5, 6, 6, 7, 8, 8, 8, 8, 9, 10, 10, 11, 12, 12, 12, 13, ...

This sequence is entry [A046699](#) in [312]. We have $a(1) = a(2) = 1$ and the sequence satisfies the peculiar recurrence

$$a(n) = a(n - a(n - 1)) + a(n - 1 - a(n - 2)) \quad \text{for } n > 2 \quad (1.26-5)$$

1.27 Fibonacci words ‡

A Fibonacci word is a word that does not contain two successive ones. Whether a given binary word is a Fibonacci word can be tested with the function [FXT: [bits/fibrep.h](#)]

```

1 static inline bool is_fibrep(ulong f)
2 {
3     return ( 0==(f&(f>>1)) );
4 }

```

The following functions convert between the binary and the Fibonacci representation:

```

1 static inline ulong bin2fibrep(ulong b)
2 // Return Fibonacci representation of b
3 // Limitation: the first Fibonacci number greater
4 // than b must be representable as ulong.
5 // 32 bit: b < 2971215073=F(47) [F(48)=4807526976 > 2^32]
6 // 64 bit: b < 12200160415121876738=F(93) [F(94) > 2^64]
7 {
8     ulong f0=1, f1=1, s=1;
9     while ( f1<=b ) { ulong t = f0+f1; f0=f1; f1=t; s<<=1; }
10    ulong f = 0;
11    while ( b )
12    {
13        s >>= 1;
14        if ( b>=f0 ) { b -= f0; f ^=s; }
15        { ulong t = f1-f0; f1=f0; f0=t; }
16    }
17    return f;
18 }

1 static inline ulong fibrep2bin(ulong f)
2 // Return binary representation of f
3 // Inverse of bin2fibrep().
4 {
5     ulong f0=1, f1=1;
6     ulong b = 0;
7     while ( f )
8     {
9         if ( f&1 ) b += f1;
10        { ulong t=f0+f1; f0=f1; f1=t; }
11        f >>= 1;
12    }
13    return b;
14 }

```

1.27.1 Lexicographic order

0:1	11: ...1.1.1..	22: .1.....1	33: .1.1.1.1.1	44: 1..1..1.1
1:1	12: ...1.1.1.1	23: .1.....1	34: 1.....1	45: 1..1.1.1..
2:1.1	13: ..1.....1	24: .1...1..1	35: 1.....1	46: 1..1.1.1.1
3:1.1.	14: ..1.....1	25: .1...1.1.1	36: 1.....1.	47: 1..1.....1
4:1.1.1	15: ..1.....1.	26: .1...1...1	37: 1.....1..	48: 1..1.....1
5:1...1	16: ..1...1...1	27: .1...1...1	38: 1.....1.1	49: 1..1...1..
6:1...1.1	17: ..1...1.1.1	28: .1...1.1.1	39: 1...1...1	50: 1..1...1..
7:1.1.1.	18: ..1...1...1	29: .1...1...1	40: 1...1...1	51: 1..1...1.1
8:1...1.	19: ..1...1...1	30: .1...1...1	41: 1...1...1.	52: 1..1...1..
9:1...1.1	20: ..1...1.1.1	31: .1...1...1	42: 1...1...1.	53: 1..1...1.1
10: ...1...1.	21: .1.....1.	32: .1.1.1.1..	43: 1..1...1.1	54: 1..1.1.1.1

Figure 1.27-A: All 55 Fibonacci words with 8 bits in lexicographic order.

The 8-bit Fibonacci words are shown in figure 1.27-A. To generate all Fibonacci words in lexicographic order, use the function [FXT: [bits/fibrep.h](#)]

```

1 static inline ulong next_fibrep(ulong x)
2 // With x the Fibonacci representation of n
3 // return Fibonacci representation of n+1.
4 {
5     // 2 examples:           // ex. 1           // ex.2
6     //                     // x == [*]0 010101 // x == [*]0 01010
7     ulong y = x | (x>>1); // y == [*]? 011111 // y == [*]? 01111
8     ulong z = y + 1;     // z == [*]? 100000 // z == [*]? 10000
9     z = z & -z;          // z == [0]0 100000 // z == [0]0 10000
10    x ^= z;               // x == [*]0 110101 // x == [*]0 11010
11    x &= ~(z-1);          // x == [*]0 100000 // x == [*]0 10000
12
13    return x;
14 }

```

The routine can be used to generate all n -bit words as shown in [FXT: bits/fibrep2-demo.cc]:

```
const ulong f = 1UL << n;
ulong t = 0;
do
{
    // visit(t)
    t = next_fibrep(t);
}
while ( t!=f );
```

The reversed order can be generated via

```
ulong f = 1UL << n;
do
{
    f = prev_fibrep(f);
    // visit(f)
}
while ( f );
```

which uses the function (64-bit version)

```
1 static inline ulong prev_fibrep(ulong x)
2 // With x the Fibonacci representation of n
3 // return Fibonacci representation of n-1.
4 {
5     // 2 examples:           // ex. 1           // ex.2
6     //                       // x == [*]0 100000 // x == [*]0 10000
7     ulong y = x & -x;       // y == [0]0 100000 // y == [0]0 10000
8     x ^= y;                  // x == [*]0 000000 // x == [*]0 00000
9     ulong m = 0x5555555555555555UL; // m == ...01010101
10    if ( m & y ) m >>= 1; // m == ...01010101 // m == ...01010101
11    m &= (y-1);             // m == [0]0 010101 // m == [0]0 010101
12    x ^= m;                  // x == [*]0 010101 // x == [*]0 010101
13    return x;
14 }
```

The forward version generates about 180 million words per second, the backward version about 170 million words per second.

1.27.2 Gray code order ‡

A Gray code for the binary Fibonacci words (shown in figure 1.27-B) can be derived from the Gray code of the radix -2 representations (see section 1.22 on page 58) of binary words whose difference is of the form

```
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213
2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2240
2241
2242
2243
2244
2245
2246
2247
2248
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2260
2261
2262
2263
2264
2265
2266
2267
2268
2269
2270
2271
2272
2273
2274
2275
2276
2277
2278
2279
2280
2281
2282
2283
2284
2285
2286
2287
2288
2289
2290
2291
2292
2293
2294
2295
2296
2297
2298
2299
2300
2301
2302
2303
2304
2305
2306
2307
2308
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2320
2321
2322
2323
2324
2325
2326
2327
2328
2329
2330
2331
2332
2333
2334
2335
2336
2337
2338
2339
2340
2341
2342
2343
2344
2345
2346
2347
2348
2349
2350
2351
2352
2353
2354
2355
2356
2357
2358
2359
2360
2361
2362
2363
2364
2365
2366
2367
2368
2369
2370
2371
2372
2373
2374
2375
2376
2377
2378
2379
2380
2381
2382
2383
2384
2385
2386
2387
2388
2389
2390
2391
2392
2393
2394
2395
2396
2397
2398
2399
2400
2401
2402
2403
2404
2405
2406
2407
2408
2409
2410
2411
2412
2413
2414
2415
2416
2417
2418
2419
2420
2421
2422
2423
2424
2425
2426
2427
2428
2429
2430
2431
2432
2433
2434
2435
2436
2437
2438
2439
2440
2441
2442
2443
2444
2445
2446
2447
2448
2449
2450
2451
2452
2453
2454
2455
2456
2457
2458
2459
2460
2461
2462
2463
2464
2465
2466
2467
2468
2469
2470
2471
2472
2473
2474
2475
2476
2477
2478
2479
2480
2481
2482
2483
2484
2485
2486
2487
2488
2489
2490
2491
2492
2493
2494
2495
2496
2497
```

j:	k(j)	k(j)-k(j-1)	x=bin2neg(k)	gray(x)	
1:11...11	...111...1	...1..1..1	= 27
2:11....1	...111....	...1...1..	= 26
3:1.11111	...111...11	...1...1.1	= 28
4:1.11..11	...11111..	...1...1.1	= 23
5:1.1.111	...1111111	...1...1...1	= 21
6:1.1.1.11	...1111111	...1...1...1	= 22
7:1.1..11	...1111...1	...1...1...1	= 25
8:1.1...1	...1111....	...1...1...1	= 24
9:1...111.1	...11.111	...1.1.1..	= 32
10:1...1.1	...11..11	...1.1.1.1	= 33
11:1...11	...11...1	...1.1...1	= 30
12:1....1	...11....	...1.1...1	= 29
13:111111	...11...11	...1.1...1	= 31
14:11...111.11	...111...11	...1.1...1	= 10
15:1.111	...111111	...1...1...1	= 8
16:1.1.1	...1111..	...1...1...1	= 9
17:1.1..11	...11...1	...1.1.1..	= 12
18:1.1	...11....	...1.1...1	= 11
19:1...111.1	...11...11	...1...1...1	= 3
20:1.1	...11...	...1...1...1	= 4
21:11	...11....	...1...1...1	= 1
22:1	...11....	...1...1...1	= 0
23:	11111111111	...11...11	...1...1...1	= 2
24:	11111111..11	...11...11	...1...1...1	= 7
25:	11111111.111	...11111	...1...1...1	= 5
26:	1111111.1.1	...111..	...1...1...1	= 6
27:	111111...11.1	...11...1	...1.1.1..	= 19
28:	111111....1	...11....	...1.1...1	= 18
29:	11111.11111	...11...11	...1.1.1..	= 20
30:	11111.11..11	...1111..	...1...1...1	= 15
31:	11111.1.111	...111111	...1...1...1	= 13
32:	11111.1.1.1	...11111..	...1...1...1	= 14
33:	11111.1..11	...111...1	...1...1...1	= 17
34:	11111.1...1	...111...1	...1...1...1	= 16

Figure 1.27-B: Gray code for the binary Fibonacci words (rightmost column).

```

16     for (ulong m=(1UL<<(n-1)); m!=0; m>>=3) fw_ |= m;
17     lw_ = fw_ >> 1;
18     if ( 0==(n&1) ) { ulong t=fw_; fw_=lw_; lw_=t; } // swap first/last
19     mw_ = ( lw_>fw_ ? lw_ : fw_ );
20     x_ = fw_;
21     k_ = inverse_gray_code(fw_);
22     k_ = neg2bin(k_);
23 }
24
25 ~bit_fibgray() {;}
26
27 ulong next()
28 // Return next word in Gray code.
29 // Return ~0 if current word is the last one.
30 {
31     if ( x_ == lw_ ) return ~0UL;
32     ulong s = n_; // shift
33     while ( 1 )
34     {
35         --s;
36         ulong c = 1 | (mw_ >> s); // possible difference for negbin word
37         ulong i = k_ - c;
38         ulong x = bin2neg(i);
39         x ^= (x>>1);
40
41         if ( 0==(x&(x>>1)) ) // is_fibrep(x)
42         {
43             k_ = i;
44             x_ = x;
45             return x;
46         }
47     }
48 }
49 };

```

About 130 million words per second are generated. The program [FXT: [bits/bitfibgray-demo.cc](https://www.sfu.ca/~edp/fxt/bits/bitfibgray-demo.cc)] shows how to use the class, figure 1.27-B was created with it. Section 14.2 on page 305 gives a recursive algorithm for Fibonacci words in Gray code order.

1.28 Binary words and parentheses strings ‡

0	P	[empty string]	[empty string]
1	...1	P	()1	()
2	..1.			...11	(())
3	..11	P	(())	..1.1	(())
4	.1..			..111	((()))
5	.1.1	P	(())	.1.11	((()))
6	.11.			.11.1	((()))
7	.111	P	((()))	.1111	(((())))
8	1...			1..11	(())
9	1..1			1.1.1	(())
10	1.1.			1.111	((()))
11	1.11	P	((()))	11.11	((()))
12	11..			111.1	(((())))
13	11.1	P	((()))	11111	((((()))))
14	111.				
15	1111	P	((((()))))		

Figure 1.28-A: Left: some of the 4-bit binary words can be interpreted as a string parentheses (marked with ‘P’). Right: all 5-bit words that correspond to well-formed parentheses strings.

A subset of the binary words can be interpreted as a (well formed) string of parentheses. The 4-bit binary words that have this property are marked with a ‘P’ in figure 1.28-A (left) [FXT: bits/parenword-demo.cc]. The strings are constructed by scanning the word from the low end and printing a ‘(’ with each one and a ‘)’ with each zero. To find out when to terminate, one adds up +1 for each opening parenthesis and -1 for a closing parenthesis. After the ones in the binary word have been scanned, the s closing parentheses have to be added where s is the value of the sum [FXT: bits/parenwords.h]:

```

1  static inline void parenword2str(ulong x, char *str)
2  {
3      int s = 0;
4      ulong j = 0;
5      for (j=0; x!=0; ++j)
6      {
7          s += ( x&1 ? +1 : -1 );
8          str[j] = "("[x&1];
9          x >>= 1;
10     }
11     while ( s-- > 0 ) str[j++] = ')'; // finish string
12     str[j] = 0; // terminate string
13 }
```

The 5-bit binary words that are valid ‘paren words’ together with the corresponding strings are shown in figure 1.28-A (right). Note that the lower bits in the word (right end) correspond to the beginning of the string (left end). If a negative value for the sums occurs at any time of the computation, the word is not a paren word. A function to determine whether a word is a paren word is

```

1  static inline bool is_parenword(ulong x)
2  {
3      int s = 0;
4      for (ulong j=0; x!=0; ++j)
5      {
6          s += ( x&1 ? +1 : -1 );
7          if ( s<0 ) break; // invalid word
8          x >>= 1;
9      }
10     return (s>=0);
11 }
```

The sequence

1, 3, 5, 7, 11, 13, 15, 19, 21, 23, 27, 29, 31, 39, 43, 45, 47, 51, 53, 55, 59, 61, 63, ...

of nonzero integers x so that `is_parenword(x)` returns true is entry A036991 in [312]. If we fix the number of paren pairs, then the following functions generate the least and biggest valid paren words. The first paren word is a block of n ones at the low end:

```

1  static inline ulong first_parenword(ulong n)
2  // Return least binary word corresponding to n pairs of parens
3  // Example, n=5: .....11111  ((((( )))))
```

```

4 {
5     return first_comb(n);
6 }

```

The last paren word is the word with a sequence of n blocks '01' at the low end:

```

1 static inline ulong last_parenword(ulong n)
2 // Return biggest binary word corresponding to n pairs of parens.
3 // Must have: 1 <= n <= BITS_PER_LONG/2.
4 // Example, n=5: .1.1.1.1.1 ( ) ( ) ( )
5 {
6     return 0x5555555555555555UL >> (BITS_PER_LONG-2*n);
7 }

```

.....11111 = (((((()))))	...1...1111 = (((()))())	..1....1111 = (((()))()()
.....1.1111 = (((()())))	...1..1.111 = (((()())())	..1...1.111 = (((()())()())
.....11.111 = (((()()()))	...1..11.11 = (((()())()())	..1...11.11 = (((()())()()())
.....111.11 = (((()((()))))	...1..111.1 = (((()())()())	..1...111.1 = (((()())()()())
.....1111.1 = (((()((()()))))	...1.1..111 = (((()())()())	..1..1..111 = (((()())()()())
.....1..1111 = (((()())())	...1.1.1.11 = (((()())()())	..1..1.1.11 = (((()())()()())
....1.1.111 = (((()())())	...1.1.11.1 = (((()())()())	..1..1.11.1 = (((()())()()())
....1.11.11 = (((()())())	...1.11..11 = (((()())()())	..1..11..11 = (((()())()()())
....1.111.1 = (((()())())	...1.11.1.1 = (((()())()())	..1..11.1.1 = (((()())()()())
....11..111 = (((()())())	...11...111 = (((()())()())	..1.1...111 = (((()())()()())
....11.1.11 = (((()())())	...11..1.11 = (((()())()())	..1.1..1.11 = (((()())()()())
....11.11.1 = (((()())())	...11..11.1 = (((()())()())	..1.1..11.1 = (((()())()()())
....111..11 = (((()())())	...11.1..11 = (((()())()())	..1.1.1..11 = (((()())()()())
....111.1.1 = (((()())())	...11.1.1.1 = (((()())()())	..1.1.1.1.1 = (((()())()()())

Figure 1.28-B: The 42 binary words corresponding to all valid pairings of 5 parentheses, in colex order.

The sequence of all binary words corresponding to n pairs of parens in colex order can be generated with the following (slightly cryptic) function:

```

1 static inline ulong next_parenword(ulong x)
2 // Next (colex order) binary word that is a paren word.
3 {
4     if ( x & 2 ) // Easy case, move highest bit of lowest block to the left:
5     {
6         ulong b = lowest_zero(x);
7         x ^= b;
8         x ^= (b>>1);
9         return x;
10    }
11    else // Gather all low "01"s and split lowest nontrivial block:
12    {
13        if ( 0==(x & (x>>1)) ) return 0;
14        ulong w = 0; // word where the bits are assembled
15        ulong s = 0; // shift for lowest block
16        ulong i = 1; // == lowest_one(x)
17        do // collect low "01"s:
18        {
19            x ^= i;
20            w <<= 1;
21            w |= 1;
22            ++s;
23            i <<= 2; // == lowest_one(x);
24        }
25        while ( 0==(x&(i<<1)) );
26
27        ulong z = x ^ (x+i); // lowest block
28        x ^= z;
29        z &= (z>>1);
30        z &= (z>>1);
31        w ^= (z>>s);
32        x |= w;
33        return x;
34    }
35 }

```

The program [FXT: bits/parenword-colex-demo.cc] shows how to create a list of binary words corresponding to n pairs of parens (code slightly shortened):

```

1     ulong n = 4; // Number of paren pairs

```


swapped. An implementation of the implied algorithm is given in [FXT: bits/bitperm1-demo.cc]. Arrays are used to give more readable code:

```

1 void perm1(uchar *a, ulong ldn, const uchar *x)
2 // Permute a[] according to the 'control word' x[].
3 // The length of a[] must be 2**ldn.
4 {
5     long n = 1L<<ldn;
6     for (long s=n/2; s>0; s/=2)
7     {
8         for (long k=0; k<n; k+=s+s)
9         {
10            if ( x[k+s-1]!='0' )
11            {
12                // swap regions [a+k,...,a+k+s-1], [a+k+s,...,a+k+2*s-1]:
13                swap(a+k, a+k+s, s);
14            }
15        }
16    }
17 }
```

The routine for the inverse permutation differs in a single line:

```
for (long s=1; s<n; s+=s)
```

No attempt has been made to optimize or parallelize the algorithm. We just explore how useful a machine instruction for the permutation of bits would be.

The program uses a fixed size of 16 bits, an 'x' is printed whenever the corresponding bit is set:

```

a=0123456789ABCDEF  bits of the input word
x=0010011000110110  control word
  8:      7
  4:      3  11x
  2:      1  5x  9  13x
  1:      0  2x  4  6x  8  10x 12  14x
a=01326754CDFEAB98  result
```

This control word leads to the Gray permutation (see 2.12 on page 128). Assume we use words with N bits. We cannot (for $N > 2$) specify all $N!$ permutations as we can choose between only 2^{N-1} control words. Now set the word length to $N := 2^n$. The reachable permutations are those where the intervals $[k \cdot 2^j, \dots, (k+1) \cdot 2^j - 1]$ contain all numbers $[p \cdot 2^j, \dots, (p+1) \cdot 2^j - 1]$ for all $j \leq n$ and $0 \leq k < 2^{n-j}$, choosing p for each interval arbitrarily ($0 \leq p < 2^{n-j}$). For example, the lower half of the permuted array must contain a permutation of either the lower or the upper half ($j = n - 1$) and each pair a_{2y}, a_{2y+1} must contain two elements $2z, 2z + 1$ ($j = 1$). The bit-reversal is computed with a control word where all bits are set. Alas, the (important!) zip permutation (bit-zip, see section 1.15 on page 38) is unreachable.

A machine instruction could choose between the two routines via the highest bit in the control word.

1.29.2 A general method

All permutations of $N = 2^n$ elements can be specified with n control words of N bits. Assume we have a machine instruction that collects bits according to a control word. An eight bit example:

```

a = abcdefgh  input data
x = ..1.11.1  control word (dots for zeros)
  abdg cefh  bits of a where x has a one
  abdg cefh  bits of a where x has a zero
  abdgcefh  result, bits separated according to x
```

We need n such instructions that work on all length- 2^k sub-words for $1 \leq k \leq n$. For example, the instruction working on half words of a 16-bit word would work as

```

a = abcdefgh ABCDEFGH  input data
x = ..1.11.1 1111....  control word (dots for zeros)
  abdg cefh ABCD  bits of a where x has a one
  abdg cefh EFGH  bits of a where x has a zero
  abdgcefh EFGHABCD  result, bits separated according to x
```

Note the bits of the different sub-words are not mixed. Now all permutations can be reached if the control word for the 2^k -bit sub-words have exactly 2^{k-1} bits set in all ranges $[j \cdot 2^k, \dots, (j+1) \cdot 2^k]$.

- A fused multiply-add instruction for floats.
- Instructions for the multiplication of complex floating-point numbers, computing $A \cdot C - B \cdot D$ and $A \cdot D + B \cdot C$ from A , B , C , and D .
- A sum-diff instruction, computing $A + B$ and $A - B$ from A and B . This can serve as a primitive for fast orthogonal transforms.
- An instruction to swap registers. Even better, a conditional version of that.

1.30.2 Nice to have

- A parity bit for the complete machine word. The parity of a word is the number of bits modulo 2, not the complement of it. Even better, an instruction for the inverse Gray code, see section 1.16 on page 41.
- A bit-count instruction, see section 1.8 on page 18. This would also give the parity at bit zero.
- An instruction for computing the index of the i -th set bit of a word, see section 1.10 on page 25. This would be useful even if execution takes a dozen cycles.
- A random number generator, LHCA's (see section 41.8 on page 878) may be candidates. At the very least: a decent entropy source.
- A conditional version of more than just the move instruction, possibly as an instruction prefix.
- A bit-zip and a bit-unzip instruction, see section 1.15 on page 38. Note this is polynomial squaring over $\text{GF}(2)$.
- Primitives for permutations of bits, see section 1.29.2 on page 81. A bit-gather and a bit-scatter instruction for sub-words of all sizes a power of 2 would allow for arbitrary permutations (see [FXT: bits/bitgather.h] and [FXT: bits/bitseparate.h] for versions working on complete words).
- Multiplication corresponding to XOR as addition. This is the multiplication without carries used for polynomials over $\text{GF}(2)$, see section 40.1 on page 822.

1.31 Some space filling curves ‡

1.31.1 The Hilbert curve

A rendering of the Hilbert curve (named after David Hilbert [182]) is shown in figure 1.31-A. An efficient algorithm to compute the direction of the n -th move of the Hilbert curve is based on the parity of the number of threes in the radix-4 representation of n (see section 38.9.1 on page 748).

Let d_x and d_y correspond to the moves at step n in the Hilbert curve. Then $d_x, d_y \in \{-1, 0, +1\}$ and exactly one of them is zero. So for both $p := d_x + d_y$ and $m := d_x - d_y$ we have $p, m \in \{-1, +1\}$.

The following function computes p and returns 0, 1 if $p = -1, +1$, respectively [FXT: bits/hilbert.h]:

```

1  static inline ulong hilbert_p(ulong t)
2  // Let dx,dy be the horizontal,vertical move
3  // with step t of the Hilbert curve.
4  // Return zero if (dx+dy)==-1, else one (then: (dx+dy)==+1).
5  // Algorithm: count number of threes in radix 4
6  {
7      ulong d = (t & 0x5555555555555555UL) & ((t & 0xaaaaaaaaaaaaaaaaUL) >> 1);
8      return parity( d );
9  }
```

If 1 is returned the step is to the right or upwards. The function can be slightly optimized as follows (64-bit version only):

```

1  static inline ulong hilbert_p(ulong t)
2  {
3      t &= ((t & 0xaaaaaaaaaaaaaaaaUL) >> 1);
```



```

11     ulong p = hilbert_p(t);
12     ulong m = hilbert_m(t);
13     ulong d = p ^ (m<<1);
14     return d;
15 }

```

To print the value of d symbolically, we can print the value of (" $>v^<$ ") [d]. The sequence of moves can also be generated by the string substitution process shown in figure 1.31-C.

```

Start: A
Rules:
A --> D>A^A<C
B --> C<BvB>D
C --> BvC<C^A
D --> A^D>DvB
> --> >
< --> <
v --> v
-----
0:   (#=1)
    A
1:   (#=7)
    D>A^A<C
2:   (#=31)
    A^D>DvB>D>A^A<C^D>A^A<C<BvC<C^A
3:   (#=127)
    D>A^A<C^A^D>DvB>A^D>DvBvC<BvB>D>A^D>DvB>D>A^A<C^D>A^A<C<BvC<C^A^A^D>DvB>D>A^A<C^D> ...

```

Figure 1.31-C: Moves of the Hilbert curve by a string substitution process, the symbols ‘A’, ‘B’, ‘C’, and ‘D’, are ignored when drawing the curve.

The turn u between steps can be computed as

```

1  static inline int hilbert_turn(ulong t)
2  // Return the turn (left or right) with the steps
3  // t and t-1 of the Hilbert curve.
4  // Returned value is
5  // 0 for no turn
6  // +1 for right turn
7  // -1 for left turn
8  {
9      ulong d1 = hilbert_dir(t);
10     ulong d2 = hilbert_dir(t-1);
11     d1 ^= (d1>>1);
12     d2 ^= (d2>>1);
13     ulong u = d1 - d2;
14     // at this point, symbolically: cout << ("+. -0+.-")[ u + 3 ];
15     if ( 0==u ) return 0;
16     if ( (long)u<0 ) u += 4;
17     return (1==u ? +1 : -1);
18 }

```

To print the value of u symbolically, we can print (" $-0+$ ") [$u+1$];.

The values of p and m , followed by the direction and turn of the Hilbert curve are shown in figure 1.31-B. The list was created with the program [FXT: bits/hilbert-moves-demo.cc]. Figure 1.31-A was created with the program [FXT: bits/hilbert-texpic-demo.cc]. The computation of a function whose series coefficients are ± 1 and $\pm i$ according to the Hilbert curve is described in section 38.9 on page 747.

A finite state machine (FSM) for the conversion from a 1-dimensional coordinate (linear coordinate of the curve) to the pair of coordinates x and y of the Hilbert curve is described in [39, item 115]. At each step two bits of input are processed. The array `htab[]` serves as lookup table for the next state and two bits of the result. The FSM has an internal state of two bits [FXT: bits/lin2hilbert.cc]:

```

1  void
2  lin2hilbert(ulong t, ulong &x, ulong &y)
3  // Transform linear coordinate t to Hilbert x and y
4  {
5      ulong xv = 0, yv = 0;
6      ulong c01 = (0<<2); // (2<<2) for transposed output (swapped x, y)
7      for (ulong i=0; i<(BITS_PER_LONG/2); ++i)
8      {
9          ulong abi = t >> (BITS_PER_LONG-2);
10         t <<= 2;

```

```

11
12     ulong st = htab[ (c01<<2) | abi ];
13     c01 = st & 3;
14
15     yv <<= 1;
16     yv |= ((st>>2) & 1);
17     xv <<= 1;
18     xv |= (st>>3);
19 }
20 x = xv;  y = yv;
21 }

```

OLD				NEW				NEW				OLD			
C	C	A	B	X	Y	C	C	C	C	X	Y	A	B	C	C
0	1	I	I	I	I	0	1	0	1	I	I	I	I	0	1
0	0	0	0	0	0	1	0	0	0	0	0	0	0	1	0
0	0	0	1	0	1	0	0	0	0	0	1	0	0	0	0
0	0	1	0	1	1	0	0	0	0	1	0	1	0	0	1
0	0	1	1	1	0	0	1	0	0	1	1	0	0	0	0
0	1	0	0	1	1	1	1	0	1	0	0	1	0	0	1
0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1
0	1	1	0	0	0	0	0	0	1	1	0	1	1	0	0
0	1	1	1	1	0	0	0	0	1	1	1	0	0	1	1
1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	0	0	1	1	0	1	0	1	0	0	1	1	1	1	1
1	0	1	0	1	1	1	0	1	0	1	0	0	1	1	0
1	0	1	1	0	1	1	1	1	1	1	1	0	1	1	0
1	1	0	0	1	1	0	1	1	1	0	0	1	0	1	1
1	1	0	1	1	0	1	1	1	1	0	1	1	1	0	0
1	1	1	0	0	0	1	1	1	1	1	0	1	1	1	1
1	1	1	1	0	1	1	0	1	1	1	1	0	0	0	1

Figure 1.31-D: The original table from [39] for the finite state machine for the 2-dimensional Hilbert curve (left). All sixteen 4-bit words appear in both the ‘OLD’ and the ‘NEW’ column. So the algorithm is invertible. Swap the columns and sort numerically to obtain the new two columns at the right, the table for the inverse function.

The table used is defined (see figure 1.31-D) as

```

1  static const ulong htab[] = {
2  #define HT(xi,yi,c0,c1) ((xi<<3)+(yi<<2)+(c0<<1)+(c1))
3  // index == HT(c0,c1,ai,bi)
4  HT( 0, 0,  1, 0 ),
5  HT( 0, 1,  0, 0 ),
6  HT( 1, 1,  0, 0 ),
7  HT( 1, 0,  0, 1 ),
8  [--snip--]
9  HT( 0, 0,  1, 1 ),
10 HT( 0, 1,  1, 0 )
11 };

```

As indicated in the code, the table maps every four bits c_0, c_1, a_i, b_i to four bits x_i, y_i, c_0, c_1 . The table for the inverse function (again, see figure 1.31-D) is

```

1  static const ulong ihtab[] = {
2  #define IHT(ai,bi,c0,c1) ((ai<<3)+(bi<<2)+(c0<<1)+(c1))
3  // index == HT(c0,c1,xi,yi)
4  IHT( 0, 0,  1, 0 ),
5  IHT( 0, 1,  0, 0 ),
6  IHT( 1, 1,  0, 1 ),
7  IHT( 1, 0,  0, 0 ),
8  [--snip--]
9  IHT( 0, 1,  1, 1 ),
10 IHT( 0, 0,  0, 1 )
11 };

```

The words have to be processed backwards:

```

1  ulong
2  hilbert2lin(ulong x, ulong y)
3  // Transform Hilbert x and y to linear coordinate t
4  {
5     ulong t = 0;
6     ulong c01 = 0;

```

```

7   for (ulong i=0; i<(BITS_PER_LONG/2); ++i)
8   {
9       t <<= 2;
10      ulong xi = x >> (BITS_PER_LONG/2-1);
11      xi &= 1;
12      ulong yi = y >> (BITS_PER_LONG/2-1);
13      yi &= 1;
14      ulong xyi = (xi<<1) | yi;
15      x <<= 1;
16      y <<= 1;
17
18      ulong st = ihtab[ (c01<<2) | xyi ];
19      c01 = st & 3;
20
21      t |= (st>>2);
22  }
23
24  return t;
25  }

```

1.31.2 The Z-order




Figure 1.31-E: The first 255 segments of the Z-order curve.

A 2-dimensional space-filling curve in *Z-order* traverses all points in each quadrant before it enters the next. Figure 1.31-E shows a rendering of the Z-order curve, created with the program [FXT: [bits/zorder-texpic-demo.cc](#)]. The conversion between a linear parameter to a pair of coordinates is done by separating the bits at the even and odd indices [FXT: [bits/zorder.h](#)]:

```
static inline void lin2zorder(ulong t, ulong &x, ulong &y) { bit_unzip2(t, x, y); }
```

The routine `bit_unzip2()` is described in section 1.15 on page 38. The inverse is

```
static inline ulong zorder2lin(ulong x, ulong y) { return bit_zip2(x, y); }
```

The next pair can be computed with the following (constant amortized time) routine:

```

1  static inline void zorder_next(ulong &x, ulong &y)
2  {

```

```

3     ulong b = 1;
4     do
5     {
6         x ^= b;  b &= ~x;
7         y ^= b;  b &= ~y;
8         b <<= 1;
9     }
10    while ( b );
11 }

```

The previous pair is computed similarly:

```

1  static inline void zorder_prev(ulong &x, ulong &y)
2  {
3      ulong b = 1;
4      do
5      {
6          x ^= b;  b &= x;
7          y ^= b;  b &= y;
8          b <<= 1;
9      }
10     while ( b );
11 }

```

The routines are written in a way that generalizes easily to more dimensions:

```

1  static inline void zorder3d_next(ulong &x, ulong &y, ulong &z)
2  {
3      ulong b = 1;
4      do
5      {
6          x ^= b;  b &= ~x;
7          y ^= b;  b &= ~y;
8          z ^= b;  b &= ~z;
9          b <<= 1;
10     }
11     while ( b );
12 }

1  static inline void zorder3d_prev(ulong &x, ulong &y, ulong &z)
2  {
3      ulong b = 1;
4      do
5      {
6          x ^= b;  b &= x;
7          y ^= b;  b &= y;
8          z ^= b;  b &= z;
9          b <<= 1;
10     }
11     while ( b );
12 }

```

Unlike with the Hilbert curve there are steps where the curve advances more than one unit.

1.31.3 Curves via paper-folding sequences

The *paper-folding sequence*, entry [A014577](#) in [312], starts as [FXT: [bits/bit-paper-fold-demo.cc](#)]:

```
11011001110010011101100011001001110110011100100011011000110010011 ...
```

The k -th element ($k > 0$) is one if $k = 2^t \cdot (4u + 1)$, entry [A091072](#) in [312]:

```
1, 2, 4, 5, 8, 9, 10, 13, 16, 17, 18, 20, 21, 25, 26, 29, 32, 33, ...
```

The k -th element of the paper-folding sequence can be computed by testing the value of the bit left to the lowest (that is, rightmost) one in the binary expansion of k [FXT: [bits/bit-paper-fold.h](#)]:

```

1  static inline bool bit_paper_fold(ulong k)
2  {
3      ulong h = k & -k;  // == lowest_one(k)
4      k &= (h<<1);
5      return ( k==0 );
6  }

```

About 550 million values per second are generated. We use `bool` as return type to indicate that only zero or one is returned. The value can be used as an integer of arbitrary type, there is no need for a cast.





Figure 1.31-F: The first 1024 segments of the dragon curve (two different renderings).


1.31.3.1 The dragon curve

Another name for the sequence is *dragon curve sequence*, because a space filling curve known as dragon curve (or *Heighway dragon*) can be generated if we interpret a one as ‘turn left’ and a zero as ‘turn right’. The top of figure 1.31-F shows the first 1024 segments of the curve (created with [FXT: [bits/dragon-curve-texpic-demo.cc](#)]). As some points are visited twice we draw the turns with cut off corners, for the (left) turn $A \rightarrow B \rightarrow C$:



The code is given in [FXT: [aux0/tex-line.cc](#)]. The first few moves of the curve can be found by repeatedly folding a strip of paper. Always pick up the right side and fold to the left. Unfold the paper and adjust all corners to be 90 degrees. This gives the first few segments of the dragon curve.

When all angles are replaced by diagonals between the midpoints of the lines



then the curve appears as shown at the bottom of figure 1.31-F.

```

Start: 0
Rules:
0 --> 01
1 --> 21
2 --> 23
3 --> 03
-----
0: 0
1: 01
2: 0121
3: 01212321
4: 0121232123032321
5: 01212321230323212303010323032321
6: 012123212303232123030103230301032303010323032321
+^-^-v^-v+v-v^-v+v+^-v+v+v-v^-v+v+^-v+^-+^-+v-v+v+^-v+v+v-v^-

```

Figure 1.31-G: Moves of the dragon curve generated by a string substitution process.

The net rotation of the dragon-curve after k steps, as multiple of the right angle, can be computed by counting the ones in the Gray code of k . Take the result modulo 4 to ignore multiples of 360 degree [FXT: [bits/bit-paper-fold.h](#)]:

```
1 static inline bool bit_dragon_rot(ulong k) { return bit_count(k ^ (k>>1)) & 3; }
```

The sequence of rotations is entry [A005811](#) in [312]:

```

seq    = 0 1 2 1 2 3 2 1 2 3 4 3 2 3 2 1 2 3 4 3 4 5 4 3 2 3 4 3 2 3 2 1 2 3 ...
mod 4  = 0 1 2 1 2 3 2 1 2 3 0 3 2 3 2 1 2 3 0 3 0 1 0 3 2 3 0 3 2 3 2 1 2 3 ...
move   = + ^ - - v - ^ - v + v - v - ^ - v + v + ^ + v - v + v - v - ^ - v ...

```

The sequence of moves (as symbols, last row) can be computed with [FXT: [bits/dragon-curve-moves-demo.cc](#)]. A function related to the paper-folding sequence is described in section 38.8.3 on page 744.

1.31.3.2 The alternate paper-folding sequence

If the strip of paper is folded alternately from the left and right, then another paper-folding sequence is obtained. It is entry [A106665](#) in [312] and it starts as [FXT: [bits/bit-paper-fold-alt-demo.cc](#)]:

```
10011100100011011001110110001100100111001000110010011101100011011 ...
```

Compute the sequence via [FXT: [bits/bit-paper-fold.h](#)]

```

1 static inline bool bit_paper_fold_alt(ulong k)
2 {
3     ulong h = k & -k; // == lowest_one(k)

```




Figure 1.31-H: The first 512 segments of the curve from the alternate paper-folding sequence.

```

Start: 0
Rules:
  0 --> 01
  1 --> 03
  2 --> 23
  3 --> 21
-----
0: 0
1: 01
2: 0103
3: 01030121
4: 0103012101032303
5: 01030121010323030103012123210121
6: 0103012101032303010301212321012101030121010323032321230301032303
  +^+v+^-^-+^+v-v+v+^+v+^-^-v-^-+^-^-+^+v+^-^-+^+v-v+v-v-^-v+v+^+v-v+v
    
```

Figure 1.31-I: Moves of the alternate curve generated by a string substitution process.




Figure 1.31-K: The first 729 segments of the terdragon (two different renderings).




Figure 1.31-O: The first 625 segments of the R5-dragon (two different renderings).

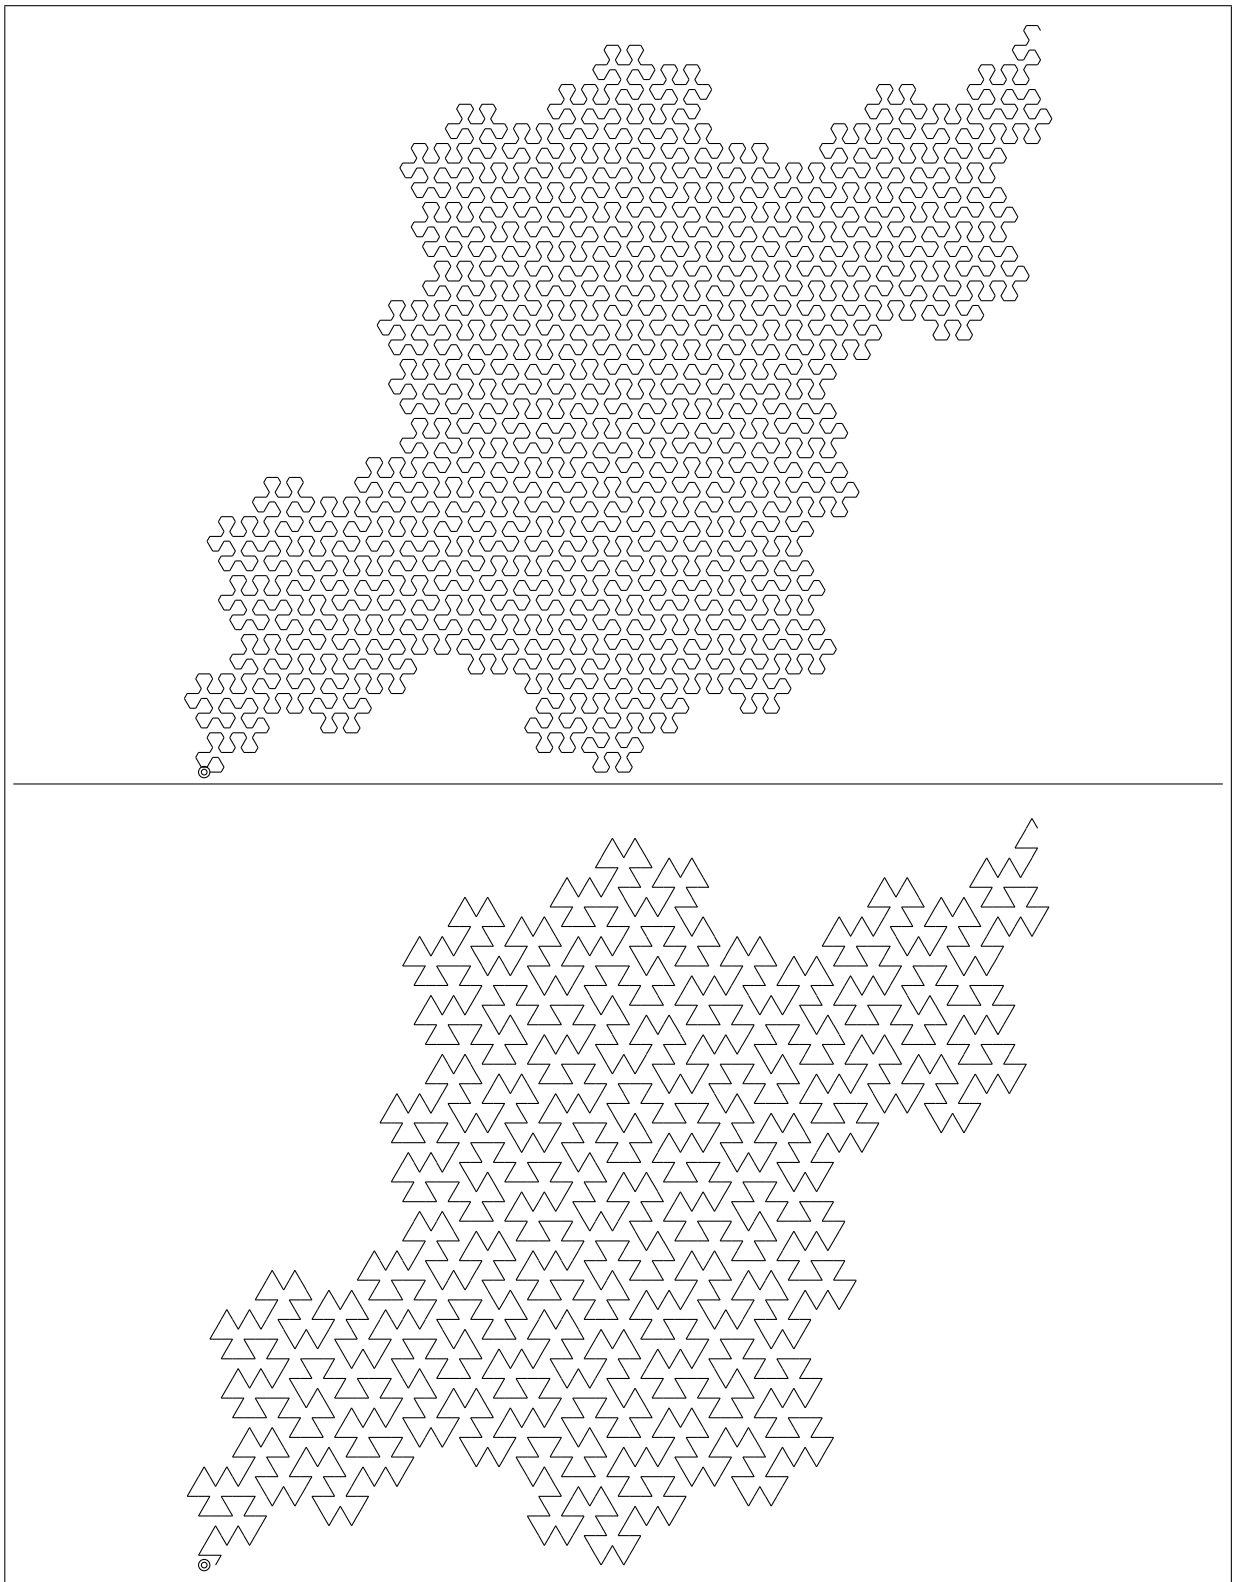


Figure 1.31-P: The first 2401 segments of the R7-dragon (two different renderings).

```

3 // return (tr) whether the lowest nonzero digit
4 // of the incremented word is either 2, 3, or 6.

```

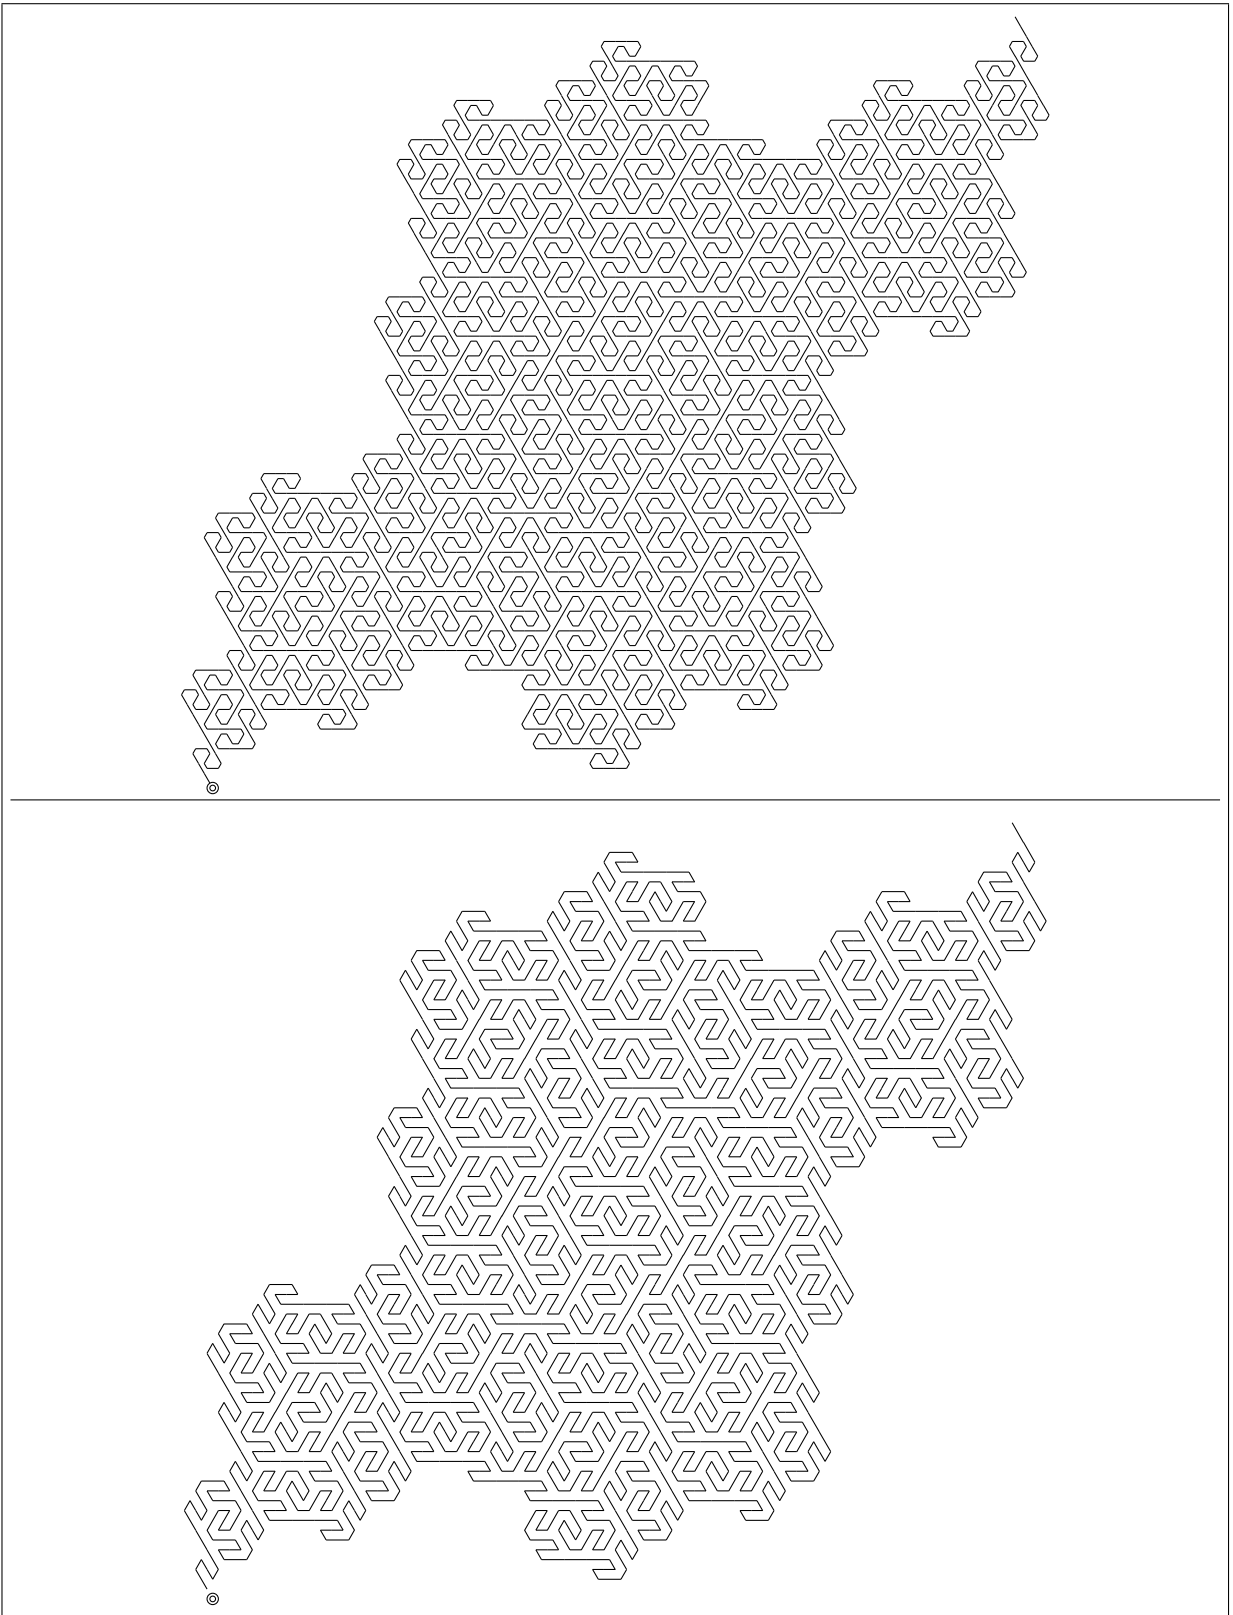



Figure 1.31-Q: The first 2401 segments of the second R7-dragon (two different renderings).

<pre> Start: 0 Rules: 0 --> 00110 1 --> 00111 ----- 0: (#=1) 0 1: (#=5) 00110 2: (#=25) 0011000110001110011100110 3: (#=125) 001100011000111001110011000110001100011100111001100011000110001110011100 \ 110001100011100111001110011000110001110011100110 </pre>
<pre> Start: 0 Rules: 0 --> 0100110 1 --> 0110110 ----- 0: (#=1) 0 1: (#=7) 0100110 2: (#=49) 0100110011011001001100100110011011001101100100110 3: (#=343) 01001100110110010011001001100110110011011001001100100110110011011001001 ... </pre>
<pre> Start: 0 Rules: 0 --> 0+---00 + --> 0+---0+ - --> 0+---0- ----- 0: (#=1) 0 1: (#=7) 0+---00 2: (#=49) 0+---000+---0+0+---0+0+---0+0+---0+0+---0+0+---000+---00 3: (#=343) 0+---000+---0+0+---0+0+---0+0+---0+0+---0+0+---000+---000+---000+---0+0+---0+0+--- ... </pre>

Figure 1.31-R: Turns of the R5-dragon (top), the R7-dragon (middle), and the second R7-dragon (bottom), generated by string substitution.

```

5  {
6      ulong s = 0;
7      while ( (x & 7) == 6 ) { x >>= 3; ++s; } // scan over nines
8      ++x; // increment next digit
9      bool tr = ( x & 2 ); // whether digit is either 2, 3, or 6
10     x <<= (3*s); // shift back
11     return tr;
12 }
13

```

Two renderings of the R7-dragon are shown in figure 1.31-P (created with [FXT: [bits/dragon-r7-texpic-demo.cc](http://bits.dragon-r7-texpic-demo.cc)]). The sequence of returned values (entry A176405 in [312]) can be computed via the string substitution shown in figure 1.31-R (middle). Turns for another curve based on radix-7 counting (shown in figure 1.31-Q, created with [FXT: bits/dragon-r7-2-texpic-demo.cc]) can be computed as follows:

```

1  static inline int bit_dragon_r7_2_turn(ulong &x)
2  // Increment the radix-7 word x and
3  // return (tr) according to the lowest nonzero digit d
4  // of the incremented word:
5  // d=[1,2,3,4,5,6] ==> rt:=[0,+1,+1,-1,-1,0]
6  // (tr * 120deg) is the turn with the second R7-dragon.
7  {
8      ulong s = 0;
9      while ( (x & 7) == 6 ) { x >>= 3; ++s; } // scan over nines
10     ++x; // increment next digit
11     int tr = 2 - ( (0x2f58 >> (2*(x&7)) ) & 3 );
12     x <<= (3*s); // shift back
13     return tr;
14 }

```

The sequence of turns can be generated by the string substitution shown in figure 1.31-R (bottom), it is

<pre> Start: F Rules: F --> F+F+F-F-F + --> + - --> - ----- 0: (#=1) F 1: (#=9) F+F+F-F-F 2: (#=49) F+F+F-F-F+F+F+F-F-F+F+F+F-F-F+F+F+F-F-F+F+F-F-F 3: (#=249) F+F+F-F-F+F+F+F-F-F+F+F+F-F-F+F+F+F-F-F+F+F-F-F+F+F-F-F+F+F-F-F-F-F+ ... </pre>
<pre> Start: F Rules: F --> F+F-F-F+F+F-F + --> + - --> - ----- 0: (#=1) F 1: (#=13) F+F-F-F+F+F-F 2: (#=97) F+F-F-F+F+F-F+F+F-F-F+F+F-F-F+F+F-F-F+F+F-F-F+F+F-F-F+F+F-F-F+F+F-F-F+ ... </pre>
<pre> Start: F Rules: F --> FOF+F+F-F-FOF + --> + - --> - 0 --> 0 ----- 0: (#=1) F 1: (#=13) FOF+F+F-F-FOF 2: (#=97) FOF+F+F-F-F-FOF-FOF+F+F-F-FOF-FOF+F+F-F-FOF-FOF+F+F-F-FOF-FOF+F+F-F-FOF-FOF+F+F-F-F-FO ... </pre>

Figure 1.31-S: String substitution processes for the turns (symbols ‘+’ and ‘-’) and moves (symbol ‘F’ is a unit move in the current direction) of the R5-dragon (top), the R7-dragon (middle), and the second R7-dragon (bottom).

entry [A176416](#) in [\[312\]](#).

Two curves respectively based on radix-9 and radix-13 counting are shown in figure [1.31-T](#). The corresponding routines are given in [\[FXT: bits/bit-dragon-r9.h\]](#)

```

1  static inline bool bit_dragon_r9_turn(ulong &x)
2  // Increment the radix-9 word x and
3  // return (tr) whether the lowest nonzero digit
4  // of the incremented word is either 2, 3, 5, or 8.
5  // tr determines whether to turn left or right (by 120 degrees)
6  // with the R9-dragon fractal.
7  // The sequence tr is the fixed point
8  // of the morphism 0 |--> 011010010, 1 |--> 011010011.
9  // Also fixed point of morphism (identify + with 0 and - with 1)
10 // F |--> F+F-F-F+F-F-F+F-F-F, + |--> +, - |--> -
11 // Also fixed point of morphism
12 // F |--> G+G-G, G |--> F-F+F, + |--> +, - |--> -
13 {
14     ulong s = 0;
15     while ( (x & 15) == 8 ) { x >>= 4; ++s; } // scan over nines
16     ++x; // increment next digit
17     bool tr = ( (0x12c >> (x&15)) & 1 ); // whether digit is either 2, 3, 5, or 8
18     x <<= (4*s); // shift back
19     return tr;
20 }

```

and [\[FXT: bits/bit-dragon-r13.h\]](#)

```

1  static inline bool bit_dragon_r13_turn(ulong &x)
2  // Increment the radix-13 word x and
3  // return (tr) whether the lowest nonzero digit
4  // of the incremented word is either 3, 6, 8, 9, 11, or 12.
5  // tr determines whether to turn left or right (by 90 degrees)
6  // with the R13-dragon fractal.
7  // The sequence tr is the fixed point
8  // of the morphism 0 |--> 0010010110110, 1 |--> 0010010110111.
9  // Also fixed point of morphism (identify + with 0 and - with 1)
10 // F |--> F+F+F-F+F-F-F+F-F-F-F, + |--> +, - |--> -
11 {
12     ulong s = 0;

```

```

13   while ( (x & 15) == 12 ) { x >>= 4; ++s; } // scan over nines
14   ++x; // increment next digit
15   bool tr = ( (0x1b48 >> (x&15)) & 1 ); // whether digit is either 3, 6, 8, 9, 11, or 12
16   x <<= (4*s); // shift back
17   return tr;
18 }

```




Figure 1.31-T: The R9-dragon (top) and the R13-dragon (bottom).