

```

FUNCTION igray(n,is)
INTEGER igray,is,n
    For zero or positive values of is, return the Gray code of n; if is is negative, return the
    inverse Gray code of n.
INTEGER idiv,ish
if (is.ge.0) then          This is the easy direction!
    igray=ieor(n,n/2)
else                      This is the more complicated direction: In hierarchical stages,
    ish=-1                starting with a one-bit right shift, cause each bit to be
    igray=n                XORed with all more significant bits.
1   continue
    idiv=ishft(igray,ish)
    igray=ieor(igray,idiv)
    if(idiv.le.1.or.ish.eq.-16)return
    ish=ish+ish           Double the amount of shift on the next cycle.
    goto 1
endif
return
END

```

In numerical work, Gray codes can be useful when you need to do some task that depends intimately on the bits of i , looping over many values of i . Then, if there are economies in repeating the task for values differing by only one bit, it makes sense to do things in Gray code order rather than consecutive order. We saw an example of this in §7.7, for the generation of quasi-random sequences.

CITED REFERENCES AND FURTHER READING:

- Horowitz, P., and Hill, W. 1989, *The Art of Electronics*, 2nd ed. (New York: Cambridge University Press), §8.02.
- Knuth, D.E. *Combinatorial Algorithms*, vol. 4 of *The Art of Computer Programming* (Reading, MA: Addison-Wesley), §7.2.1. [Unpublished. Will it be always so?]

20.3 Cyclic Redundancy and Other Checksums

When you send a sequence of bits from point A to point B, you want to know that it will arrive without error. A common form of insurance is the “parity bit,” attached to 7-bit ASCII characters to put them into 8-bit format. The parity bit is chosen so as to make the total number of one-bits (versus zero-bits) either always even (“even parity”) or always odd (“odd parity”). Any *single bit* error in a character will thereby be detected. When errors are sufficiently rare, and do not occur closely bunched in time, use of parity provides sufficient error detection.

Unfortunately, in real situations, a single noise “event” is likely to disrupt more than one bit. Since the parity bit has two possible values (0 and 1), it gives, on average, only a 50% chance of detecting an erroneous character with more than one wrong bit. That probability, 50%, is not nearly good enough for most applications. Most communications protocols [1] use a multibit generalization of the parity bit called a “cyclic redundancy check” or CRC. In typical applications the CRC is 16 bits long (two bytes or two characters), so that the chance of a random error going undetected is 1 in $2^{16} = 65536$. Moreover, M -bit CRCs have the mathematical property of detecting *all* errors that occur in M or fewer *consecutive* bits, for any

length of message. (We prove this below.) Since noise in communication channels tends to be “bursty,” with short sequences of adjacent bits getting corrupted, this consecutive-bit property is highly desirable.

Normally CRCs lie in the province of communications software experts and chip-level hardware designers — people with bits under their fingernails. However, there are at least two kinds of situations where some understanding of CRCs can be useful to the rest of us. First, we sometimes need to be able to communicate with a lower-level piece of hardware or software that expects a valid CRC as part of its input. For example, it can be convenient to have a program generate XMODEM or Kermit [2] packets directly into the communications line rather than having to store the data in a local file.

Second, in the manipulation of large quantities of (e.g., experimental) data, it is useful to be able to tag aggregates of data (whether numbers, records, lines, or whole files) with a statistically unique “key,” its CRC. Aggregates of any size can then be compared for identity by comparing only their short CRC keys. Differing keys imply nonidentical records. Identical keys imply, to high statistical certainty, identical records. If you can’t tolerate the very small probability of being wrong, you can do a full comparison of the records when the keys are identical. When there is a possibility of files or data records being inadvertently or irresponsibly modified (for example, by a computer virus), it is useful to have their prior CRCs stored externally on a physically secure medium, like a floppy disk.

Sometimes CRCs can be used to compress data as it is recorded. If identical data records occur frequently, one can keep sorted in memory the CRCs of previously encountered records. A new record is archived in full if its CRC is different, otherwise only a pointer to a previous record need be archived. In this application one might desire a 4- or 8-byte CRC, to make the odds of mistakenly discarding a different data record be tolerably small; or, if previous records can be randomly accessed, a full comparison can be made to decide whether records with identical CRCs are in fact identical.

Now let us briefly discuss the theory of CRCs. After that, we will give implementations of various (related) CRCs that are used by the official or de facto standard protocols [1-3] listed in the accompanying table.

The mathematics underlying CRCs is “polynomials over the integers modulo 2.” Any binary message can be thought of as a polynomial with coefficients 0 and 1. For example, the message “1100001101” is the polynomial $x^9 + x^8 + x^3 + x^2 + 1$. Since 0 and 1 are the only integers modulo 2, a power of x in the polynomial is either present (1) or absent (0). Polynomials over the integers modulo 2 are a so-called unique factorization domain. This means that any polynomial has a unique factorization into so-called irreducible or “primitive” polynomials — analogs of the prime integers. The polynomial $x^2 + x + 1$ is primitive, while the polynomial $x^2 + 1$ is not: $x^2 + 1 = (x + 1)(x + 1)$. (Remember that the integer arithmetic is done modulo two!) A related theorem says that if a polynomial p divides another polynomial q , then it must divide at least one of q ’s factors.

An M -bit long CRC is based on a particular primitive polynomial of degree M , called the generator polynomial. The choice of which primitive polynomial to use is only a matter of convention. For 16-bit CRC’s, the CCITT (Comité Consultatif International Télégraphique et Téléphonique) has anointed the “CCITT polynomial,” which is $x^{16} + x^{12} + x^5 + 1$. This polynomial is used by all of the

Conventions and Test Values for Various CRC Protocols						
	icrc args		Test Values (C_2C_1 in hex)		Packet	
Protocol	jinit	jrev	T	CatMouse987654321	Format	CRC
XMODEM	0	1	1A71	E556	$S_1S_2 \dots S_N C_2 C_1$	0
X.25	255	-1	1B26	F56E	$S_1S_2 \dots S_N \overline{C_1 C_2}$	FOB8
(no name)	255	-1	1B26	F56E	$S_1S_2 \dots S_N C_1 C_2$	0
SDLC (IBM)	same as X.25					
HDLC (ISO)	same as X.25					
CRC-CCITT	0	-1	14A1	C28D	$S_1S_2 \dots S_N C_1 C_2$	0
(no name)	0	-1	14A1	C28D	$S_1S_2 \dots S_N \overline{C_1 C_2}$	FOB8
Kermit	same as CRC-CCITT			see Notes		
Notes: Overbar denotes bit complement. $S_1 \dots S_N$ are character data. C_1 is CRC's least significant 8 bits, C_2 is its most significant 8 bits, so $CRC = 256C_2 + C_1$ (shown in hex). Kermit (block check level 3) sends the CRC as 3 printable ASCII characters (sends value +32). These contain, respectively, 4 most significant bits, 6 middle bits, 6 least significant bits.						

protocols listed in the table. Another common choice is the “CRC-16” polynomial $x^{16} + x^{15} + x^2 + 1$, which is used for EBCDIC messages in IBM's BISYNCH [1]. A common 12-bit choice, “CRC-12,” is $x^{12} + x^{11} + x^3 + x + 1$. A common 32-bit choice, “AUTODIN-II,” is $x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^8 + x^7 + x^5 + x^4 + x^2 + x + 1$. For a table of some other primitive polynomials, see §7.4.

Given the generator polynomial G of degree M (which can be written either in polynomial form or as a bit-string, e.g., 10001000000100001 for CCITT), here is how you compute the CRC for a sequence of bits S : First, multiply S by x^M , that is, append M zero bits to it. Second divide — by long division — G into Sx^M . Keep in mind that the subtractions in the long division are done modulo 2, so that there are never any “borrows”: Modulo 2 subtraction is the same as logical exclusive-or (XOR). Third, ignore the quotient you get. Fourth, when you eventually get to a remainder, it is the CRC, call it C . C will be a polynomial of degree $M - 1$ or less, otherwise you would not have finished the long division. Therefore, in bit string form, it has M bits, which may include leading zeros. (C might even be all zeros, see below.) See [3] for a worked example.

If you work through the above steps in an example, you will see that most of what you write down in the long-division tableau is superfluous. You are actually just left-shifting sequential bits of S , from the right, into an M -bit register. Every time a 1 bit gets shifted off the left end of this register, you zap the register by an XOR with the M low order bits of G (that is, all the bits of G except its leading 1). When a 0 bit is shifted off the left end you don't zap the register. When the last bit that was originally part of S gets shifted off the left end of the register, what remains is the CRC.

You can immediately recognize how efficiently this procedure can be implemented in hardware. It requires only a shift register with a few hard-wired XOR taps into it. That is how CRCs are computed in communications devices, by a single chip (or small part of one). In software, the implementation is not so elegant, since

bit-shifting is not generally very efficient. One therefore typically finds (as in our implementation below) table-driven routines that pre-calculate the result of a bunch of shifts and XORs, say for each of 256 possible 8-bit inputs [4].

We can now see how the CRC gets its ability to detect all errors in M consecutive bits. Suppose two messages, S and T , differ only within a frame of M bits. Then their CRCs differ by an amount that is the remainder when G is divided into $(S - T)x^M \equiv D$. Now D has the form of leading zeros (which can be ignored), followed by some 1's in an M -bit frame, followed by trailing zeros (which are just multiplicative factors of x). Since factorization is unique, G cannot possibly divide D : G is primitive of degree M , while D is a power of x times a factor of (at most) degree $M - 1$. Therefore S and T have inevitably different CRCs.

In most protocols, a transmitted block of data consists of some N data bits, directly followed by the M bits of their CRC (or the CRC XORed with a constant, see below). There are two equivalent ways of validating a block at the receiving end. Most obviously, the receiver can compute the CRC of the data bits, and compare it to the transmitted CRC bits. Less obviously, but more elegantly, the receiver can simply compute the CRC of the total block, with $N + M$ bits, and verify that a result of zero is obtained. Proof: The total block is the polynomial $Sx^M + C$ (data left-shifted to make room for the CRC bits). The definition of C is that $Sx^m = QG + C$, where Q is the discarded quotient. But then $Sx^M + C = QG + C + C = QG$ (remember modulo 2), which is a perfect multiple of G . It remains a multiple of G when it gets multiplied by an additional x^M on the receiving end, so it has a zero CRC, q.e.d.

A couple of small variations on the basic procedure need to be mentioned [1,3]: First, when the CRC is computed, the M -bit register need not be initialized to zero. Initializing it to some other M -bit value (e.g., all 1's) in effect prefaces all blocks by a phantom message that would have given the initialization value as its remainder. It is advantageous to do this, since the CRC described thus far otherwise cannot detect the addition or removal of any number of initial zero bits. (Loss of an initial bit, or insertion of zero bits, are common "clocking errors.") Second, one can add (XOR) any M -bit constant K to the CRC before it is transmitted. This constant can either be XORed away at the receiving end, or else it just changes the expected CRC of the whole block by a known amount, namely the remainder of dividing G into Kx^M . The constant K is frequently "all bits," changing the CRC into its ones complement. This has the advantage of detecting another kind of error that the CRC would otherwise not find: deletion of an initial 1 bit in the message with spurious insertion of a 1 bit at the end of the block.

The accompanying function `icrc` implements the above CRC calculation, including the possibility of the mentioned variations. Input to the function is the starting address of an array of characters, and the length of that array. (In practice, FORTRAN allows you to use the address of *any* data structure; `icrc` will treat it as a byte array.) Output is in both of two formats. The function value returns the CRC as a 4-byte integer in the range 0 to 65535. The character array `crc`, of length 2, returns the CRC as two 8-bit characters. `icrc` has two "switch" arguments that specify variations in the CRC calculation. A zero or positive value of `jinit` causes the 16-bit register to have each byte initialized with the value `jinit`. A negative value of `jrev` causes each input character to be interpreted as its bit-reverse image, and a similar bit reversal to be done on the output CRC. You do not have to understand this; just use the values of `jinit` and `jrev` specified in the table.

(If you *insist* on knowing, the explanation is that serial data ports send characters *least-significant bit first* (!), and many protocols shift bits into the CRC register in exactly the order received.) The table shows how to construct a block of characters from the input array and output CRC of `icrc`. You should not need to do any additional bit-reversal outside of `icrc`.

The switch `jinit` has one additional use: When negative it causes the input value of the array `crc` to be used as initialization of the register. If `crc` is unmodified since the last call to `icrc`, this in effect appends the current input array to that of the previous call or calls. Use this feature, for example, to build up the CRC of a whole file a line at a time, without keeping the whole file in memory.

At initialization, the routine `icrc` figures out the order in which the bytes occur when a 4-byte character array is equivalenced to a 4-byte integer. This is not strictly portable FORTRAN, but it should work on all machines with 32-bit word lengths. `icrc` is loosely based on a more portable C function in [4], a good place to turn if you have trouble running the program here.

Here is how to understand the operation of `icrc`: First look at the function `icrc1`. This incorporates one input character into a 16-bit CRC register. The only trick used is that character bits are XORED into the most significant bits, eight at a time, instead of being fed into the least significant bit, one bit at a time, at the time of the register shift. This works because XOR is associative and commutative — we can feed in character bits *any* time before they will determine whether to zap with the generator polynomial. (The decimal constant 4129 has the generator's bits in it.)

```

FUNCTION icrc1(crc,onech,ib1,ib2,ib3)
INTEGER icrc1,ib1,ib2,ib3
    Given a remainder up to now, return the new CRC after one character is added. This routine
    is functionally equivalent to icrc(, ,1,-1,1), but slower. It is used by icrc to initialize
    its table.
INTEGER i,ichr,ireg
CHARACTER*1 onech,crc(4),creg(4)
EQUIVALENCE (creg,ireg)
ireg=0
creg(ib1)=crc(ib1)           Here is where the character is folded into the register.
creg(ib2)=char(ieor(ichar(crc(ib2)),ichar(onech)))
do 11 i=1,8                 Here is where 8 one-bit shifts, and some XORs with the gen-
    ichr=ichar(creg(ib2))     erator polynomial, are done.
    ireg=ireg+ireg
    creg(ib3)=char(0)
    if(ichr.gt.127)ireg=ieor(ireg,4129)
enddo 11
icrc1=ireg
return
END

```

Now look at `icrc`. There are two parts to understand, how it builds a table when it initializes, and how it uses that table later on. Go back to thinking about a character's bits being shifted into the CRC register from the least significant end. The key observation is that while 8 bits are being shifted into the register's low end, all the generator zapping is being determined by the bits already in the high end. Since XOR is commutative and associative, all we need is a table of the result of all this zapping, for each of 256 possible high-bit configurations. Then we can play catch-up and XOR an input character into the result of a lookup into this table. The routine makes repeated use of an equivalenced 4-byte integer and 4-byte

character array to get at different 8-bit chunks. The only other content to `icrc` is the construction at initialization time of an 8-bit bit-reverse table from the 4-bit table stored in `it`, and the logic associated with doing the bit reversals. References [4-6] give further details on table-driven CRC computations.

```

FUNCTION icrc(crc,bufptr,len,jinit,jrev)
INTEGER icrc,jinit,jrev,len
CHARACTER*1 bufptr(*),crc(2)
C USES icrc1
  Computes a 16-bit Cyclic Redundancy Check for an array bufptr of length len bytes,
  using any of several conventions as determined by the settings of jinit and jrev (see
  accompanying table). The result is returned both as an integer icrc and as a 2-byte array
  crc. If jinit is negative, then crc is used on input to initialize the remainder register, in
  effect concatenating bufptr to the previous call.
INTEGER ich,init,ireg,j,icrctb(0:255),it(0:15),icrc1,ib1,ib2,ib3
CHARACTER*1 creg(4),rchr(0:255)
SAVE icrctb,rchr,init,it,ib1,ib2,ib3
EQUIVALENCE (creg,ireg)      Used to get at the 4 bytes in an integer.
DATA it/0,8,4,12,2,10,6,14,1,9,5,13,3,11,7,15/, init /0/
  Table of 4-bit bit-reverses, and flag for initialization.
if (init.eq.0) then          Do we need to initialize tables?
  init=1
  ireg=256*(256*ichar('3')+ichar('2'))+ichar('1')
  do 11 j=1,4                Figure out which component of creg addresses which
    if (creg(j).eq.'1') ib1=j   byte of ireg.
    if (creg(j).eq.'2') ib2=j
    if (creg(j).eq.'3') ib3=j
  enddo 11
  do 12 j=0,255              The two tables are: CRCs of all characters, and bit-reverses
    ireg=j*256                of all characters.
    icrctb(j)=icrc1(creg,char(0),ib1,ib2,ib3)
    ich=it(mod(j,16))*16+it(j/16)
    rchr(j)=char(ich)
  enddo 12
endif
if (jinit.ge.0) then        Initialize the remainder register.
  crc(1)=char(jinit)
  crc(2)=char(jinit)
else if (jrev.lt.0) then    If not initializing, do we reverse the register?
  ich=ichar(crc(1))
  crc(1)=rchr(ichar(crc(2)))
  crc(2)=rchr(ich)
endif
do 13 j=1,len              Main loop over the characters in the array.
  ich=ichar(bufptr(j))
  if(jrev.lt.0)ich=ichar(rchr(ich))
  ireg=icrctb(ieor(ich,ichar(crc(2))))
  crc(2)=char(ieor(ichar(creg(ib2)),ichar(crc(1))))
  crc(1)=creg(ib1)
enddo 13
if (jrev.ge.0) then        Do we need to reverse the output?
  creg(ib1)=crc(1)
  creg(ib2)=crc(2)
else
  creg(ib2)=rchr(ichar(crc(1)))
  creg(ib1)=rchr(ichar(crc(2)))
  crc(1)=creg(ib1)
  crc(2)=creg(ib2)
endif
icrc=ireg
return
END

```

World Wide Web sample page from NUMERICAL RECIPES IN FORTRAN 77: THE ART OF SCIENTIFIC COMPUTING (ISBN 0-521-43064-X)
 Copyright (C) 1988-1992 by Cambridge University Press. Programs Copyright (C) 1988-1992 by Numerical Recipes Software.
 Permission is granted for internet users to make one paper copy for their own personal use. Further reproduction, or any copying of machine-readable files (including this one), to any server computer, is strictly prohibited. To order Numerical Recipes books, diskettes, or CDROMs visit website <http://www.nr.com> or call 1-800-872-7423 (North America only), or send email to trade@cup.cam.ac.uk (outside North America).

What if you need a 32-bit checksum? For a true 32-bit CRC, you will need to rewrite the routines given to work with a longer generating polynomial. For example, $x^{32} + x^7 + x^5 + x^3 + x^2 + x + 1$ is primitive modulo 2, and has nonleading, nonzero bits only in its least significant byte (which makes for some simplification). The idea of table lookup on only the most significant byte of the CRC register goes through unchanged. Pay attention to the fact that FORTRAN does not have unsigned integers, so half of your CRCs will appear to be negative in integer format.

If you do not care about the M -consecutive bit property of the checksum, but rather only need a statistically random 32 bits, then you can use `icrc` as given here: Call it once with `jrev = 1` to get 16 bits, and *again* with `jrev = -1` to get another 16 bits. The internal bit reversals make these two 16-bit CRCs in effect totally independent of each other.

Other Kinds of Checksums

Quite different from CRCs are the various techniques used to append a decimal “check digit” to numbers that are handled by human beings (e.g., typed into a computer). Check digits need to be proof against the kinds of highly structured errors that humans tend to make, such as transposing consecutive digits. Wagner and Putter [7] give an interesting introduction to this subject, including specific algorithms.

Checksums now in widespread use vary from fair to poor. The 10-digit ISBN (International Standard Book Number) that you find on most books, including this one, uses the check equation

$$10d_1 + 9d_2 + 8d_3 + \cdots + 2d_9 + d_{10} = 0 \pmod{11} \quad (20.3.1)$$

where d_{10} is the right-hand check digit. The character “X” is used to represent a check digit value of 10. Another popular scheme is the so-called “IBM check,” often used for account numbers (including, e.g., MasterCard). Here, the check equation is

$$2\#d_1 + d_2 + 2\#d_3 + d_4 + \cdots = 0 \pmod{10} \quad (20.3.2)$$

where $2\#d$ means, “multiply d by two and add the resulting decimal digits.” United States banks code checks with a 9-digit processing number whose check equation is

$$3a_1 + 7a_2 + a_3 + 3a_4 + 7a_5 + a_6 + 3a_7 + 7a_8 + a_9 = 0 \pmod{10} \quad (20.3.3)$$

The bar code put on many envelopes by the U.S. Postal Service is decoded by removing the single tall marker bars at each end, and breaking the remaining bars into 6 or 10 groups of five. In each group the five bars signify (from left to right) the values 7,4,2,1,0. Exactly two of them will be tall. Their sum is the represented digit, except that zero is represented as 7 + 4. The 5- or 9-digit Zip Code is followed by a check digit, with the check equation

$$\sum d_i = 0 \pmod{10} \quad (20.3.4)$$

None of these schemes is close to optimal. An elegant scheme due to Verhoeff is described in [7]. The underlying idea is to use the ten-element *dihedral group* D_5 ,

which corresponds to the symmetries of a pentagon, instead of the cyclic group of the integers modulo 10. The check equation is

$$a_1 * f(a_2) * f^2(a_3) * \dots * f^{n-1}(a_n) = 0 \quad (20.3.5)$$

where $*$ is (noncommutative) multiplication in D_5 , and f^i denotes the i th iteration of a certain fixed permutation. Verhoeff's method finds *all* single errors in a string, and *all* adjacent transpositions. It also finds about 95% of twin errors ($aa \rightarrow bb$), jump transpositions ($acb \rightarrow bca$), and jump twin errors ($aca \rightarrow bcb$). Here is an implementation:

```

LOGICAL FUNCTION decchk(string,n,ch)
INTEGER n
CHARACTER string*(*),ch*1
    Decimal check digit computation or verification. Returns as ch a check digit for appending
    to string(1:n), that is, for storing into string(n+1:n+1). In this mode, ignore the
    returned logical value. If string(1:n) already ends with a check digit (string(n:n)),
    returns the function value .true. if the check digit is valid, otherwise .false. In this
    mode, ignore the returned value of ch. Note that string and ch contain ASCII characters
    corresponding to the digits 0-9, not byte values in that range. Other ASCII characters are
    allowed in string, and are ignored in calculating the check digit.
INTEGER ij(10,10),ip(10,8),i,j,k,m
SAVE ij,ip                                Group multiplication and permutation tables.
DATA ip/0,1,2,3,4,5,6,7,8,9,1,5,7,6,2,8,3,0,9,4,
* 5,8,0,3,7,9,6,1,4,2,8,9,1,6,0,4,3,5,2,7,9,4,5,3,1,2,6,8,7,0,
* 4,2,8,6,5,7,3,9,0,1,2,7,9,3,8,0,6,4,1,5,7,0,4,6,9,1,3,2,5,8/,
* ij/0,1,2,3,4,5,6,7,8,9,1,2,3,4,0,9,5,6,7,8,2,3,4,0,1,8,9,5,6,
* 7,3,4,0,1,2,7,8,9,5,6,4,0,1,2,3,6,7,8,9,5,5,6,7,8,9,0,1,2,3,
* 4,6,7,8,9,5,4,0,1,2,3,7,8,9,5,6,3,4,0,1,2,8,9,5,6,7,2,3,4,0,
* 1,9,5,6,7,8,1,2,3,4,0/
k=0
m=0
do 11 j=1,n                                Look at successive characters.
    i=ichar(string(j:j))
    if (i.ge.48.and.i.le.57)then            Ignore everything except digits.
        k=ij(k+1,ip(mod(i+2,10)+1,mod(m,8)+1)+1)
        m=m+1
    endif
enddo 11
decchk=(k.eq.0)
do 12 i=0,9                                Find which appended digit will check properly.
    if (ij(k+1,ip(i+1,mod(m,8)+1)+1).eq.0) goto 1
enddo 12
1 ch=char(i+48)                            Convert to ASCII.
return
end

```

CITED REFERENCES AND FURTHER READING:

- McNamara, J.E. 1982, *Technical Aspects of Data Communication*, 2nd ed. (Bedford, MA: Digital Press). [1]
- da Cruz, F. 1987, *Kermit, A File Transfer Protocol* (Bedford, MA: Digital Press). [2]
- Morse, G. 1986, *Byte*, vol. 11, pp. 115–124 (September). [3]
- LeVan, J. 1987, *Byte*, vol. 12, pp. 339–341 (November). [4]
- Sarwate, D.V. 1988, *Communications of the ACM*, vol. 31, pp. 1008–1013. [5]
- Griffiths, G., and Stones, G.C. 1987, *Communications of the ACM*, vol. 30, pp. 617–620. [6]
- Wagner, N.R., and Putter, P.S. 1989, *Communications of the ACM*, vol. 32, pp. 106–110. [7]

20.4 Huffman Coding and Compression of Data

A lossless data compression algorithm takes a string of symbols (typically ASCII characters or bytes) and translates it *reversibly* into another string, one that is *on the average* of shorter length. The words “on the average” are crucial; it is obvious that no reversible algorithm can make all strings shorter — there just aren’t enough short strings to be in one-to-one correspondence with longer strings. Compression algorithms are possible only when, on the input side, some strings, or some input symbols, are more common than others. These can then be encoded in fewer bits than rarer input strings or symbols, giving a net average gain.

There exist many, quite different, compression techniques, corresponding to different ways of detecting and using departures from equiprobability in input strings. In this section and the next we shall consider only *variable length codes* with *defined word* inputs. In these, the input is sliced into fixed units, for example ASCII characters, while the corresponding output comes in chunks of variable size. The simplest such method is Huffman coding [1], discussed in this section. Another example, *arithmetic compression*, is discussed in §20.5.

At the opposite extreme from defined-word, variable length codes are schemes that divide up the *input* into units of variable length (words or phrases of English text, for example) and then transmit these, often with a fixed-length output code. The most widely used code of this type is the Ziv-Lempel code [2]. References [3-6] give the flavor of some other compression techniques, with references to the large literature.

The idea behind Huffman coding is simply to use shorter bit patterns for more common characters. We can make this idea quantitative by considering the concept of *entropy*. Suppose the input alphabet has N_{ch} characters, and that these occur in the input string with respective probabilities p_i , $i = 1, \dots, N_{ch}$, so that $\sum p_i = 1$. Then the fundamental theorem of information theory says that strings consisting of independently random sequences of these characters (a conservative, but not always realistic assumption) require, on the average, at least

$$H = - \sum p_i \log_2 p_i \quad (20.4.1)$$

bits per character. Here H is the entropy of the probability distribution. Moreover, coding schemes exist which approach the bound arbitrarily closely. For the case of equiprobable characters, with all $p_i = 1/N_{ch}$, one easily sees that $H = \log_2 N_{ch}$, which is the case of no compression at all. Any other set of p_i ’s gives a smaller entropy, allowing some useful compression.

Notice that the bound of (20.4.1) would be achieved if we could encode character i with a code of length $L_i = -\log_2 p_i$ bits: Equation (20.4.1) would then be the average $\sum p_i L_i$. The trouble with such a scheme is that $-\log_2 p_i$ is not generally an integer. How can we encode the letter “Q” in 5.32 bits? Huffman coding makes a stab at this by, in effect, approximating all the probabilities p_i by integer powers of $1/2$, so that all the L_i ’s are integral. If all the p_i ’s are in fact of this form, then a Huffman code does achieve the entropy bound H .

The construction of a Huffman code is best illustrated by example. Imagine a language, Vowellish, with the $N_{ch} = 5$ character alphabet A, E, I, O, and U, occurring with the respective probabilities 0.12, 0.42, 0.09, 0.30, and 0.07. Then the construction of a Huffman code for Vowellish is accomplished in the following table:

World Wide Web sample page from NUMERICAL RECIPES IN FORTRAN 77: THE ART OF SCIENTIFIC COMPUTING (ISBN 0-521-43064-X)
 Copyright (C) 1988-1992 by Cambridge University Press. Programs Copyright (C) 1988-1992 by Numerical Recipes Software.
 Permission is granted for internet users to make one paper copy for their own personal use. Further reproduction, or any copying of machine-readable files (including this one), to any server computer, is strictly prohibited. To order Numerical Recipes books, diskettes, or CDROMs visit website <http://www.nr.com> or call 1-800-872-7423 (North America only), or send email to trade@cup.cam.ac.uk (outside North America).