

A RAPID ENTROPY-CODING ALGORITHM

WM. DOUGLAS WITHERS

Department of Mathematics
United States Naval Academy
Annapolis, MD 21402
and
Pegasus Imaging Corporation

ABSTRACT. We present a new algorithm for entropy coding, competitive in both speed and compression efficiency with the Q-coder and QM-coder.

key words and phrases: entropy coding, data compression, algorithms

Copyright © 1996 by Wm. Douglas Withers

1. ENTROPY CODING

Truth is the most valuable thing we have. Let us economize it.

Mark Twain

The problem of compressing digital data of a certain type (text, images, audio data, whatever) can be decoupled into two subproblems: *modeling* and *entropy coding*. Whatever the given data may represent in the real world, in digital form it exists as a sequence of symbols, such as bits. The modeling problem is to choose a suitable symbolic representation for the data and to predict for each symbol of the representation the probability that it takes each of the allowable values for that symbol. The entropy-coding problem is to code each symbol as compactly as possible, given this knowledge of probabilities. (In the realm of lossy compression, there is a third subproblem: evaluating the relative importance of various kinds of errors.)

For example, suppose we want to transmit messages composed of the four letters **a**, **b**, **c**, and **d**. A straightforward scheme for coding these messages in bits would be to represent **a** by “00”, **b** by “01”, **c** by “10” and **d** by “11”. However, suppose we know that for any letter of the message (independent of all other letters), **a** occurs with probability .5, **b** occurs with probability .25, and **c** or **d** occur with probability .125 each. Then we might choose a shorter representation for **a**, at the necessary cost of accepting longer representations for the other letters. We could represent **a** by “0”, **b** by “10”, **c** by “110”, and **d** by “111”. This representation is more compact on average than the first one; indeed, it is the most compact representation possible (though not uniquely so). In this simple example, the modeling part of the problem is determining the probabilities for each possible symbol value; the entropy-coding part of the problem is determining the representations in bits from those probabilities; let us emphasize that *the probabilities associated with the symbol values play a fundamental role in entropy coding*.

In general, entropy coding is an abstract problem weakly related to the type of data being compressed, while the modeling aspect of data compression depends intimately on the type of data being compressed. Entropy coding is well understood theoretically—known algorithms provide the greatest compression possible *for a given modeling method*—while for many real-world types of data, the modeling issue is as yet mysterious and only somewhat tractable. This article treats only entropy coding and the most rudimentary aspect of modeling; thus it provides a partial solution to almost any data compression problem but a complete solution to none.

One well-known method of entropy coding is *Huffman coding*, which yields an optimal coding provided all symbol-value probabilities are integer powers of .5. Another method, yielding optimal compression performance for any set of probabilities, is *arithmetic coding*. In spite of the superior compression given by arithmetic coding, so far it has not been a dominant presence in real data-compression applications. This is most likely due to concerns over speed and complexity, as well as patent issues; a rapid, simple algorithm for arithmetic coding is therefore potentially very useful.

An algorithm is known which allows rapid encoding and decoding in a fashion akin to arithmetic coding. Developed by G. G. Langdon, J. L. Mitchell, W. B. Pennebaker, and J. J. Rissanen [4], [6], [7], it is known as the *Q-coder*. The *QM-coder* is a subsequent variant.

However, these algorithms being protected by patents, new algorithms with competitive performance continue to be of interest. The algorithm described here is one such; we call it the ELS-coder (for *Entropy Logarithmic-Scale*). Though the ELS-coder is likewise protected by a pending patent, Pegasus Imaging Corporation intends to license it on a royalty-free basis.

Due to space limitations, we do not provide a general introduction to data compression theory, nor a description of other algorithms for entropy coding. Witten and Cleary [5] provides an introduction to the “standard” arithmetic compression algorithm. Pennebaker and Mitchell [3] provides a good introduction to the operation of the QM-coder. Nelson [1] and Cover and Thomas [2] are general practical and theoretical introductions to data compression.

2. THE DECODING ALGORITHM

I were better to be eaten to death with rust than to be scoured to nothing with perpetual motion.

Shakespeare

The ELS-coder works only with a binary alphabet $\{0, 1\}$. One can certainly encode symbols from larger alphabets; but they must be converted to a binary format first. The example programs COMPRESS and DECOMPRESS illustrate how this can be done. The necessity for this conversion is a disadvantage, but restriction to a binary alphabet facilitates rapid coding and rapid probability estimation.

The ELS-coder was developed with an eye toward a particular application which would generally involve compressing a dataset once and subsequently decompressing it many times. Thus decoding speed was of somewhat greater concern to us than encoding speed. We therefore designed the decoder first, and then designed the encoder to fit. We likewise describe the decoder first.

Let us consider first the procedures given in Figure 1. (We use the C programming language throughout to describe algorithms.) This decoder does not as yet provide any compression; it simply extracts bits from a file (and not in the most efficient manner, either). The most-significant bits of each byte in the file are extracted first. We give it here purely as an illustration of our paradigm for the operation of a decoder. (Note: throughout this article, we assume that a **char** has eight bits of precision, a **short** has sixteen bits of precision, and a **long** has thirty-two bits of precision.)

In order for the decoder to hold data, it must be furnished with several states. It is convenient to represent the state of the decoder by two components: one (b) indicating the quantity of data (in bits) held by the decoder and another (x) indicating the content of that data. When called upon to decode a symbol (the procedure *decode_bit()*), the decoder is in one of several possible states. One subset of these states represents the symbol value **0**, and another represents the symbol value **1**. The decoder must determine to which subset its state belongs (by comparing x to *Threshold[b]*). It then transits to a new state to reflect the depletion of its store of data (by decrementing b and, if necessary, decreasing x). At times (when b reaches the value 0) the decoder must replenish its store of data by reading from an external file (the procedure *decode_import()*). This *import* operation modifies both x and b .

```

int b;
int Threshold[10] = {0, 1, 2, 4, 8, 16, 32, 64, 128, 256};
int A[ ] = Threshold + 1; /* For expository purposes only */
FILE *external_in;
unsigned char x;

int decode_bit() {
    if (x ≥ Threshold[b]) {
        x − = Threshold[b];
        if (− − b ≤ 0)
            decode_import();
        return 1;
    } else {
        if (− − b ≤ 0)
            decode_import();
        return 0;
    }
}

void decode_import() {
    x = fgetc(external_in);
    b + = 8;
}

```

FIGURE 1. A pair of procedures which extract bits from a file with no compression involved.

A few enhancements to this simple example will produce the ELS decoding algorithm. We first call attention to the table $A[\]$, not directly used by the decoder but a convenient reference for discussion. In Figure 1, $A[b]$ gives the number of allowable states of a system holding b bits of data. Thus $A[b] = 2^b$. There are $A[b]$ allowable values for x at any time. The insight underpinning arithmetic coding is that this relationship between the number of bits of data in a system and the number of allowable states remains valid even when b is not an integer and 2^b not a power of two.

Note also: we apportioned the allowable values corresponding to **0** and **1**, respectively, so that all of the former values were less than all of the latter. Thus we can determine to which subset of values the value of x belongs by comparing it to $Threshold[b]$, which is the minimum value of x corresponding to **1**. When the symbol value **0** is decoded, the number of allowable values for x necessarily becomes the value of $Threshold[b]$ before decoding the symbol.

Figure 2 describes a full-fledged ELS-decoder, although at a crude level of accuracy. (The procedures for initializing and terminating the decoding process are not shown; straightforward, they can be found in the accompanying appendices.) Let us note the

```

#define F 15

unsigned short A[2 * F] = {1, 2, 3, 4, 5, 7, 10, 14, 20, 28, 41, 59, 85, 123, 177, 256, 371,
536, 776, 1123, 1625, 2353, 3405, 4928, 7132, 10321, 14938, 21619, 31288, 45283};

struct{
    unsigned short *Threshold;
    short c0;
    short c1;
    } Ladder[3] = {{A + 14, 1, 4}, {A + 13, 2, 2}, {A + 11, 4, 1}};

unsigned short x;
int j;
FILE *external_in;

int decode_bit(unsigned char rung){
    if (x ≥ Ladder[rung].Threshold[j]){
        x -= Ladder[rung].Threshold[j];
        if ((j -= Ladder[rung].c1) ≤ 0)
            decode_import();
        return 1;
    }else{
        if ((j -= Ladder[rung].c0) ≤ 0)
            decode_import();
        return 0;
    }
}

void decode_import(){
    j += F;
    x <<= 8;
    x |= fgetc(external_in);
}

```

FIGURE 2. The ELS-decoder.

changes from Figure 1 to Figure 2:

First: to compress symbols from a binary alphabet, we must work with quantities of data smaller than a bit. We therefore define a *jot* to be a unit of data equal to $1/F$ of a byte, F being an integer larger than 8. In Figure 2, F is given the value 15; in practice F normally takes much larger values. The decoder in Figure 2 measures data in jots rather than bits.

For example, $A[j]$ now gives the number of allowable values for x for a given number of

jots. This is determined by the same relation as previously: $A[j]$ takes the value $2^{8j/F}$ with appropriate rounding (“appropriate” to be discussed fully later). For example, 23 jots is equal to 23/15 bytes or $8 \cdot 23/15$ bits. The number of corresponding allowable values is $2^{8 \cdot 23/15} \approx 4927.59$. We round this to give a value of 4928 for $A[23]$.

Second: we must now consider the probability associated with the given symbol value. Thus *decode_bit()* now has a parameter *rung* (whose relation to the probability will later be made explicit). This is used as an index into a table *Ladder*[] of structures with three elements: c_0 and c_1 , indicating the number of jots required to code **0** and **1**, respectively; and *Threshold*, which (as before) is the lower bound for allowable values of x corresponding to the symbol value **1**. (The aptness of our terminology *Ladder* and *rung* will become apparent in Section 4, when we discuss probability estimation.)

Third: unlike b in Figure 1, j is not decreased by a single predictable quantity for each symbol decoded. Thus we can not depend on j hitting the value 0 exactly as it is decreased. Therefore we have expanded x from one byte to two; we call the higher-order byte the *working* byte and the lower-order byte the *reserve* byte. We try to keep the reserve byte completely filled with data and the working byte at least partially filled with data—this is equivalent to saying that we maintain at least 256 allowable values for x . We let j indicate the amount of data in the working byte of x , negative values of j indicating that the reserve byte is at least partially depleted. Meaningful values for j range from $-F$ when the decode is completely empty of data (though in actual operation the minimum value attained by j is $1 - F$), to F when the decoder is completely full. The decoder calls *decode_import()* when the value of j dips to zero or lower. Moreover, j is never decreased by more than F in a single call to *decode_bit()*, lest we exhaust the reserve byte.

Fourth: the operation of importing a byte is complicated somewhat by the expansion of x to two bytes. It involves a shift and OR rather than a simple assignment as in Figure 1.

Consider as an example a call to *decode_bit()* with a value of 0 for *rung*, the value of j being 3. This indicates that x contains $F + 3 = 18$ meaningful jots. Note that $A[18] = 776$; thus the value of x must be one of $0, 1, \dots, 775$. If the decoded symbol has the value **0**, then j will be decreased to 2 by subtracting *Ladder*[0]. c_0 , and x will then contain $F + 2 = 17$ meaningful jots. If the decoded symbol has the value **1**, then j will be decreased to -1 by subtracting *Ladder*[0]. c_1 , and x will then contain $F - 1 = 14$ meaningful jots. Note $A[17] = 536$ and $A[14] = 177$; therefore, out of the 776 allowable values for x , the 536 values $0, 1, \dots, 535$ represent **0** and the 177 values $536, 537, \dots, 712$ represent **1**. Suppose the value of x is 600. The first step in decoding the symbol is to compare 600 to 536 (the value of *Ladder*[0].*Threshold*[3]) and see that the symbol has the value **1**. We then subtract 4 from j , giving it the value -1 . The allowable values for x are now $0, 1, \dots, 177$; we subtract 536 from x (making it 64) to bring it within this range.

Since $-1 \leq 0$, we have exhausted the data in the working byte of x ; we call *decode_import()* to import a byte from the external file. Suppose the next byte in the external file has the value 137. We update the value of x to $(64 \cdot 256) + 137 = 16521$ (actually accomplished by shifting and ORing) and update the value of j to $-1 + 15 = 14$. Note that the number of allowable values for x is now $A[14 + 15]$, or 45283.

This is an opportune moment to remark on some fundamental principles for the design

of a decoder according to our paradigm. For example, note that of the 776 allowable values for x , 536 represent **0** and 177 represent **1**. The other $776 - 536 - 177 = 63$ values are wasted. This is a defect of the decoder; ideally every allowable value should represent either **0** or **1**, but the restriction of j and A to integer values makes such waste unavoidable at least some of the time.

Our definition of “allowable” states does not consider the future development of a state. Thus some allowable states may not lead to allowable states in the future; such states are unusable in practice. We call the situation where such states exist a *data leak*. In the presence of a data leak, there are possible states for the decoder which are not allowable. Another characteristic of a data leak is that some possible coded data streams are illegal or redundant.

Data leaks form one of two main sources of coding inefficiency in the ELS-coder, the other being inaccuracy in the proportion of values allocated to **0** and **1**, which will be discussed in Section 4. However, for larger values of F , the inefficiency is quite modest. As discussed in Section 5, a good working value for F is 754; in this case the data leak described causes a coding inefficiency of less than .008 bits per symbol.

A data leak exists when some allowable states of the decoder do not lead to allowable states of the decoder at a later time. Since these allowable states embody the data present in the decoder, a data leak implies that data is somehow dissipating into nothing. The situation is analogous to an engine failing to convert all the energy input into usable energy output. It is a virtue for an engine to convert a large portion of energy input to energy output. On the other hand, one who expects energy output to exceed energy input for any engine has made some sort of fundamental error.

The analogous principle for the ELS-coder (or for any decoder viewed in the light of our paradigm) is that all allowable states of the decoder *must* be derivable from allowable states of the decoder at any *earlier* time. Our name for a violation of this principle is *perpetual motion*. Perpetual motion implies the existence of allowable states which are not possible. Data leaks and perpetual motion are dual evils; while the former is regrettable, the latter must be avoided at all costs.

For example, the components c_0 and c_1 for each entry of the array *Ladder*[] in Figure 2 must satisfy the constraint that

$$A[j - c_0] + A[j - c_1] \leq A[j]$$

for all values of j between 1 and F , inclusive; i.e., of $A[j]$ allowable values for x at any time, those representing **0** and those representing **1** can total to no more than $A[j]$.

Our example decoding process illustrates a second data leak in the decoder, which may be less immediately apparent. This occurs while importing a byte from the external file. Immediately before importing the byte, x has 177 allowable values. Importing a byte makes the number of values available $256 \cdot 177 = 45312$. However, the number of allowable values after importing is specified as 45283; thus 29 allowable values have been lost. In constructing the table $A[]$, we must take care to avoid perpetual motion while importing a byte. Recall that $A[j]$ takes the value $2^{8j/F}$ with “appropriate” rounding: “appropriate” means that $A[j]$ takes the value $2^{8j/F}$ rounded to the nearest integer when $F \leq j < 2F$,

but for $0 \leq j < \mathbf{F}$, to avoid perpetual motion, we calculate $A[j]$ as $(A[j + \mathbf{F}] + 255)/256$ (i.e. rounded *up*).

The choice of the value of *rung* to be used in a call to *decode_bit()* is dictated by the probabilities that the symbol to be decoded will have value **0** or **1**. Let p denote the probability that the symbol takes the value **1**. For a given element of the array *Ladder*[], the expected number of jots used to code the symbol is $c_0(1 - p) + c_1p$. For example, with a value of 0 for *rung*, the expected number of jots is $(1 - p) + 4p = 1 + 3p$; with a value of 1 for *rung*, the expected number of jots is $2(1 - p) + 2p = 2$. For purposes of data compression we would of course prefer the smaller of these two values; we can solve the inequality $1 + 3p \leq 2$ for p to find that the value 0 is to be preferred to 1 for *rung* if $p \leq 1/3$. With similar calculations we find that 1 is the preferred value if $1/3 \leq p \leq 2/3$ and 2 is the preferred value if $2/3 \leq p$. (At the boundary points between these intervals we may equally well choose the preferred value for either side.)

Incidentally, the theoretical optimum compression can be calculated by setting

$$(2.1) \quad c_0 = \mathbf{F} \log_{256}(1 - p), \quad c_1 = \mathbf{F} \log_{256} p;$$

the corresponding expected number of jots

$$-((1 - p) \log_{256}(1 - p) + p \log_{256} p)\mathbf{F}$$

is known in data compression theory as the *entropy* of the symbol (usually measured in bits rather than jots).

All entries of the table *Ladder*[] are calculated so that the values of c_0 and c_1 satisfy the following constraints, termed the *ladder* constraints::

- (1) $A[j - c_0] + A[j - c_1] \leq A[j]$ for any value of j between 0 and $\mathbf{F} - 1$, inclusive (to avoid perpetual motion);
- (2) $c_0 > 0$ and $c_1 > 0$ (so that each symbol value **0** or **1** corresponds to at least some allowable values);
- (3) $c_0 \leq \mathbf{F}$ and $c_1 \leq \mathbf{F}$ (to avoid running beyond the end of the reserve byte while decoding a symbol);
- (4) subject to the above criteria, $A[J] - (A[J - c_0] + A[J - c_1])$ should be minimized (to minimize data leaks).

In all cases *Ladder*[*i*].*Threshold* takes the value *allowable* - *Ladder*[*i*].*c*₀ + **F**. With $\mathbf{F} = 15$ there are but three combinations of c_0 and c_1 satisfying all the above criteria; these yield the three entries of the table *Ladder* in Figure 2.

3. THE ENCODING ALGORITHM

Now for an instant their eyes met in the mirror; and the woman's face he saw there, or seemed to see there, yearned toward him, and was unutterably loving, and compassionate, and yet was resolute in its denial. For it denied him, no matter with what wistful tenderness, or with what wonder at his folly. Just for a moment he seemed to see that; and then he doubted, for Kathleen's lips lifted complaisantly to his, and Kathleen's matter-of-fact face was just as he was used to seeing it.

And thus, with no word uttered, Felix Kennaston understood that his wife must disclaim any knowledge of the sigil of Scoteia, should he be bold enough to speak of it.

James Branch Cabell, *The Cream of the Jest*

Such a moment must be considered a formidable accomplishment from the standpoint of data compression. The reader can most likely cite other occasions when a glance or a handful of words sufficed to carry volumes of meaning. Communication of such high efficiency requires profound empathy between sender and receiver. This is true no less of machines than people; indeed, for machines such a high level of understanding is far more easily attained. One computer program can contain another in entirety, if necessary.

The ELS-coder decoding algorithm has already been described. The encoder uses its knowledge of the decoder's inner workings to create a data stream which will manipulate the decoder into producing the desired sequence of decoded symbols; the encoder plays Rasputin to the decoder's Alexandra.

```

unsigned long x_min;
int j;
int n;
FILE *external_out;

void encode_bit(unsigned char rung, int bit){
    if (bit){
        /* Encode a 1 */
        x_min += Ladder[rung].Threshold[j];
        j -= Ladder[rung].c1;
    }else{
        /* Encode a 0. */
        j -= Ladder[rung].c0;
    }
    if (j ≤ 0)
        encode_export();
}

```

FIGURE 3. ELS-encoder: basic procedure.

Figures 3 and 4 show procedures for encoding compatible with the decoding procedures of Figure 2. The tables $A[\]$ and $Ladder[\]$ and the macro F are identical to those used by the decoder.

In rough terms, the ELS-encoder operates by considering *all* possible coded data streams and gradually eliminating those inconsistent with the current state of the decoder (“current” and other adverbs of time used similarly in this discussion should be understood to refer to position in the data stream rather than actual physical time). For the decoder the “allowable” values of the internal buffer form a convenient reference point for discussion; for the encoder we are concerned with the set of values for the coded data stream *consistent* with the current set of allowable values in the decoder.

As a practical matter, the encoder need not actually consider the entire coded data stream at one time. We can partition the coded data stream at any time into three portions; from end to beginning of the data stream they are: *preactive* bytes, which as yet exert no influence over the current state of the decoder; *active* bytes, which affect the current state of the decoder and have more than one consistent value; and *postactive* bytes, which affect the current state of the decoder and have converged to a single consistent value. Each byte of the coded data stream goes from preactive to active to postactive; the earlier a byte’s position in the stream, the earlier these transitions occur. A byte is not actually moved to the external file until it becomes postactive. Only the active portion of the data stream need be considered at any time.

Since the internal buffer of the decoder contains two bytes, there are always at least two active bytes. The variable n counts the number of active bytes in excess of two. In theory n can take arbitrarily high values, but higher values become exponentially less likely. We number the active bytes $0, 1, \dots, n + 1$ from latest to earliest. The encoder has a variable j matching in value the decoder’s variable j at the same point in the data stream.

Regarding the arithmetic operations used for encoding we can think of the $n + 2$ active bytes as forming a single unsigned number with $n + 2$ bytes of precision, byte 0 being least significant and byte $n + 1$ being most significant. The set of allowable values in the decoder at any time form a continuous range; it can be shown that the consistent values of the active bytes in the encoder at any time likewise form a continuous range. Thus we can describe this range simply by its minimum value m and maximum value M . Each of these is a nonnegative integer with $n + 2$ bytes of precision; we write m_k or M_k to refer specifically to byte k of m or M , respectively. Moreover, since the number of elements of the range is given by $A[F + j]$ only the minimum value need be specified.

Suppose $n > 0$. Consider the most-significant active byte, byte $(n + 1)$. The minimum consistent value for this byte is m_{n+1} and the maximum consistent value is M_{n+1} . Since this byte is active and not yet postactive, it has more than one consistent value; thus $m_{n+1} < M_{n+1}$. The current number of allowable values for the decoder’s internal buffer is $A[j]$; and $M = m + (A[j] - 1)$. Recall that $A[j]$ is a two-byte value. If we consider the operation of adding $(A[j] - 1)$ to m byte-by-byte to obtain M , we see that a carry must occur from byte 1 to byte 2 and on upward to byte $(n + 1)$ in order for byte $(n + 1)$ to take differing values for m and M . We see also that we must have $m_{n+1} + 1 = M_{n+1}$. And furthermore, if $n > 1$, then m_2, m_3, \dots, m_n must all take the maximum possible value 255 in order for the carry to propagate from byte 1 to byte $(n + 1)$. (And hence

M_2, M_3, \dots, M_n all take the minimum possible value 0.)

This is convenient; for we need not separately store all $n + 2$ bytes of m , but can make the four-byte variable x_min serve. Bytes 0, 1, and 2 of x_min represent m_0, m_1 , and (when n is positive) m_2 , respectively. When $n > 1$, then byte 3 of x_min represents m_{n+1} . Since bytes m_2, \dots, m_n all hold the same value, x_min and n together suffice to describe m in entirety. Moreover, most of the arithmetic operations to be performed on m can be performed on x_min in a completely ordinary way. The exceptions can be recognized in the code because they require the manipulation of n .

```

void encode_export() {
    unsigned long diff_bits;

    /* First check for bytes becoming postactive;
    diff_bits marks differences between max and min consistent values */
    j += F;
    diff_bits = (x_min + A[j] - 1) ^ x_min;
    switch (n) {
    default: /* 2 or greater */
        if (diff_bits & 0xFF000000)
            break;
        fputc(x_min >> 24, external_out);
        while (--n > 1)
            fputc((x_min >> 16) & 0xFF, external_out);
    case 1:
        if (diff_bits & 0x00FF0000)
            break;
        fputc((x_min >> 16) & 0xFF, external_out);
        n--;
    case 0:
        if (diff_bits & 0x0000FF00)
            break;
        fputc((x_min >> 8) & 0xFF, externalOut);
        n--;
    }
}
/* Move a byte from preactive to active. */
if (++n > 2)
    x_min = (x_min & 0xFF000000) | ((x_min & 0x0000FFFF) << 8);
else
    x_min <<= 8;
}

```

FIGURE 4. Encoding procedure to export data.

Having concluded these rather lengthy preliminaries, let us now consider the arithmetic operations entailed in encoding a symbol. Recall the sequence of events in the decoder attending the decoding of the symbol value $\mathbf{0}$. Initially, x has one of $A[\mathbf{F} + j]$ values. The decoder compares x to $Ladder[rung].Threshold[j]$; a $\mathbf{0}$ is decoded if x holds the lesser value. Then $Ladder[rung].c_0$ is subtracted from j , reflecting the newly decreased number of allowable values for x . Note that the allowable values eliminated are all taken from the top of the range (those greater than $Ladder[rung].Threshold[j]$). Thus m does not change; and the encoder need not modify x_min or n but only j .

To encode the symbol value $\mathbf{1}$: The chief difference in the sequence of events in the decoder as compared to the case for $\mathbf{0}$ is that those values representing $\mathbf{0}$, numbering $Ladder[rung].Threshold[j]$, are eliminated from the *bottom* of the range of consistent values for the coded data stream. Some additional values may be eliminated from the top of the range as well, if there is a data leak. Thus the encoder, as well as changing the value of j to track its value in the decoder, must add $Ladder[rung].Threshold[j]$ to x_min to raise m . The encoder not representing M directly, the operation of eliminating consistent values from the top of the range takes care of itself. The unusual format by which x_min represents m never becomes an issue in *encode_bit()*.

When the value of j dips to zero or lower, the encoder calls the procedure *encode_export()*. In contrast to *decode_import()*, which invariably reads a single byte from the external file, *encode_export()* writes from zero to several bytes in a single call. One of its tasks is to determine whether the most-significant active bytes have yet converged to a unique consistent value, thus becoming postactive and hence ready to be exported to the encoder's external file. These most-significant bytes may actually converge to a unique consistent value well before *encode_export()* is called, but there is no harm in waiting until then to check. On each call, *encode_export()* moves a single byte from the preactive to the active portion of the coded data stream. The format by which x_min represents m does become an issue in *encode_export()*; we must manipulate n and x_min together.

4. PROBABILITY ESTIMATION

As for a future life, every man must judge for himself between conflicting vague probabilities.

Charles Darwin

The foregoing sections describe an encoder and decoder for entropy coding; however, powerful operations of data modeling can be incorporated in a natural way into the encoder and decoder using the developer's own sophisticated model for the particular type of data involved.

Earlier we commented that the two main sources of compression inefficiency in the ELS-coder are data leaks and misestimates of probabilities. We now provide some rules of thumb for estimating the compression inefficiency due to the latter in any entropy coder.

Data compression theory teaches that the shortest possible expected code length for a binary symbol taking the value $\mathbf{1}$ or $\mathbf{0}$ with probability p or $(1 - p)$, respectively is achieved by using a code of length $\log_{.5} p$ bits for $\mathbf{1}$ and a code of length $\log_{.5}(1 - p)$ bits for $\mathbf{0}$. The resulting optimum expected code length (the entropy of the symbol) is

$[(1-p)\log_{.5}(1-p) + p\log_{.5}p]$ bits. Suppose we estimate p by $p' = p + \epsilon$ with an error of ϵ ; then we would use code lengths for **0** and **1** of $\log_{.5}(1-p-\epsilon)$ and $\log_{.5}(p+\epsilon)$; the resulting expected code length is $[(1-p)\log_{.5}(1-p-\epsilon) + p\log_{.5}(p+\epsilon)]$. Using a power series representation for the logarithm, it can be shown that for small values of ϵ the expected code length is approximately

$$(4.1) \quad \frac{\epsilon^2}{p(1-p)\log 2}$$

bits per symbol greater than optimum.

This depends on p ; at times it is more convenient to work with the *probability angle* α such that $p = \sin^2 \alpha$ and $1-p = \cos^2 \alpha$. Then a small error of η radians in the probability angle entails a compression inefficiency of approximately

$$(4.2) \quad \frac{2\eta^2}{\log 2}$$

bits per symbol—independent of α .

In a real-world application, it can be quite difficult to estimate the probabilities for a particular symbol. One approach is to examine the history of symbols which appeared in similar contexts. (This determination of “context” is part of the modeling problem; the example programs COMPRESS and EXPAND provide an illustration.) On the $(t+1)$ st occurrence of a particular context, we create an estimate p_{t+1} for the probability p that the symbol takes the value **1** by averaging its values over previous occurrences:

$$(4.3) \quad p_{t+1} = \frac{1}{t}(\sigma_1 + \sigma_2 + \cdots + \sigma_t),$$

where σ_u represents the value of the symbol on the u th occurrence of this context.

Such an approach may be improved upon if the value of p shifts over time. For example, in compressing the Bible, one would find that “Moses” occurs more frequently than “Peter” in the Old Testament, while the opposite is true in the New Testament. Therefore we might choose to weight recent occurrences of a given context more heavily. For example, we could use geometrically decreasing weights to create an estimate P_t for p :

$$(4.4) \quad P_{t+1} = s(\sigma_t + (1-s)\sigma_{t-1} + (1-s)^2\sigma_{t-2} + \cdots),$$

s being a parameter between 0 and 1. We have simplified matters by assuming in (4.4) that the context has occurred infinitely many times already; the weights given to occurrences in the distant past are very small anyway. The greater the value of s , the more heavily more recent occurrences of a context are weighted, and thus the more rapidly P_t reflects changes in p . We therefore call s the *speed* of adaption.

Higher speeds of adaptation better recognize shifting probabilities, but there is an attendant disadvantage. Small speeds of adaptation are akin to calculating an average over a large sample, while large values of speed calculate an average over a small sample. Aside from shifts in the value of p , we would expect our estimate to be more accurate as the

sample is larger. Indeed, assuming the symbols σ_u are independent (a reasonable assumption, since dependence indicates an unexploited opportunity for the modeling scheme), and ignoring shifts in the value of p , statistics theory indicates that the mean-squared error in our estimate P_{t+1} for p is given by

$$\epsilon^2 = \frac{s}{2-s} p(1-p),$$

increasing as s increases. In light of formula (4.1), the corresponding coding inefficiency is

$$(4.5) \quad \frac{s}{(2-s) \log 2}$$

bits per symbol (conveniently independent of p).

This dichotomy is of fundamental importance in real-world data compression problems. In estimating probabilities associated with a particular type of data, often one's only recourse is to judge empirically based on the data previously seen. One must decide how much additional importance should be accorded to more recent data. If recent data is weighted heavily, then the accuracy of the estimate suffers because one is estimating from a small sample. If recent data is not weighted heavily, then the accuracy of the estimate may suffer because one's estimate will be slow to reflect shifts in the true probability. The state of the art in modeling many real-world classes of data is such that the probability estimate presented to the entropy coder should often be regarded as rather haphazard.

This gloomy state of affairs is not completely without its consolations. Suppose the probability estimate presented to an entropy coder (of whatever variety) has some significant error. If, due to its own limitations, the entropy coder cannot reproduce the requested probability value exactly, this is as likely to improve compression as to worsen it. Other types of coding inaccuracies may also pale in severity in the face of this fundamental difficulty of probability estimation. Thus the real world can be surprisingly forgiving of minor inefficiencies in an entropy coder, particularly of those related to probability values.

Though more sophisticated than (4.3), formula (4.4) can be more rapidly computed in practice. Note that from (4.4) we can derive:

$$P_{t+1} = s\sigma_t + (1-s)P_t;$$

i.e.,

$$(4.6) \quad P_{t+1} = \begin{cases} (1-s)P_t & \text{if } \sigma_t = \mathbf{0}, \\ (1-s)P_t + s & \text{if } \sigma_t = \mathbf{1}. \end{cases}$$

In software, this can be implemented by a pair of lookup tables giving P_{t+1} as a function of P_t , depending on whether the last symbol for the given context had the value $\mathbf{0}$ or $\mathbf{1}$. We cannot treat P_t as a floating-point value but must restrict it to a discrete set. However, a relatively modest collection of values for P_t can provide excellent coding efficiency. For example, consider 90 values, distributed between 0 and 1 so that the corresponding probability angles are evenly spaced between 0 and 90°. Then the error in probability

```

struct{
    unsigned short *Threshold;
    short c0;
    short c1;
    unsigned char next0;
    unsigned char next1;
    } Ladder[...] = {...};
...
int decode_bit(unsigned char*rung){
    if (x ≥ Ladder[*rung].Threshold[j]){
        x − = Ladder[*rung].Threshold[j];
        if ((j − = Ladder[*rung].c1) ≤ 0)
            decode_import();
        *rung = Ladder[*rung].next1;
        return 1;
    }else{
        if ((j − = Ladder[*rung].c0) ≤ 0)
            decode_import();
        *rung = Ladder[*rung].next0;
        return 0;
    }
}

void encode_bit(unsigned char*rung, intbit){
    if (bit){
        xmin + = Ladder[*rung].Threshold[j];
        j − = Ladder[*rung].c1;
        *rung = Ladder[*rung].next1;
    }else{
        j − = Ladder[*rung].c0;
        *rung = Ladder[*rung].next0;
    }
    if (j ≤ 0)
        encode_export();
}

```

FIGURE 5. Entropy decoding and encoding procedures incorporating probability estimation.

angle need never exceed $.5^\circ$. From (4.2), we find the corresponding coding inefficiency to be 0.0002197 bits per symbol. This is the error due to the discretization of p ; as noted above, this is likely to be insignificant compared to coding inefficiencies from other effects.

Figure 5 illustrates the incorporation of these lookup tables into the ELS-coder. We have

added two elements *next0* and *next1* to the structures making up the table *Ladder*[], indicating the next value of the first argument to the procedures *decode_bit*() and *encode_bit*() if the current symbol has the value **0** or **1**, respectively. Here can be seen the origin of our nomenclature *Ladder*[]; we think of the elements of this table as rungs of a ladder, which we ascend and descend as the probability associated with the context rises and falls. The procedures *decode_bit*() and *encode_bit*() now must modify the value of their first argument, so it is no longer appropriate to pass it by value. The procedures *decode_bit*() and *encode_bit*() now modify the value of the index into *Ladder*[]; thus the parameter *rung* is now of type **unsigned char*** rather than **unsigned char** and is a pointer to this index rather than representing the index directly.

The foregoing assumes that probabilities are updated once for every symbol coded. A slight modification of the described technique is to update the probabilities only when the decoder imports data or the encoder exports data (i.e., in the procedures *decode_import*() and *encode_export*() of our examples). This should provide faster coding, since the update operation is performed less often, but could decrease coding efficiency, since the probabilities are estimated from a more restricted sample. This approach is appealing from the standpoint that probabilities are modified more often when compression is poor, since in that case bytes are imported more frequently. Our name for this variant is *inalacritous* updating, the other method being of course *alacritous*. The appendices illustrate both approaches, selectable by compile switches.

Inalacritous updating requires a different set of look-up tables for the probability ladder, since bytes are more likely to be imported during the coding of certain symbol values than others. The updating formula (4.6) indicates that the probability estimate increases by $\Delta p_1 = s(1 - P_t)$ when the coded symbol has the value **1** and decreases by $\Delta p_0 = -sP_t$ when the coded symbol has the value **0**—note $\Delta p_1 - \Delta p_0 = s$. Thus the ratio of these two step sizes is $\Delta p_1/\Delta p_0 = -(1 - P_t)/P_t$. For a given element of the table *Ladder*[], the probability of importing a byte when **0** is coded is c_0/F and the probability of importing a byte when **1** is coded is c_1/F . To compensate, we therefore adjust the ratio of the two step sizes to

$$\frac{\Delta p'_1}{\Delta p'_0} = \left(\frac{\Delta p_1}{\Delta p_0} \right) \left(\frac{c_0/F}{c_1/F} \right) = -\frac{(1 - P_t)c_0}{P_t c_1}.$$

If we define the speed s to be $\Delta p'_1 - \Delta p'_0$ as previously, then we obtain

$$\Delta p'_0 = \frac{-sP_t c_1}{(1 - P_t)c_0 + P_t c_1},$$

$$\Delta p'_1 = \frac{s(1 - P_t)c_0}{(1 - P_t)c_0 + P_t c_1}.$$

For P_t near 0 or 1, we may need to restrict s to small values lest $P_t + \Delta p'_0 < 0$ or $P_t + \Delta p'_1 > 1$. The appendices include complete examples of both types of probability ladder.

5. THE CODE SAMPLES

The appendices contain four files ELSCODER.H, ELSCODER.C, COMPRESS.C and EXPAND.C. The files ELSCODER.H and ELSCODER.C contain declarations and defini-

tions for an implementation of the ELS-coder, respectively. The files COMPRESS.C and EXPAND.C are sample programs which use the ELS-coder to compress and expand a file using a one-byte model.

The code fragments given as examples in the previous sections were designed more for pedagogical value than utility. The sample programs represent a compromise between these two goals with more emphasis on utility. The reader will note the following changes:

- (1) The variables describing the states of the decoder and encoder have been collected into structures named *d* and *e*, respectively.
- (2) The procedures *decode_import()* and *encode_export()* have been eliminated and the corresponding operations incorporated directly into *els_decode_bit()* and *els_encode_bit()*.
- (3) The value of *F* has been increased from 15 to 754. The value 754 represents a sort of local optimum. Although higher values for *F* generally yield better coding efficiency, the vagaries of the ladder constraints listed in Section 2 dictate that higher compression ratios are attainable for the value 754 than for 755 and subsequent higher values up to around 1508.

The file ELSCODER.C includes straightforward procedures *els_decode_start()* and *els_encode_start()* for initializing and *els_decode_end()* and *els_encode_end()* for terminating encoding and decoding.

One can choose alacritous or inalacritous probability updating in the coder by defining or not defining *ALACRITOUS* when compiling ELSCODER.C.

The sample files also incorporate some enhancements to the basic ELS-coder described in the previous sections:

First: Although the entropy-coding algorithm works with units of data smaller than a bit, most likely the source and destination files must have lengths that are multiples of a byte. Thus a file coded with the ELS-coder is likely to end with a fraction of a byte not containing meaningful data. The sample ELS-coder attempts to make a virtue of necessity by using this empty portion to encode something akin to a checksum. When encoding is concluded, a certain number of consistent values remain in the encoder; these are used to store the value of *j* to whatever precision is possible. This is done automatically in the procedure *els_encode_end()*. If desired, the user can call the function *els_decode_ok()* when finished decoding but before calling *els_decode_end()*; this will verify that the value of *j* matches that sent by the encoder. The ELS-coder is such that any corruption of the compressed data will most likely profoundly alter the course of execution of the decoder; the probability of ending with the same value of *j* is quite small. Of course, this value is only encoded to the precision possible with the fraction of a byte remaining at the end of coding; thus the probability of hitting the correct value by chance ranges from 1 in 255 to 1, depending on what that fraction is.

Second: Recall that our design of the probability state ladder (in both the alacritous and non-alacritous cases) assumes that any given context has already occurred infinitely many times. This assumption is least tenable near the beginning of the coding process. It seems appropriate to use smaller samples to estimate probabilities (i.e., to use greater speeds of adaptation) early in the coding process and larger samples (lesser speeds of adaptation) later in the process. This has been achieved by incorporating *transient* states into the

probability ladder; these are states which can be visited at most once early in the coding process and never again thereafter.

Consider for example the alacritous probability ladder. The initial rung is at index 0 and has probability .5 (for convenience, the index and probability value corresponding to each rung of the probability ladder are listed in a comment preceding the initializer for that rung). State 0 is followed by rung 1 or 2, with probabilities of .265 or .735, respectively. The speed of rung 0 is equal to the difference between these values, or .47. Rungs 1 and 2 are followed by rungs 3,4,5,6; the speed for rungs 1 and 2 is given by, for example, the difference between the probability values for rungs 3 and 4, or .41185. Similarly, rungs 3,4,5,6, are followed by rungs 7, . . . , 14, which are in turn followed by rungs 15, . . . , 30, which are in turn followed by rungs 31, . . . , 62, which are in turn followed by rungs 63, . . . , 126. The speed decreases at each step, from .47 to .41185 to .35371 to .29557 to .23743 to .17928 to .12113. From one of rungs 63, . . . , 126 the coder transits to one of rungs 127, . . . , 254; these form the permanent part of the ladder; these rungs can be visited repeatedly. The speed for the permanent rungs is .063.

The inalacritous probability ladder likewise consists of 127 transient rungs and 128 permanent rungs. The speed ranges from an initial speed of .39 to a final speed of .047.

The values of speed used here are by no means sacred; the best value of speed for the permanent rungs depends on the nature of the data being compressed. Nor is it required to have 127 transient rungs and 128 permanent rungs; we simply chose values totaling to almost 256 in order to get the most mileage out of the **unsigned char** serving as an index into the ladder.

The example shell programs COMPRESS.C and EXPAND.C illustrate the use of various contexts for modeling. In this case the model is a one-byte model. For example, if compressing English text, the model exploits the fact that ‘e’ occurs more frequently than ‘u’, but not the fact that ‘u’ occurs more frequently than ‘e’ following ‘q’; it represents frequencies for each byte value but not relationships between different bytes. This model is described by a state machine, each state holding a probability rung. Each state of the machine corresponds to a particular context. For example, the first bit of a byte is one of these contexts. The last bit of a byte corresponds to 128 different contexts, depending on the values of the preceding seven bits. Any model to be used with the ELS-coder or similar coders handling only binary symbols must be expressed as such a state machine.

This one-byte model requires 255 states (stored in the table *context*[]), organized in a tree fashion similar to the transient rungs of the probability ladder—we number these from 1 to 255. State 1 is used for the first bit of any byte. State 2 is used for the second bit of a byte if the first bit was **0**; state 3 is used for the second bit of a byte if the first bit was **1**. In general, state $2k$ is used for the next bit of a byte if the preceding bit corresponding to state k was **0**, while state $2k + 1$ is used for the next bit of a byte if the preceding bit corresponding to state k was **1**. This arrangement of states also makes for straightforward calculation of state transitions, so storing the values of successor states is unnecessary. Such a tree of probabilities is equivalent to storing the probabilities for all 256 byte values.

The entry *context*[0] has not been used in any of the above; we use it to identify a 257th symbol value serving as an end-of-file marker. We represent the byte value 255 by the sequence of *nine* bits **111111110**; the end-of-file marker is the nine-bit sequence **111111111**.

The state *context*[0] is used for the final bit of this sequence. Both COMPRESS.C and EXPAND.C must take special action in the case of a byte with value 255.

One can compress a file *raw* to one *compressed* by the command

COMPRESS *raw compressed*.

Subsequently one can expand the file *compressed* to one *expanded* by the command

EXPAND *compressed expanded*.

The files *raw* and *expanded* should then be identical.

6. COMPARISON WITH THE Q-CODER

We now point out similarities and differences between the ELS-coder and the Q-coder. These may prove useful for a later study of the Q-coder.

The Q-coder is easily interpreted in terms of our paradigm for a decoder as given in Section 2. (Actually we would claim that *any* decoder can be described by this paradigm, but the description is more transparent for some than for others.) Like the ELS-coder, the Q-coder uses two components to describe the state of the decoder at any time: one, the “augend”, denoted A , describes the amount of data at any time (as does j); the other, denoted X , gives the content of that data (as does x). While j in the ELS-coder directly measures the quantity of data—the corresponding number of allowable states being given by an exponential relation as embodied in the table $A[\]$ —, the augend A in the Q-coder directly measures the number of allowable states in the decoder—the corresponding quantity of data being proportional to the logarithm of A . Most points of difference between the two coders are consequences of this fundamental difference.

Ideally, while decoding a symbol, the proportion of allowable states corresponding to the symbol value $\mathbf{0}$ should match the probability $p_{\mathbf{0}}$ that the symbol takes the value $\mathbf{0}$. Determining the number of allowable states corresponding to each symbol value therefore requires a multiplication or division to be accomplished in some way. In the ELS-coder, the number of allowable states allocated to the symbol value $\mathbf{0}$ is given by an entry of a lookup table *Threshold*[] depending on $p_{\mathbf{0}}$. Since the ELS-coder represents the number of states indirectly, via an exponential table, a simple addition performed on the index into the table is equivalent to a multiplication performed on the number of allowable states. It is interesting to note that in spite of the fundamental role played by probabilities in entropy coding, the ELS-coder operates without representing any probabilities directly. By contrast the Q-coder represents each probability state by a single number (called a *Qe-value*) proportional to the probability. The Q-coder approximates the desired multiplication operation either by subtracting the Qe-value directly from the augend A or by assigning the Qe-value to A .

This approximation provides (perhaps surprising) coding efficiency, but it does impose some constraints on the design of the Q-coder:

- (1) Maintaining sufficient accuracy of approximation requires that the number of allowable states in the decoder be restricted to a much smaller range than that used

in the ELS-coder; in general the number of allowable states is not permitted to vary by more than a factor of two (whereas in the ELS-coder as described here, this number varies by a factor of 256). As a consequence, data must be imported much more frequently in the Q-coder than in the ELS-coder. Moreover, this data import operation is more complex in the Q-coder: one or several bits may be imported at one time, whereas in the ELS-coder, the quantity of data imported is always exactly one byte.

- (2) The approximation is most accurate for small values of the probability, and is acceptable only for values less than one-half. Thus the Q-coder does not work with symbol values **0** and **1** but with the so-called “MPS” (more probable symbol) and “LPS” (less probable symbol). In encoding or decoding a symbol, the Q-coder must determine which of **0** or **1** is the MPS. Such an extra level of indirection might be useful in an implementation of the ELS-coder, but can be left to the discretion of the implementer.

Unlike the ELS-coder, the Q-coder has no data leaks. On the other hand, under the best of conditions, the approximation used in the Q-coder is less accurate than that used in the ELS-coder even with data leaks taken into consideration. However, in the light of the practical difficulties of accurate probability estimation discussed in Section 4, and the effect of such on coding efficiency, not too much should be made of this distinction.

The Q-coder uses a probability ladder as does the ELS-coder, but derived on different principles. This ladder relies strictly on inalcritous probability updating; an alacritous version would certainly be possible but does not seem to have been the subject of much study.

The question of which of the Q-coder and ELS-coder codes more efficiently and which operates more rapidly depends on many issues of architecture and nature of the data.

The Q-coder performs best in the realm of profound compression: that is, when the probability of the LPS is very close to 0 and each symbol can be coded very compactly. The approximation used in the Q-coder provides its best coding efficiency in such a case, and a high compression ratio entails less frequent importing of data, so that operation is faster. Indeed the highest compression level attainable by the Q-coder ($2^{15} : 1$) exceeds that of the ELS-coder (754 : 8 if **F** takes the value 754). In relative terms the difference is large (a ratio of 43.5) but in absolute terms it is small (a difference of 0.00130 bits per symbol). Our design philosophy is that the absolute measure of coding efficiency is more important.

In the realm of mild compression, the superior accuracy and less frequent data import of the ELS-coder probably give it an edge. It is worth noting that data compression applications can often be designed so that profound compression is unnecessary; such a design is likely to improve both speed and compression efficiency. For example, in a simple binary model, if **0** is overwhelmingly more probable than **1**, one is likely to find long runs of symbols with value **0** in the data. Run-length-encoding symbols with value **0** would reduce the number of calls to the entropy coder while reducing the importance of profound compression for overall compression efficiency.

Another example of such a design decision is given in the sample programs COMPRESS.C and EXPAND.C with respect to the coding of the end-of-file marker. The bit

indicating end-of-file could have been inserted at the base of the state tree. However, this would require every byte of the compressed file to carry an extra **0** bit indicating that it was not end-of-file; profound compression would be needed to compress these bits collectively to a tiny portion of the file. By attaching this bit rather to a node of the tree, we insured that such bits would usually appear much less often; thus these bits usually represent a small portion of the compressed file even without profound compression.

We have compared the QM-coder and the ELS-coder empirically in the context of an image-compression application. True to expectations, the alacritous version of the ELS-coder performed better than the QM-coder on not-easily-compressible images, while the QM-coder performed better on easily-compressible images. The non-alacritous version of the ELS-coder entailed an overall loss in compression efficiency of about 15% (though the relative performance could be manipulated by choosing the images to be compressed); however the the non-alacritous ELS-coder might be the method of choice where speed is an overriding concern.

ACKNOWLEDGMENTS

I would like to thank Stephen Mann for reading an early draft of this article and proposing many improvements in both form and content.

REFERENCES

1. T. M. Cover and J. A. Thomas, *Elements of Information Theory*, John Wiley and Sons, 1991.
2. M. Nelson, *The Data Compression Book*, M&T Books, Redwood City, CA, 1991.
3. W. B. Pennebaker and J. L. Mitchell, *JPEG: Still Image Data Compression Standard*, Van Nostrand Reinhold, New York, 1993.
4. W. B. Pennebaker, J.L. Mitchell, G.G. Langdon, and R.B. Arps, *An overview of the basic principles of the Q-coder adaptive binary arithmetic coder*, IBM Journal of Research and Development **32** (1988), 771–726.
5. Ian H. Witten, Radford M. Neal, and John H. Cleary, *Arithmetic coding for data compression*, Communications of the Association for Computing Machinery **30**, No. 6 (1987), 520–540.
6. , *Arithmetic coding encoder and decoder system*, United States Patent Number 4,905,297 (1990).
7. , *Probability Adaptation for Arithmetic Coders*, United States Patent Number 4,935,882 (1990).