

Coding with Asymmetric Numeral Systems

Jeremy Gibbons

University of Oxford

Abstract. Asymmetric Numeral Systems (ANS) are an entropy-based encoding method introduced by Jarek Duda, combining the Shannon-optimal compression *effectiveness* of arithmetic coding with the execution *efficiency* of Huffman coding. Existing presentations of the ANS encoding and decoding algorithms are somewhat obscured by the lack of suitable presentation techniques; we present here an equational derivation, calculational where it can be, and highlighting the creative leaps where it cannot.

1 Introduction

Entropy encoding techniques compress symbols according to a model of their expected frequencies, with common symbols being represented by fewer bits than rare ones. The best known entropy encoding technique is *Huffman coding* (HC) [20], taught in every undergraduate course on algorithms and data structures: a classic greedy algorithm uses the symbol frequencies to construct a trie, from which an optimal *prefix-free binary code* can be read off. For example, suppose an alphabet of $n = 3$ symbols $s_0 = \text{'a'}$, $s_1 = \text{'b'}$, $s_2 = \text{'c'}$ with respective expected relative frequencies $c_0 = 2$, $c_1 = 3$, $c_2 = 5$ (that is, 'a' is expected $2/2+3+5 = 20\%$ of the time, and so on); then HC might construct the trie and prefix-free code shown in Figure 1. A text is then encoded as the concatenation of its symbol codes; thus, the text "cbcacbcacb" encodes to 1 01 1 00 1 01 1 00 1 01. This is optimal, in the sense that no prefix-free binary code yields a shorter encoding of any text that matches the expected symbol frequencies.

But HC is only 'optimal' among encodings that use a whole number of bits per symbol; if that constraint is relaxed, more effective encoding becomes possible. Note that the two symbols 'a' and 'b' were given equal-length codes 00 and 01 by HC, despite having unequal frequencies—indeed, any expected frequencies in the same order $c_0 < c_1 < c_2$ will give the same code. More starkly, if the



Fig. 1. A Huffman trie and the corresponding prefix-free code

alphabet has only two symbols, HC can do no better than to code each symbol as a single bit, whatever their expected frequencies; that might be acceptable when the frequencies are similar, but is unacceptable when they are not.

Arithmetic coding (AC) [25, 29] is an entropy encoding technique that allows for a fractional number of bits per symbol. In a nutshell, a text is encoded as a half-open subinterval of the unit interval; the model assigns disjoint subintervals of the unit interval to each symbol, in proportion to their expected frequencies (as illustrated on the left of Figure 2); encoding starts with the unit interval, and narrows this interval by the model subinterval for each symbol in turn (the narrowing operation is illustrated on the right of Figure 2). The encoding is the shortest binary fraction in the final interval, without its final ‘1’. For example, with the model illustrated in Figure 2, the text "abc" gets encoded via the narrowing sequence of intervals

$$[0, 1) \xrightarrow{\text{'a'}} [0, 1/5) \xrightarrow{\text{'b'}} [1/25, 1/10) \xrightarrow{\text{'c'}} [7/100, 1/10)$$

from which we pick the binary fraction $3/32$ (since $7/100 \leq 3/32 < 1/10$) and output the bit sequence 0001. We formalize this sketched algorithm in Section 3.

This doesn’t look like much saving: this particular example is only one bit shorter than with HC; and similarly, the arithmetic coding of the text "cbcacbcacb" is 14 bits, where HC uses 15 bits. But AC can do much better; for example, it encodes the permutation "cabbacbcc" of that text in 7 bits, whereas of course HC uses the same 15 bits as before.

In fact, AC is *Shannon-optimal*: the number of bits used tends asymptotically to the Shannon entropy of the message—the sum $\sum_i -\log_2 p_i$ of the negative logarithms of the symbol probabilities. Moreover, AC can be readily made *adaptive*, whereby the model evolves as the text is read, whereas HC entails separate modelling and encoding phases.

However, AC does have some problems. One problem is a historical accident: specific applications of the technique became mired in software patents in the 1980s, and although those patents have now mostly expired, the consequences are still being felt (for example, Seward’s `bzip` compressor [26] switched in 1996




Fig. 2. A text model in interval form. Narrowing interval $[l, r)$ by interval $[p, q)$ yields the interval marked in bold on the left, which stands in relation to $[l, r)$ as $[p, q)$ does to $[0, 1)$.

from AC to HC because of patents, and has not switched back since). A more fundamental problem is that AC involves a lot of arithmetic, and even after slightly degrading coding effectiveness in order to use only single-word fixed-precision rather than arbitrary-precision arithmetic, state-of-the-art implementations are still complicated and relatively slow.

A recent development that addresses both of these problems has been Jarek Duda’s introduction of *asymmetric numeral systems* (ANS) [11, 12, 14]. This is another entropy encoding technique; in a nutshell, rather than encoding longer and longer messages as narrower and narrower subintervals of the unit interval, they are simply encoded as larger and larger integers. Concretely, with the same frequency counts $c_0 = 2, c_1 = 3, c_2 = 5$ as before, and cumulative totals $t_0 = 0, t_1 = t_0 + c_0 = 2, t_2 = t_1 + c_1 = 5, t = t_2 + c_2 = 10$, encoding starts with an accumulator at 0, and for each symbol s_i (traditionally from right to left in the text) maps the current accumulator x to $(x \text{ div } c_i) \times t + t_i + (x \text{ mod } c_i)$, as illustrated in Figure 3. Thus, the text "abc" gets encoded via the increasing (read from right to left) sequence of integers:

$$70 \xleftarrow{'a'} 14 \xleftarrow{'b'} 5 \xleftarrow{'c'} 0$$

It is evident even from this brief sketch that the encoding process is quite simple, with just a single division and multiplication per symbol; it turns out that decoding is just as simple. The encoding seems quite mysterious, but it is very cleverly constructed, and again achieves Shannon-optimal encoding; ANS combines the *effectiveness* of AC with the *efficiency* of Huffman coding, addressing the more fundamental concern with AC. The purpose of this paper is to motivate and justify the development, using calculational techniques where possible.

As it happens, Duda is also fighting to keep ANS in the public domain, despite corporate opposition [22], thereby addressing the more accidental concern too. These benefits have seen ANS recently adopted by large companies for products such as Facebook Zstd [9], Apple LZFS [10], Google Draco [7], and




Fig. 3. The start of the coding table for alphabet 'a', 'b', 'c' with counts 2, 3, 5. The indices 0.. are distributed across the alphabet, in proportion to the counts: two for 'a', three for 'b', and so on. Encoding symbol s with current accumulator x yields the index of the x th blob in row s as the new accumulator. For example, with $x = 4$ and next symbol $s = 'b' = s_i$ with $i = 1$, we have $c_i = 3, t_i = 2, t = 10$ so $x' = (x \text{ div } c_i) \times t + t_i + (x \text{ mod } c_i) = 13$, and indeed the 4th blob in row 'b' (counting from zero) is in column 13.

Dropbox DivANS [24], and ANS is expected [1] to be featured in the forthcoming JPEG XL standard [21].

One disadvantage of ANS is that, whereas AC acts in a first-in first-out manner, ANS acts last-in first-out, in the sense that the decoded text comes out in the reverse order to which it went in. Our development will make clear where this happens. This reversal makes ANS unsuitable for encoding a communications channel, and also makes it difficult to employ adaptive text models. (DivANS [24] processes the input forwards for statistical modelling, and then uses this information backwards to encode the text; one could alternatively batch process the text in fixed-size blocks. In some contexts, such as encoding the video stream of a movie for distribution to set-top boxes, it is worth expending more effort in offline encoding in order to benefit online decoding.)

The remainder of this paper is structured as follows. Section 2 recaps various well-known properties of folds and unfolds on lists. Section 3 presents the relevant basics of AC, and Section 4 a proof of correctness of this basic algorithm. Section 5 presents the key step from AC to ANS, namely the switch from accumulating fractions to accumulating integers. Section 6 shows how to modify this naive ANS algorithm to work in bounded precision, and Section 7 shows how to make the resulting program ‘stream’ (to start generating output before consuming all the input). Section 8 discusses related work and concludes.

We use Haskell [23] as an algorithmic notation. Note that function application binds tightest of all binary operators, so that for example $f\ x + y$ means $(f\ x) + y$; apart from that, we trust that the notation is self-explanatory. We give definitions of functions from the Haskell standard library as we encounter them. The code from the paper is available online [16], as is a longer version [17] of the paper including proofs and other supporting material.

2 Origami programming

In this section, we recap some well-studied laws of folds

$$\begin{aligned} foldr &:: (a \rightarrow b \rightarrow b) \rightarrow b \rightarrow [a] \rightarrow b \\ foldr\ f\ e\ [] &= e \\ foldr\ f\ e\ (x : xs) &= f\ x\ (foldr\ f\ e\ xs) \\ foldl &:: (b \rightarrow a \rightarrow b) \rightarrow b \rightarrow [a] \rightarrow b \\ foldl\ f\ e\ [] &= e \\ foldl\ f\ e\ (x : xs) &= foldl\ f\ (f\ e\ x)\ xs \end{aligned}$$

and unfolds

$$\begin{aligned} unfoldr &:: (b \rightarrow Maybe\ (a, b)) \rightarrow b \rightarrow [a] \\ unfoldr\ f\ y = \mathbf{case}\ f\ y\ \mathbf{of}\ \mathit{Nothing} &\rightarrow [] \\ &\quad \mathit{Just}\ (x, y') \rightarrow x : unfoldr\ f\ y' \end{aligned}$$

on lists.

Folds. The *First Duality Theorem* of *foldl* and *foldr* [5, §3.5.1] states that

$$\text{foldr } f \ e = \text{foldl } f \ e$$

when f and e form a monoid. The *Third Duality Theorem*, from the same source, says:

$$\text{foldr } f \ e \cdot \text{reverse} = \text{foldl } (\text{flip } f) \ e$$

where $\text{flip } f \ a \ b = f \ b \ a$ swaps the arguments of a binary function. (The published version [5, §3.5.1] has the *reverse* on the other side, and holds only for finite lists.)

We will also use the *Fusion Law* for *foldr* [4, §4.6.2]:

$$h \cdot \text{foldr } f \ e = \text{foldr } f' \ e' \iff h \ e = e' \wedge h \ (f \ x \ y) = f' \ x \ (h \ y)$$

(at least, on finite lists), and its corollaries the *Map Fusion* laws:

$$\begin{aligned} \text{foldr } f \ e \cdot \text{map } g &= \text{foldr } (\lambda x \ y \rightarrow f \ (g \ x) \ y) \ e \\ \text{foldl } f \ e \cdot \text{map } g &= \text{foldl } (\lambda x \ y \rightarrow f \ x \ (g \ y)) \ e \end{aligned}$$

Unfolds. The sole law of *unfoldr* stated in the Haskell 98 standard [23, §17.4] gives conditions under which it inverts a *foldr*: if

$$\begin{aligned} g \ (f \ x \ z) &= \text{Just } (x, z) \\ g \ e &= \text{Nothing} \end{aligned}$$

for all x and z , then

$$\text{unfoldr } g \ (\text{foldr } f \ e \ xs) = xs$$

for all finite lists xs . We call this the *Unfoldr-Foldr Theorem*. (The proof is a straightforward induction on xs .)

We make two generalisations to this theorem. The first, the *Unfoldr-Foldr Theorem with Junk*, allows the unfold to continue after reconstructing the original list: if only

$$g \ (f \ x \ z) = \text{Just } (x, z)$$

holds, for all x and z , then

$$\exists ys \ . \ \text{unfoldr } g \ (\text{foldr } f \ e \ xs) = xs \ ++ \ ys$$

for all finite lists xs —that is, the *unfoldr* inverts the *foldr* except for appending some (possibly infinite) junk ys to the output. This too can be proved by induction on xs .

The second generalisation is the *Unfoldr-Foldr Theorem with Invariant*. We say that predicate p is an invariant of *foldr* $f \ e$ and *unfoldr* g if

$$\begin{aligned} p \ (f \ x \ z) &\iff p \ z \\ p \ z' &\iff p \ z \wedge g \ z = \text{Just } (x, z') \end{aligned}$$

for all x, z, z' . The theorem states that if p is such an invariant, and the conditions

$$\begin{aligned} g (f x z) = Just (x, z) &\Leftarrow p z \\ g e = Nothing &\Leftarrow p e \end{aligned}$$

hold for all x and z , then

$$unfoldr g (foldr f e xs) = xs \Leftarrow p e$$

for all finite lists xs . Again, there is a straightforward proof by induction. And of course, there is an *Unfoldr–Foldr Theorem with Junk and Invariant*, incorporating both generalisations; this is the version we will actually use.

3 Arithmetic coding

We start from a simplified version of arithmetic coding: we use a fixed rather than adaptive model, and rather than picking the shortest binary fraction within the final interval, we simply pick the lower bound of the interval.

Intervals and symbols. We represent intervals as pairs of rationals,

$$\text{type Interval} = (\text{Rational}, \text{Rational})$$

so the unit interval is $unit = (0, 1)$ and the lower bound is obtained by fst . We suppose a symbol table

$$counts :: [(Symbol, Integer)]$$

that records a positive count for every symbol in the alphabet; in the interests of brevity, we omit the straightforward definitions of functions

$$\begin{aligned} encodeSym &:: Symbol \rightarrow Interval \\ decodeSym &:: Rational \rightarrow Symbol \end{aligned}$$

that work on this fixed global model, satisfying the central property: for $x \in unit$,

$$decodeSym x = s \Leftrightarrow x \in encodeSym s$$

For example, with the same alphabet of three symbols 'a', 'b', 'c' and counts 2, 3, and 5 as before, we have $encodeSym 'b' = (1/5, 1/2)$ and $decodeSym (1/3) = 'b'$.

We have operations on intervals:

$$\begin{aligned} weight, scale &:: Interval \rightarrow Rational \rightarrow Rational \\ weight (l, r) x &= l + (r - l) \times x \\ scale (l, r) y &= (y - l) / (r - l) \\ narrow &:: Interval \rightarrow Interval \rightarrow Interval \\ narrow i (p, q) &= (weight i p, weight i q) \end{aligned}$$

that satisfy

$$\begin{aligned} weight i x \in i &\Leftrightarrow x \in unit \\ weight i x = y &\Leftrightarrow scale i y = x \end{aligned}$$

Informally, *weight* $(l, r) x$ is ‘fraction x of the way between l and r ’, and conversely *scale* $(l, r) y$ is ‘the fraction of the way y is between l and r ’; and *narrow* is illustrated in Figure 2.

Encoding and decoding. Now we can specify arithmetic encoding and decoding by:

$$\begin{aligned} \text{encode}_1 &:: [\text{Symbol}] \rightarrow \text{Rational} \\ \text{encode}_1 &= \text{fst} \cdot \text{foldl} \text{ estep}_1 \text{ unit} \\ \text{estep}_1 &:: \text{Interval} \rightarrow \text{Symbol} \rightarrow \text{Interval} \\ \text{estep}_1 \ i \ s &= \text{narrow} \ i \ (\text{encodeSym} \ s) \\ \text{decode}_1 &:: \text{Rational} \rightarrow [\text{Symbol}] \\ \text{decode}_1 &= \text{unfoldr} \ \text{dstep}_1 \\ \text{dstep}_1 &:: \text{Rational} \rightarrow \text{Maybe} \ (\text{Symbol}, \text{Rational}) \\ \text{dstep}_1 \ x &= \text{let} \ s = \text{decodeSym} \ x \ \text{in} \ \text{Just} \ (s, \text{scale} \ (\text{encodeSym} \ s) \ x) \end{aligned}$$

For example, with the same alphabet and counts, the input text "abc" gets encoded symbol by symbol, from left to right (because of the *foldl*), starting with the unit interval $(0, 1)$, via the narrowing sequence of intervals

$$\begin{aligned} \text{estep}_1 \ (0, 1) \ 'a' &= (0, 1/5) \\ \text{estep}_1 \ (0, 1/5) \ 'b' &= (1/25, 1/10) \\ \text{estep}_1 \ (1/25, 1/10) \ 'c' &= (7/100, 1/10) \end{aligned}$$

from which we select the lower bound $7/100$ of the final interval. Conversely, decoding starts with $7/100$, and proceeds as follows:

$$\begin{aligned} \text{dstep}_1 \ (7/100) &= \text{Just} \ ('a', 7/20) \\ \text{dstep}_1 \ (7/20) &= \text{Just} \ ('b', 1/2) \\ \text{dstep}_1 \ (1/2) &= \text{Just} \ ('c', 0) \\ \text{dstep}_1 \ 0 &= \text{Just} \ ('a', 0) \\ &\dots \end{aligned}$$

Note that decoding runs forever; but the finite encoded text is a prefix of the decoded output—for any input text xs , there is an infinite sequence of junk ys such that

$$\text{decode}_1 \ (\text{encode}_1 \ xs) = xs \ ++ \ ys$$

(indeed, $ys = \text{repeat} \ 'a'$ when we pick the *fst* of an interval).

4 Correctness of arithmetic coding

Using the laws of folds, we can ‘fission’ the symbol encoding out of encode_1 , turn the *foldl* into a *foldr* (because *narrow* and *unit* form a monoid), fuse the *fst* with the *foldr*, and then re-fuse the symbol encoding with the fold:

$$\begin{aligned} &\text{encode}_1 \\ &= \{ \text{definition} \} \end{aligned}$$

$$\begin{aligned}
&fst \cdot foldl \ estep_1 \ unit \\
= &\{ \text{map fusion for } foldl, \text{ backwards} \} \\
&fst \cdot foldl \ narrow \ unit \cdot map \ encodeSym \\
= &\{ \text{duality: } narrow \ \text{and} \ unit \ \text{form a monoid} \} \\
&fst \cdot foldr \ narrow \ unit \cdot map \ encodeSym \\
= &\{ \text{fusion for } foldr \text{ (see below)} \} \\
&foldr \ weight \ 0 \cdot map \ encodeSym \\
= &\{ \text{map fusion; let } estep_2 \ s \ x = weight \ (encodeSym \ s) \ x \} \\
&foldr \ estep_2 \ 0
\end{aligned}$$

For the fusion step, we have

$$\begin{aligned}
fst \ unit &= 0 \\
fst \ (narrow \ i \ (p, q)) &= weight \ i \ (fst \ (p, q))
\end{aligned}$$

as required. So we have calculated $encode_1 = encode_2$, where

$$\begin{aligned}
encode_2 &:: [Symbol] \rightarrow Rational \\
encode_2 &= foldr \ estep_2 \ 0 \\
estep_2 &:: Symbol \rightarrow Rational \rightarrow Rational \\
estep_2 \ s \ x &= weight \ (encodeSym \ s) \ x
\end{aligned}$$

Now encoding is a simple *foldr*, which means that it is easier to manipulate.

Inverting encoding. Let us turn now to decoding, and specifically the question of whether it faithfully decodes the encoded text. We use the Unfoldr–Foldr Theorem. Of course, we have to accept junk, because our decoder runs indefinitely. We check the inductive condition:

$$\begin{aligned}
&dstep_1 \ (estep_2 \ s \ x) \\
= &\{ \ estep_2; \ \text{let } x' = weight \ (encodeSym \ s) \ x \} \\
&dstep_1 \ x' \\
= &\{ \ dstep_1; \ \text{let } s' = decodeSym \ x' \} \\
&Just \ (s', \ scale \ (encodeSym \ s') \ x')
\end{aligned}$$

Now, we hope to recover the first symbol; that is, we require $s' = s$:

$$\begin{aligned}
&s' = s \\
\Leftrightarrow &\{ \ s' = decodeSym \ x'; \ \text{central property} \} \\
&x' \in encodeSym \ s \\
\Leftrightarrow &\{ \ \text{definition of } x' \} \\
&weight \ (encodeSym \ s) \ x \in encodeSym \ s \\
\Leftrightarrow &\{ \ \text{property of } weight \} \\
&x \in unit
\end{aligned}$$

Fortunately, it is an invariant of the computation that the state x is in the unit interval, as is easy to check; so indeed $s' = s$. Continuing:

$$\begin{aligned}
&dstep_1 \ (estep_2 \ s \ x) \\
= &\{ \ \text{above} \}
\end{aligned}$$

$$\begin{aligned}
& \text{Just } (s, \text{scale } (\text{encodeSym } s) (\text{weight } (\text{encodeSym } s) x)) \\
= & \{ \text{scale } i \cdot \text{weight } i = \text{id} \} \\
& \text{Just } (s, x)
\end{aligned}$$

as required. Therefore, by the Unfoldr–Foldr Theorem with Junk and Invariant, decoding inverts encoding, up to junk: for all finite xs ,

$$\exists ys . \text{decode}_1 (\text{encode}_2 xs) = xs \# ys$$

But we can discard the junk, by pruning to the desired length:

$$\text{take } (\text{length } xs) (\text{decode}_1 (\text{encode}_2 xs)) = xs$$

for all finite xs . Alternatively, we can use an ‘end of text’ marker ω that is distinct from all proper symbols:

$$\text{takeWhile } (\neq \omega) (\text{decode}_1 (\text{encode}_2 (xs \# [\omega]))) = xs \iff \omega \notin xs$$

for all finite xs . Either way, arithmetic decoding does indeed faithfully invert arithmetic coding.

5 From fractions to integers

We now make the key step from AC to ANS. Whereas AC encodes longer and longer messages as more and more precise fractions, ANS encodes them as larger and larger integers. Given the symbol table *counts* as before, we can easily derive definitions of the following functions (for example, by tabulating *cumul* using a scan)—we again omit the definitions for brevity:

$$\begin{aligned}
\text{count} & :: \text{Symbol} \rightarrow \text{Integer} \\
\text{cumul} & :: \text{Symbol} \rightarrow \text{Integer} \\
\text{find} & :: \text{Integer} \rightarrow \text{Symbol}
\end{aligned}$$

such that *count* s gives the count of symbol s , *cumul* s gives the cumulative counts of all symbols preceding s in the symbol table, and *find* x looks up an integer $0 \leq x < t$:

$$\text{find } x = s \iff \text{cumul } s \leq x < \text{cumul } s + \text{count } s$$

(where $t = \text{sum } (\text{map } \text{snd } \text{counts})$ is the total count).

The integer encoding step. We encode a text as an integer x , containing $\log_2 x$ bits of information. The next symbol s to encode has probability $p = \text{count } s / t$, and so requires an additional $\log_2 (1/p)$ bits; in total, that makes $\log_2 x + \log_2 (1/p) = \log_2 (x/p) = \log_2 (x \times t / \text{count } s)$ bits. So entropy considerations tell us that, roughly speaking, to incorporate symbol s into state x we want to map x to $x' \simeq x \times t / \text{count } s$. Of course, in order to decode, we need to be able to invert this transformation, to extract s and x from x' ; this suggests that we should do the division by *count* s first:

$$x' = (x \text{ div } \text{count } s) \times t \quad \text{-- not final}$$

so that the multiplication by the known value t can be undone first:

$$x \text{ div } \textit{count } s = x' \text{ div } t$$

(we will refine this definition shortly). How do we reconstruct s ? Well, there is enough headroom in x' to add any value u with $0 \leq u < t$ without affecting the division; in particular, we can add $\textit{cumul } s$ to x' , and then we can use \textit{find} on the remainder:

$$x' = (x \text{ div } \textit{count } s) \times t + \textit{cumul } s \quad \text{-- still not final}$$

so that

$$\begin{aligned} x \text{ div } \textit{count } s &= x' \text{ div } t \\ \textit{cumul } s &= x' \text{ mod } t \\ s &= \textit{find } (\textit{cumul } s) = \textit{find } (x' \text{ mod } t) \end{aligned}$$

(this version still needs to be refined further). We are still missing some information from the lower end of x through the division, namely $x \text{ mod } \textit{count } s$; so we can't yet reconstruct x . Happily,

$$\textit{find } (\textit{cumul } s) = \textit{find } (\textit{cumul } s + r)$$

for any r with $0 \leq r < \textit{count } s$; of course, $x \text{ mod } \textit{count } s$ is in this range, so there is still precisely enough headroom in x' to add this lost information too, without affecting the \textit{find} , allowing us also to reconstruct x :

$$x' = (x \text{ div } \textit{count } s) \times t + \textit{cumul } s + x \text{ mod } \textit{count } s \quad \text{-- final}$$

so that

$$\begin{aligned} x \text{ div } \textit{count } s &= x' \text{ div } t \\ s &= \textit{find } (\textit{cumul } s + x \text{ mod } \textit{count } s) \\ &= \textit{find } (x' \text{ mod } t) \\ x &= \textit{count } s \times (x \text{ div } \textit{count } s) + x \text{ mod } \textit{count } s \\ &= \textit{count } s \times (x' \text{ div } t) + x' \text{ mod } t - \textit{cumul } s \end{aligned}$$

This is finally the transformation we will use for encoding one more symbol.

Integer ANS. We define

$$\begin{aligned} \textit{encode}_3 &:: [\textit{Symbol}] \rightarrow \textit{Integer} \\ \textit{encode}_3 &= \textit{foldr } \textit{estep}_3 \textit{ } 0 \\ \textit{estep}_3 &:: \textit{Symbol} \rightarrow \textit{Integer} \rightarrow \textit{Integer} \\ \textit{estep}_3 \textit{ } s \textit{ } x &= \textit{let } (q, r) = x \text{ divMod } \textit{count } s \textit{ in } q \times t + \textit{cumul } s + r \\ \textit{decode}_3 &:: \textit{Integer} \rightarrow [\textit{Symbol}] \\ \textit{decode}_3 &= \textit{unfoldr } \textit{dstep}_3 \\ \textit{dstep}_3 &:: \textit{Integer} \rightarrow \textit{Maybe } (\textit{Symbol}, \textit{Integer}) \\ \textit{dstep}_3 \textit{ } x &= \textit{let } (q, r) = x \text{ divMod } t \\ &\quad s = \textit{find } r \\ &\textit{in } \textit{Just } (s, \textit{count } s \times q + r - \textit{cumul } s) \end{aligned}$$

Correctness of integer ANS. Using similar reasoning as for AC, we can show that a decoding step inverts an encoding step:

$$\begin{aligned}
& dstep_3 (estep_3 s x) \\
= & \{ estep_3; \text{let } (q, r) = x \text{ divMod } count \ s, x' = q \times t + cumul \ s + r \} \\
& dstep_3 x' \\
= & \{ dstep_3; \text{let } (q', r') = x' \text{ divMod } t, s' = find \ r' \} \\
& Just (s', count \ s' \times q' + r' - cumul \ s') \\
= & \{ r' = cumul \ s + r, 0 \leq r < count \ s, \text{ so } s' = find \ r' = s \} \\
& Just (s, count \ s \times q' + r' - cumul \ s) \\
= & \{ r' - cumul \ s = r, q' = x' \text{ div } t = q \} \\
& Just (s, count \ s \times q + r) \\
= & \{ (q, r) = x \text{ divMod } count \ s \} \\
& Just (s, x)
\end{aligned}$$

Therefore decoding inverts encoding, modulo pruning, by the Unfoldr–Foldr Theorem with Junk:

$$take (length \ xs) (decode_3 (encode_3 \ xs)) = xs$$

for all finite xs . For example, with the same alphabet and symbol counts as before, encoding the text "abc" proceeds (now from right to left, because of the *foldr* in $encode_3$) as follows:

$$\begin{aligned}
estep_3 \ 'c' \ 0 &= 5 \\
estep_3 \ 'b' \ 5 &= 14 \\
estep_3 \ 'a' \ 14 &= 70
\end{aligned}$$

and the result is 70. Decoding inverts this:

$$\begin{aligned}
dstep_3 \ 70 &= Just \ ('a', 14) \\
dstep_3 \ 14 &= Just \ ('b', 5) \\
dstep_3 \ 5 &= Just \ ('c', 0) \\
dstep_3 \ 0 &= Just \ ('a', 0) \\
&\dots
\end{aligned}$$

Huffman as an instance of ANS. Incidentally, we can see here that ANS is in fact a generalisation of HC. If the symbol counts and their sum are all powers of two, then the arithmetic in $estep_3$ amounts to simple manipulation of bit vectors by shifting and insertion. For example, with an alphabet of four symbols 'a', 'b', 'c', 'd' with counts 4, 2, 1, 1, encoding operates on a state x with binary expansion $\dots x_3 x_2 x_1 x_0$ (written most significant bit first) as follows:

$$\begin{aligned}
estep_3 \ 'a' \ (\dots x_3 x_2 x_1 x_0) &= \dots x_3 x_2 0 x_1 x_0 \\
estep_3 \ 'b' \ (\dots x_3 x_2 x_1 x_0) &= \dots x_3 x_2 x_1 1 0 x_0 \\
estep_3 \ 'c' \ (\dots x_3 x_2 x_1 x_0) &= \dots x_3 x_2 x_1 x_0 1 1 0 \\
estep_3 \ 'd' \ (\dots x_3 x_2 x_1 x_0) &= \dots x_3 x_2 x_1 x_0 1 1 1
\end{aligned}$$

That is, the symbol codes 0, 10, 110, 111 are inserted into rather than appended onto the state so far; the binary expansion of the ANS encoding of a text yields some permutation of the HC encoding of that text.

A different starting point. As it happens, the inversion property of $encode_3$ and $decode_3$ holds, whatever value we use to start encoding with (since this value is not used in the proof); in Section 6, we start encoding with a certain lower bound l rather than 0. Moreover, $estep_3$ is strictly increasing on states strictly greater than zero, and $dstep_3$ strictly decreasing; which means that the decoding process can stop when it returns to the lower bound. That is, if we pick some $l > 0$ and define

$$\begin{aligned} encode_4 &:: [Symbol] \rightarrow Integer \\ encode_4 &= foldr estep_3 l \\ decode_4 &:: Integer \rightarrow [Symbol] \\ decode_4 &= unfoldr dstep_4 \\ dstep_4 &:: Integer \rightarrow Maybe (Symbol, Integer) \\ dstep_4 x &= \mathbf{if} \ x == l \ \mathbf{then} \ \mathit{Nothing} \ \mathbf{else} \ dstep_3 \ x \end{aligned}$$

then the stronger version of the Unfoldr–Foldr Theorem (without junk) holds, and we have

$$decode_4 (encode_4 \ xs) = xs$$

for all finite xs .

6 Bounded precision

The previous versions all used arbitrary-precision arithmetic, which is expensive. We now change the approach slightly to use only bounded-precision arithmetic. As usual, there is a trade-off between effectiveness (a bigger bound on the numbers involved means more accurate approximations to ideal entropies) and efficiency (a smaller bound generally means faster operations). Fortunately, the reasoning does not depend much on the actual bounds. We will pick a base b and a lower bound l , and represent the integer accumulator x as a pair (w, ys) which we call a *Number*:

type $Number = (Int, [Int])$

such that ys is a list of digits in base b , and w is an integer in the range $l \leq w < u$ (where we define $u = l \times b$ for the upper bound), under the abstraction relation $x = abstract \ (w, ys)$ induced by

$$abstract \ (w, ys) = foldl \ inject \ w \ ys$$

where

$$inject \ w \ y = w \times b + y$$

We call w the ‘window’ and ys the ‘remainder’. For example, with $b = 10$ and $l = 100$, the pair $(123, [4, 5, 6])$ represents the value 123456.

Properties of the window. Specifying a range of the form $l \leq w < l \times b$ induces nice properties. If we introduce an operation

$$\text{extract } w = w \text{ divMod } b$$

as an inverse to *inject*, then we have

$$\begin{aligned} \text{inject } w < u &\Leftrightarrow w < l \\ l \leq \text{fst}(\text{extract } w) &\Leftrightarrow u \leq w \end{aligned}$$

(we omit the straightforward proofs, using the universal property

$$u < v \times w \Leftrightarrow u \text{ div } w < v$$

of integer division). That is, given an in-range window value w , injecting another digit will take it outside (above) the range; but if w is initially below the range, injecting another digit will keep it below the upper bound. So starting below the range and repeatedly injecting digits will eventually land within the range (it cannot hop right over), and injecting another digit would take it outside the range again. Conversely, given an in-range window value w , extracting a digit will take it outside (below) the range; but if w is initially above the range, extracting a digit will keep it at least the lower bound. So starting above the range and repeatedly extracting digits will also eventually land within the range (it cannot hop right over), and extracting another digit would take it outside the range again. This is illustrated in Figure 4. In particular, for any $x \geq l$ there is a unique representation of x under *abstract* that has an in-range window.

For fast execution, b should be a power of two, so that multiplication and division by b can be performed by bit shifts; and arithmetic on values up to u should fit within a single machine word. It is also beneficial for t to divide evenly into l , as we shall see shortly.

Encoding with bounded arithmetic. The encoding step acts on the window in the accumulator using *estep*₃, which risks making it overflow the range; we




Fig. 4. ‘Can’t miss it’ properties of the range: injecting an extra digit can only land within the range $[l, u)$ when starting below it, and will land above the range when starting within it; and conversely, extracting a digit can only land within the range when starting above it, and will land below the range when starting within it.

therefore renormalize with $enorm_5$ by shifting digits from the window to the remainder until this overflow would no longer happen, before consuming the symbol.

$$\begin{aligned}
econsume_5 &:: [Symbol] \rightarrow Number \\
econsume_5 &= foldr estep_5 (l, []) \\
estep_5 &:: Symbol \rightarrow Number \rightarrow Number \\
estep_5 s (w, ys) &= \mathbf{let} (w', ys') = enorm_5 s (w, ys) \mathbf{in} (estep_3 s w', ys') \\
enorm_5 &:: Symbol \rightarrow Number \rightarrow Number \\
enorm_5 s (w, ys) &= \mathbf{if} \quad estep_3 s w < u \\
&\quad \mathbf{then} (w, ys) \\
&\quad \mathbf{else} \mathbf{let} (q, r) = extract w \mathbf{in} enorm_5 s (q, r : ys)
\end{aligned}$$

That is, $enorm_5$ preserves the abstract value of a $Number$:

$$abstract \cdot enorm_5 = abstract$$

and leaves the window safe for $estep_3$ to incorporate the next symbol.

Note that if t divides l , then we can rearrange the guard in $enorm_5$:

$$\begin{aligned}
&estep_3 s w < u \\
\Leftrightarrow &\{ estep_3; \mathbf{let} (q, r) = w \mathit{divMod} count s \} \\
&q \times t + cumul s + r < u \\
\Leftrightarrow &\{ t \text{ divides } l, \text{ so } u = (u \mathit{div} t) \times t \} \\
&q \times t + cumul s + r < (u \mathit{div} t) \times t \\
\Leftrightarrow &\{ \text{universal property of division: } u < v \times w \Leftrightarrow u \mathit{div} w < v \} \\
&(q \times t + cumul s + r) \mathit{div} t < u \mathit{div} t \\
\Leftrightarrow &\{ 0 \leq r < count s, \text{ so } 0 \leq cumul s + r < t \} \\
&q < u \mathit{div} t \\
\Leftrightarrow &\{ q = w \mathit{div} count s \} \\
&w \mathit{div} count s < u \mathit{div} t \\
\Leftrightarrow &\{ \text{universal property of } \mathit{div} \text{ again} \} \\
&w < (u \mathit{div} t) \times count s \\
\Leftrightarrow &\{ u = l \times b, t \text{ divides } l \} \\
&w < b \times (l \mathit{div} t) \times count s
\end{aligned}$$

This is worthwhile because $b \times (l \mathit{div} t)$ is a constant, independent of s , so the comparison can be done with a single multiplication, whereas the definition of $estep_3$ involves a division by $count s$.

For example, consider again encoding the text "abc", with $b = 10$ and $l = 100$. The process is again from right to left, with the accumulator starting at $(100, [])$. Consuming the 'c' then the 'b' proceeds as before, because the window does not overflow u :

$$\begin{aligned}
estep_5 'c' (100, []) &= (205, []) \\
estep_5 'b' (205, []) &= (683, [])
\end{aligned}$$

Now directly consuming the 'a' would make the window overflow, because $estep_3 \text{'a'} 683 = 3411 \geq u$; so we must renormalize to $(68, [3])$ before consuming the 'a', leading to the final state $(340, [3])$:

$$\begin{aligned} enorm_5 \text{'a'} (683, []) &= (68, [3]) \\ estep_5 \text{'a'} (683, []) &= (estep_3 \text{'a'} 68, [3]) = (340, [3]) \end{aligned}$$

Note that the move from arbitrary to fixed precision is not just a data refinement—it is not the case that $econsume_5 xs$ computes some representation of $encode_4 xs$. For example, $encode_4 \text{"abc"} = 3411$, whereas $econsume_5 \text{"abc"} = (340, [3])$, which is not a representation of 3411. We have really sacrificed some effectiveness in encoding in return for the increased efficiency of fixed precision arithmetic.

Decoding with bounded arithmetic. Decoding is an unfold using the accumulator as state. We repeatedly output a symbol from the window; this may make the window underflow the range, in which case we renormalize if possible by injecting digits from the remainder (and if this is not possible, because there are no more digits to inject, it means that we have decoded the entire text).

$$\begin{aligned} dproduce_5 &:: Number \rightarrow [Symbol] \\ dproduce_5 &= unfoldr dstep_5 \\ dstep_5 &:: Number \rightarrow Maybe (Symbol, Number) \\ dstep_5 (w, ys) &= \text{let } Just (s, w') = dstep_3 w \\ &\quad (w'', ys'') = dnorm_5 (w', ys) \\ &\quad \text{in if } w'' \geq l \text{ then } Just (s, (w'', ys'')) \text{ else Nothing} \\ dnorm_5 &:: Number \rightarrow Number \\ dnorm_5 (w, y : ys) &= \text{if } w < l \text{ then } dnorm_5 (inject w y, ys) \text{ else } (w, y : ys) \\ dnorm_5 (w, []) &= (w, []) \end{aligned}$$

Note that decoding is of course symmetric to encoding; in particular, when encoding we renormalize before consuming a symbol; therefore when decoding we renormalize *after* emitting a symbol. For example, decoding the final encoding $(340, [3])$ starts by computing $dstep_3 340 = Just \text{'a'}, 68$; the window value 68 has underflowed, so renormalization consumes the remaining digit 3, restoring the accumulator to $(683, [])$; then decoding proceeds to extract the 'b' and 'c' in turn, returning the accumulator to $(100, [])$ via precisely the same states as for encoding, only in reverse order.

$$\begin{aligned} dstep_5 (340, [3]) &= Just \text{'a'}, (683, []) \\ dstep_5 (683, []) &= Just \text{'b'}, (205, []) \\ dstep_5 (205, []) &= Just \text{'c'}, (100, []) \\ dstep_5 (100, []) &= Nothing \end{aligned}$$

Correctness of decoding. We can prove that decoding inverts encoding, again using the Unfoldr–Foldr Theorem with Invariant. Here, the invariant p is that the window w is in range ($l \leq w < u$), which is indeed maintained by the consumer

$estep_5$ and producer $dstep_5$. As for the conditions of the theorem: in the base case, $dstep_3 l = Just (s, w')$ with $w' < l$, and $dnorm_5 (w', []) = (w', [])$, so indeed

$$dstep_5 (l, []) = Nothing$$

For the inductive step, suppose that $l \leq w < u$; then we have:

$$\begin{aligned} & dstep_5 (estep_5 s (w, ys)) \\ = & \{ estep_5; \text{let } (w', ys') = enorm_5 s (w, ys) \} \\ & dstep_5 (estep_3 s w', ys') \\ = & \{ dstep_5, dstep_3; \text{let } (w'', ys'') = dnorm_5 (w', ys') \} \\ & \text{if } w'' \geq l \text{ then } Just (s, (w'', ys'')) \text{ else } Nothing \\ = & \{ \text{see below: } dnorm_5 \text{ inverts } enorm_5 s, \text{ so } (w'', ys'') = (w, ys) \} \\ & \text{if } w \geq l \text{ then } Just (s, (w, ys)) \text{ else } Nothing \\ = & \{ \text{invariant holds, so in particular } w \geq l \} \\ & Just (s, (w, ys)) \end{aligned}$$

The remaining proof obligation is to show that

$$dnorm_5 (enorm_5 s (w, ys)) = (w, ys)$$

when $l \leq w < u$. We prove this in several steps. First, note that $dnorm_5$ is idempotent:

$$dnorm_5 \cdot dnorm_5 = dnorm_5$$

Second, when $l \leq w$ holds,

$$dnorm_5 (w, ys) = (w, ys)$$

Finally, the key lemma is that, for $w < u$ (but not necessarily $w \geq l$), $dnorm_5$ is invariant under $enorm_5$:

$$dnorm_5 (enorm_5 s (w, ys)) = dnorm_5 (w, ys)$$

When additionally $w \geq l$, the second property allows us to conclude that $dnorm_5$ inverts $enorm_5$:

$$dnorm_5 (enorm_5 s (w, ys)) = (w, ys)$$

The ‘key lemma’ is proved by induction on w . For $w = 0$, we clearly have

$$\begin{aligned} & dnorm_5 (enorm_5 s (w, ys)) \\ = & \{ estep_3 s 0 = cumul s \leq t \leq l, \text{ so } enorm_5 s (w, ys) = (w, ys) \} \\ & dnorm_5 (w, ys) \end{aligned}$$

For the inductive step, we suppose that the result holds for all $q < w$, and consider two cases for w itself. In case $estep_3 s w < u$, we have:

$$\begin{aligned} & dnorm_5 (enorm_5 s (w, ys)) \\ = & \{ \text{assumption; } enorm_5 \} \\ & dnorm_5 (w, ys) \end{aligned}$$

as required. And in case $estep_3 s w \geq u$, we have:

$$\begin{aligned}
& dnorm_5 (enorm_5 s (w, ys)) \\
= & \{ \text{assumption; } enorm_5; \text{ let } (q, r) = \text{extract } w \} \\
& dnorm_5 (enorm_5 s (q, r : ys)) \\
= & \{ q < w; \text{induction} \} \\
& dnorm_5 (q, r : ys) \\
= & \{ w < u, \text{ so } q = w \text{ div } b < l; dnorm_5 \} \\
& dnorm_5 (\text{inject } q \ r, ys) \\
= & \{ q, r \} \\
& dnorm_5 (w, ys)
\end{aligned}$$

Note that we made essential use of the limits of the range: $w < u \Rightarrow w \text{ div } b < l$. Therefore decoding inverts encoding:

$$dproduce_5 (consume_5 xs) = xs$$

for all finite xs .

7 Streaming

The version of encoding in the previous section yields a *Number*, that is, a pair consisting of an integer window and a digit-sequence remainder. It would be more conventional for encoding to take a sequence of symbols to a sequence of digits alone, and decoding to take the sequence of digits back to a sequence of symbols. For encoding, we have to flush the remaining digits out of the window at the end of the process, reducing the window to zero:

$$\begin{aligned}
& eflush_5 :: \text{Number} \rightarrow [\text{Int}] \\
& eflush_5 (0, ys) = ys \\
& eflush_5 (w, y) = \mathbf{let} (w', y) = \text{extract } w \mathbf{in} eflush_5 (w', y : ys)
\end{aligned}$$

For example, $eflush_5 (340, [3]) = [3, 4, 0, 3]$. Then we can define

$$\begin{aligned}
& encode_5 :: [\text{Symbol}] \rightarrow [\text{Int}] \\
& encode_5 = eflush_5 \cdot consume_5
\end{aligned}$$

Correspondingly, decoding should start by populating an initially-zero window from the sequence of digits:

$$\begin{aligned}
& dstart_5 :: [\text{Int}] \rightarrow \text{Number} \\
& dstart_5 ys = dnorm_5 (0, ys)
\end{aligned}$$

For example, $dstart_5 [3, 4, 0, 3] = (340, [3])$. Then we can define

$$\begin{aligned}
& decode_5 :: [\text{Int}] \rightarrow [\text{Symbol}] \\
& decode_5 = dproduce_5 \cdot dstart_5
\end{aligned}$$

One can show that $dstart_5$ inverts $eflush_5$ on in-range values:

$$dstart_5 (eflush_5 (w, ys)) = (w, ys) \quad \Leftarrow \quad l \leq w < u$$

and therefore again decoding inverts encoding:

$$\begin{aligned}
& \text{decode}_5 (\text{encode}_5 \text{ } xs) \\
= & \{ \text{decode}_5, \text{encode}_5 \} \\
& \text{dproduce}_5 (\text{dstart}_5 (\text{eflush}_5 (\text{econsume}_5 \text{ } xs))) \\
= & \{ \text{econsume}_5 \text{ yields in-range, on which } \text{dstart}_5 \text{ inverts } \text{eflush}_5 \} \\
& \text{dproduce}_5 (\text{econsume}_5 \text{ } xs) \\
= & \{ \text{dproduce}_5 \text{ inverts } \text{econsume}_5 \} \\
& xs
\end{aligned}$$

for all finite xs .

Introducing streaming. We would now like to *stream* the encoding and decoding processes, so that each can start generating output before having consumed all its input. With some effort, it is possible to persuade the definitions of encode_5 and decode_5 into *metamorphism* form [15]; however, that turns out to be rather complicated. Here, we take a more direct route instead.

For encoding, we have

$$\text{encode}_5 = \text{eflush}_5 \cdot \text{foldr } \text{estep}_5 (l, [])$$

A first step for streaming is to make as much of this as possible tail-recursive. The best we can do is to apply the Third Duality Theorem to transform the *foldr* into a *foldl*:

$$\text{encode}_5 = \text{eflush}_5 \cdot \text{foldl } (\text{flip } \text{estep}_5) (l, []) \cdot \text{reverse}$$

Now we note that the remainder component of the *Number* behaves like a queue, in the sense that already-enqueued digits simply pass through without being further examined:

$$\begin{aligned}
\text{eflush}_5 (w, ys \ ++ \ zs) &= \text{eflush}_5 (w, ys) \ ++ \ zs \\
\text{enorm}_5 \ s (w, ys \ ++ \ zs) &= \mathbf{let} (w', ys') = \text{enorm}_5 \ s (w, ys) \ \mathbf{in} (w', ys' \ ++ \ zs) \\
\text{estep}_5 \ s (w, ys \ ++ \ zs) &= \mathbf{let} (w', ys') = \text{estep}_5 \ s (w, ys) \ \mathbf{in} (w', ys' \ ++ \ zs)
\end{aligned}$$

If we then introduce the auxilliary functions e_1, e_2 specified by

$$\begin{aligned}
\text{reverse } (e_1 \ w \ ss) \ ++ \ ys &= \text{eflush}_5 (\text{foldl } (\text{flip } \text{estep}_5) (w, ys) \ ss) \\
\text{reverse } (e_2 \ w) \ ++ \ ys &= \text{eflush}_5 (w, ys)
\end{aligned}$$

and unfold definitions, exploiting the queueing properties, we can synthesize $\text{encode}_5 = \text{encode}_6$, where:

$$\begin{aligned}
\text{encode}_6 &:: [\text{Symbol}] \rightarrow [\text{Int}] \\
\text{encode}_6 &= \text{reverse} \cdot e_1 \ l \cdot \text{reverse} \ \mathbf{where} \\
e_1 \ w \ (s : ss) &= \mathbf{let} (q, r) = w \ \text{divMod} \ \text{count} \ s \\
& \quad w' = q \times t + \text{cumul} \ s + r \ \mathbf{in} \\
& \quad \mathbf{if} \ w' < u \ \mathbf{then} \ e_1 \ w' \ ss \\
& \quad \quad \mathbf{else} \ \mathbf{let} (q', r') = w \ \text{divMod} \ b \ \mathbf{in} \ r' : e_1 \ q' \ (s : ss) \\
e_1 \ w \ [] &= e_2 \ w \\
e_2 \ w &= \mathbf{if} \ w == 0 \ \mathbf{then} \ [] \ \mathbf{else} \ \mathbf{let} (w', y) = w \ \text{divMod} \ b \ \mathbf{in} \ y : e_2 \ w'
\end{aligned}$$

In this version, the accumulator w simply maintains the window, and digits in the remainder are output as soon as they are generated. Note that the two *reverses* mean that encoding effectively reads its input and writes its output from right to left; that seems to be inherent to ANS.

Streaming decoding. Decoding is easier, because $dnorm_5$ is already tail-recursive. Similarly specifying functions d_0, d_1, d_2 by

$$\begin{aligned} d_0 w ys &= dproduce_5 (dnorm_5 (w, ys)) \\ d_1 w ys &= dproduce_5 (w, ys) \\ d_2 s w' ys &= \mathbf{let} (w'', ys'') = dnorm_5 (w', ys) \mathbf{in} \\ &\quad \mathbf{if} w'' \geq l \mathbf{then} s : d_1 w'' ys'' \mathbf{else} [] \end{aligned}$$

and unfolding definitions allows us to synthesize directly that $decode_5 = decode_6$, where:

$$\begin{aligned} decode_6 &:: [Int] \rightarrow [Symbol] \\ decode_6 &= d_0 0 \mathbf{where} \\ d_0 w (y : ys) \mid w < l &= d_0 (w \times b + y) ys \\ d_0 w ys &= d_1 w ys \\ d_1 w ys &= \mathbf{let} (q, r) = w \mathit{divMod} t \\ &\quad s = \mathit{find} r \\ &\quad w' = \mathit{count} s \times q + r - \mathit{cumul} s \\ &\quad \mathbf{in} d_2 s w' ys \\ d_2 s w (y : ys) \mid w < l &= d_2 s (w \times b + y) ys \\ d_2 s w ys \mid w \geq l &= s : d_1 w ys \\ d_2 s w [] &= [] \end{aligned}$$

Ignoring additions and subtractions, encoding involves one division by $\mathit{count} s$ and one multiplication by t for each input symbol s , plus one division by b for each output digit. Conversely, decoding involves one multiplication by b for each input digit, plus one division by t and one multiplication by $\mathit{count} s$ for each output symbol s . Encoding and decoding are both tail-recursive. The arithmetic in base b can be simplified to bit shifts by choosing b to be a power of two. They therefore correspond rather directly to simple imperative implementations [19].

8 Conclusion

We have presented a development using the techniques of constructive functional programming of the encoding and decoding algorithms involved in asymmetric numeral systems, including the step from arbitrary- to fixed-precision arithmetic and then to streaming processes. The calculational techniques depend on the theory of folds and unfolds for lists, especially the duality between *foldr* and *foldl*, fusion, and the Unfoldr–Foldr Theorem. We started out with the hypothesis that the theory of streaming developed by the author together with Richard Bird for arithmetic coding [3, 15] would be a helpful tool; but although it can be applied, it seems here to be more trouble than it is worth.

To be precise, what we have described is the *range* variant (rANS) of ANS. There is also a *tabled* variant (tANS), used by Zstd [9] and LZFS [10], which essentially tabulates the functions $estep_5$ and $dstep_5$; for encoding this involves a table of size $n \times (u - l)$, the product of the alphabet size and the window width, and for decoding two tables of size $u - l$. Tabulation makes even more explicit that HC is a special case of ANS, with the precomputed table corresponding to the Huffman trie. Tabulation also allows more flexibility in the precise allocation of codes, which slightly improves effectiveness [14]. For example, the coding table in Figure 3 corresponds to the particular arrangement "aabbcc" of the three symbols in proportion to their counts, and lends itself to implementation via arithmetic; but any permutation of this arrangement would still work, and a permutation such as "cbcacbcacb" which distributes the symbols more evenly turns out to be slightly more effective and no more difficult to tabulate.

One nice feature of AC is that the switch from arbitrary-precision to fixed-precision arithmetic can be expressed in terms of a carefully chosen adaptive model, which slightly degrades the ideal distribution in order to land on convenient rational endpoints [27]. We do not have that luxury with ANS, because of the awkwardness of incorporating adaptive coding; consequently, it is not clear that there is any simple relationship between the arbitrary-precision and fixed-precision versions. But even with AC, that nice feature only applies to encoding; the approximate arithmetic seems to preclude a correctness argument in terms of the Unfoldr–Foldr Theorem, and therefore a completely different (and more complicated) approach is required for decoding [3].

The ANS algorithms themselves are of course not novel here; they are due to Duda [13, 14]. Our development in Section 5 of the key bit of arithmetic in ANS encoding was informed by a very helpful commentary on Duda’s paper published in a series of twelve blog posts [6] by Charles Bloom. The illustration in Figure 3 derives from Duda, and was also used by Roman Cheplyaka [8] as the basis of a (clear but very inefficient) prototype Haskell implementation.

Acknowledgements: I am grateful to participants of IFIP WG2.1 Meeting 78 in Xitou and IFIP WG2.11 Meeting 19 in Salem, who gave helpful feedback on earlier presentations of this work, and to Jarek Duda and Richard Bird who gave much helpful advice on content and presentation. I thank the Programming Research Laboratory at the National Institute of Informatics in Tokyo, for providing a Visiting Professorship during which some of this research was done; in particular, Zhixuan Yang and Josh Ko of NII explained to me the significance of t dividing evenly into l , as exploited in Section 6.

References

1. Jyrki Alakuijala. Google’s compression projects. <https://encode.ru/threads/3108-Google-s-compression-projects#post60072>, May 2019.
2. Richard Bird. Proof of equivalence between huffman trees and certain bags of fractions. Personal communication, April 2019.

3. Richard Bird and Jeremy Gibbons. Arithmetic coding with folds and unfolds. In Johan Jeuring and Simon Peyton Jones, editors, *Advanced Functional Programming 4*, volume 2638 of *Lecture Notes in Computer Science*, pages 1–26. Springer-Verlag, 2003.
4. Richard S. Bird. *Introduction to Functional Programming Using Haskell*. Prentice-Hall, 1998.
5. Richard S. Bird and Philip L. Wadler. *An Introduction to Functional Programming*. Prentice-Hall, 1988.
6. Charles Bloom. Understanding ANS. <http://cbloomrants.blogspot.com/2014/01/1-30-14-understanding-ans-1.html>, January 2014.
7. Jamieson Brette and Frank Galligan. Introducing Draco: Compression for 3D graphics. <https://opensource.googleblog.com/2017/01/introducing-draco-compression-for-3d.html>, January 2017.
8. Roman Cheplyaka. Understanding asymmetric numeral systems. <https://ro-che.info/articles/2017-08-20-understanding-ans>, August 2017.
9. Yann Collet and Chip Turner. Smaller and faster data compression with Zstandard. <https://code.fb.com/core-data/smaller-and-faster-data-compression-with-zstandard/>, August 2016.
10. Sergio De Simone. Apple open-sources its new compression algorithm LZFBSE. *InfoQ*, July 2016. <https://www.infoq.com/news/2016/07/apple-lzfbse-lossless-opensource>.
11. Jarosław Duda. Kodowanie i generowanie układówstatystycznych za pomocą algorytmów probabilistycznych (Coding and generation of statistical systems using probabilistic algorithms). Master’s thesis, Faculty of Physics, Astronomy and Applied Computer Science, Jagiellonian University, 2006.
12. Jarosław Duda. Optimal encoding on discrete lattice with translational invariant constraints using statistical algorithms. *CoRR*, abs/0710.3861, 2007. English translation of [11].
13. Jarosław Duda. Asymmetric numeral systems. *CoRR*, 0902.0271v5, May 2009.
14. Jarosław Duda. Asymmetric numeral systems: Entropy coding combining speed of Huffman coding with compression rate of arithmetic coding. *CoRR*, 1311.2540v2, January 2014.
15. Jeremy Gibbons. Metamorphisms: Streaming representation-changers. *Science of Computer Programming*, 65(2):108–139, 2007.
16. Jeremy Gibbons. Coding with asymmetric numeral systems (Haskell code). <http://www.cs.ox.ac.uk/jeremy.gibbons/publications/asymm.hs>, July 2019.
17. Jeremy Gibbons. Coding with asymmetric numeral systems (long version). <http://www.cs.ox.ac.uk/jeremy.gibbons/publications/asymm-long.pdf>, July 2019.
18. Jeremy Gibbons and Graham Hutton. Proof methods for corecursive programs. *Fundamenta Informaticae*, 66(4):353–366, April/May 2005.
19. Fabien Giesen. Simple rANS encoder/decoder. https://github.com/rygorous/ryg_rans, February 2014.
20. David A. Huffman. A method for the construction of minimum-redundancy codes. *Proceedings of the IRE*, 40(9):1098–1101, September 1952.
21. Joint Photographic Experts Group. Overview of JPEG XL. <https://jpeg.org/jpegxl/>.
22. Timothy B. Lee. Inventor says Google is patenting work he put in the public domain. *Ars Technica*, October 2018.

- <https://arstechnica.com/tech-policy/2018/06/inventor-says-google-is-patenting-work-he-put-in-the-public-domain/>.
23. Simon Peyton Jones. *The Haskell 98 Language and Libraries: The Revised Report*. Cambridge University Press, 2003.
 24. Daniel Reiter Horn and Jongmin Baek. Building better compression together with DivANS. <https://blogs.dropbox.com/tech/2018/06/building-better-compression-together-with-divans/>, June 2018.
 25. Jorma J. Rissanen and Glen G. Langdon. Arithmetic coding. *IBM Journal of Research and Development*, 23(2):149–162, March 1979.
 26. Julian Seward. *bzip2*. <https://en.wikipedia.org/wiki/Bzip2>, 1996.
 27. Barney Stratford. *A Formal Treatment of Lossless Data Compression Algorithms*. DPhil thesis, University of Oxford, 2005.
 28. Varmo Vene and Tarmo Uustalu. Functional programming with apomorphisms (corecursion). *Proceedings of the Estonian Academy of Sciences: Physics, Mathematics*, 47(3):147–161, 1998. 9th Nordic Workshop on Programming Theory.
 29. Ian H. Witten, Radford M. Neal, and John G. Cleary. Arithmetic coding for data compression. *Communications of the ACM*, 30(6):520–540, June 1987.