

The ZPAQ Open Standard Format for Highly Compressed Data - Level 1

Draft version 0.34

Last updated Feb. 21, 2009 by Matt Mahoney

Abstract

This document specifies the ZPAQ open standard format for highly compressed byte string data. The format supports memory to memory compression, single file compressors, and archivers, either solid or with independently compressed files. The compression algorithm uses a bitwise context mixing model (like PAQ8 [1]) followed by arithmetic decoding, packing into bytes, and post-processing transforms. The format supports future improvements in the compression of arbitrarily complex data types without loss of compatibility because the compressed stream contains instructions for specifying the model architecture, and byte-code programs to compute arbitrarily complex contexts and transforms. The standard is open, in that no license is required to develop or use software that reads or writes ZPAQ compliant data.

Scope

This document describes a proposed level 1 compression, which has not been finalized. The format is currently experimental and subject to change. When finalized, all compliant level 1 implementations will be compatible with each other. Higher levels will be backward compatible back to level 1. A file or array produced by any level 1 implementation will be readable by any other level 1 implementation. Level n implementations, where $n > 1$, will be able to read data produced by any level m implementation, where $1 \leq m \leq n$.

The current implementation is level 0. Different versions of level 0 are not compatible with each other or with level 1 in either direction.

1. Introduction

The ZPAQ open standard specifies a compressed representation for one or more byte (8 bit value) sequences. A ZPAQ stream consists of a sequence of *blocks* that can be decompressed independently. A block consists of a sequence of *segments* that must be decompressed sequentially from the beginning of the block. Each segment might represent an array of bytes in memory, a file, or a contiguous portion of a file.

ZPAQ uses a context mixing data compression algorithm based on the PAQ series (PAQ8, PAQ9, LPAQ) [1]. The decompressed stream is decoded one bit at a time, packed into bytes, and then transformed through a post-processor to undo transforms that were intended to make the data more compressible. A bit is decoded by a *model*, which predicts (assigns a probability to) the next bit based on previously decoded bits, an arithmetic decoder which takes the prediction and the compressed data and outputs the bit. The bit is fed back to the model so that it can refine future predictions.

The decoded data for each block starts with a flag to indicate whether it should be output directly or post-processed. In the latter case, it consists of a program followed by its input data. The output of this program is the output of the decompressor. In either case, the output may then be divided into separate arrays or files as described in the segment headers.

A model is a set of *components* that make independent predictions given a *context* and/or by combining the predictions of other components. Each component has a context that is computed from the previously decoded bits by a program described in the block header. For example, the context could be a hash of the last 20 bits. Up to 255 components of the following types may be connected in an arbitrary manner for each block:

- **CONST** - The prediction is a fixed value.
- **CM** - Context Model - A table maps the context to a prediction (initially equal for 0 and 1) and a counter. After a bit is decoded, the prediction is adjusted in proportion to the prediction error and inversely proportional to the count, and the count is incremented up to a specified limit.
- **ICM** - Indirect Context Model - A hash table maps the context to a bit history, a state representing bounded counts of previously seen 0 and 1 bits (initially both 0) and the exact sequence of recently seen bits (initially empty). A second table maps the history to a prediction and a count like in a CM. After a bit is decoded, the history is updated in addition to the prediction and count.
- **MATCH** - An index maps the context to the most recent occurrence of the same context in the output buffer. The bits following the match are predicted with a confidence that depends on the length of the match. The index is updated every 8 bits.
- **AVG** - Two predictions are combined by weighted averaging. The model specifies the weights.
- **MIX2** - Two predictions are combined by weighted averaging. The weights (initially 1/2) are selected from a table by context. After decoding, the weights are adjusted in proportion to the prediction error times the input times a specified learning rate. This has the effect of favoring the most accurate components in each possible context.
- **MIX** - A mixer like MIX2 but with input from a contiguous block of m components. The weights are initially $1/m$.
- **IMIX2** - Mixing with indirect contexts. A table maps the context to a bit history as with an ICM. The history is used as the context to a MIX2. An IMIX2 with one input CONST improves the prediction of the other input, similar to SSE (secondary symbol estimation). The PAQ9A [1] model is a cascade of these IMIX2 with increasingly higher context orders. After decoding, the bit history and weights are updated as with an ICM and MIX2.
- **SSE** - Secondary Symbol Estimation. SSE takes a quantized input prediction and a context and outputs a new (interpolated) prediction from a table. After decoding, the nearest table entry is adjusted to reduce the prediction error as with a CM.

One possible architecture is shown below. This example is similar to PAQ8.




Fig. 1. A typical ZPAQ decompression architecture

Component input and output predictions are expressed as log odds. If a component predicts that a 0 or 1 will occur with probability p_0 and p_1 respectively, then the output is $p = \text{stretch}(p_1) = \ln p_1/p_0$. Predictions are computed along a fixed sequence of components with input from earlier components. The input to the arithmetic decoder is p_1 from the last component, where $p_1 = \text{squash}(p) = 1/(1 + e^{-p})$ is the inverse of $\text{stretch}(p_1)$. This has the effect of weighting models with high confidence predictions (large magnitudes) more heavily.

Contexts hashes are computed by a program described in the block header and interpreted once every 8 bits during decompression, and combined with the partially decoded current byte to form a complete context. The program runs on a virtual machine which takes the last decoded byte as input and writes the context hashes to the components. The language, called ZPAQL, resembles an assembly language so that it can be implemented efficiently while allowing for arbitrarily complex contexts. The same language is used for post-processing. Both programs are called once for each byte of input data.

The compressed data represents a high precision binary number in the range (0...1) with the most significant bits of the fraction first. The arithmetic decoder maintains a range (low...high), initially (0...1), which bounds the data and shrinks as decoding proceeds. The decoder receives a prediction and splits the range into two parts in proportion to p_0 and p_1 . Whichever part now contains the compressed data determines the decoded bit and the new range. Arithmetic coding effectively codes each bit Y at a cost very close to the theoretical limit of $\log_2 1/p_Y$ bits. To mark the end of the data, each decoded byte is preceded by a EOS (end of segment) bit, which is 1 after the last byte, coded with p_1 very near 0. The decoder is designed so that the end of the coded segment can also be found by quickly scanning without decompressing the data.

The model and post-processor are initialized at the start of each block and maintain state information across segments. The arithmetic coder is initialized at the start of each segment. Segment boundaries are invisible to the model; the block appears as a continuous stream of bytes. The purpose of segments is to allow decompression to different destinations (e. g. to different files) and to signal the post-processor at the end of each file.

2. Syntax

A ZPAQ level 1 stream shall have the syntax described in this section. In this description, the notation " $X ::= Y Z$ " means that symbol X is composed of Y followed by Z . Terminal symbols (those not expanded further) are single bytes with range (0...255) inclusive. Symbols in parenthesis are nonterminal. The notation $X[0...n]$ means an array of X repeated n times, individually $X[0]...X[n-1]$. When X is used as a number, it means an n -byte number in base 256, least significant byte (LSB) first, in the range (0...256 ^{n} -1). The notation $X[0...]$ means X repeated 0 or more times. A string enclosed in double quotes is interpreted as a sequence of ASCII bytes, e.g. "zPQ" means 122 80 81. The notation $X=n$ means that X is a symbolic constant with value n (in 0...255). The notation $(X | Y)$ means either X or Y .

ZPAQ ::= (block)[0...]

block ::= "zPQ" LEVEL=1 HPROG=1 hsize[0..1] (header) (segment)[0...] EOB=255

(Also, header must be hsize bytes long.)

header ::= hh hm ph pm n(1...255) (comp)[0...n-1] END=0 (hcomp) END=0

comp[i] ::= (

CONST=1 c

| CM=2 sizebits limit

| ICM=3 sizebits

| MATCH=4 sizebits

| AVG=5 j(0...i-1) k(0...i-1) wt

| MIX2=6 sizebits j(0...i-1) k(0...i-1) rate mask

| MIX=7 sizebits j(0...i-1) m(1...i-j) rate mask

| IMIX2=8 sizebits j(0...i-1) k(0...i-1) wt rate

| SSE=9 sizebits j start limit mask)

segment ::= 1 (filename) 0 (comment) 0 RESERVED=0 (ecd) EOS=254

filename ::= c(1...255)[0...]

comment ::= c(1...255)[0...]

ecd ::= c[0...] 0 0 0 0

(Also, there are never more than 3 consecutive bytes of c = 0)

(ecd) (entropy coded data) is arithmetic decoded using the context mixing model described by (comp) with contexts computed by the program (hcomp), which is executed once for each decoded byte with that byte as input. That program is allocated $4*2^{hh} + 2^{hm}$ bytes of memory. The decoded data (dd) for each block has the format:

dd ::= (PASS=0 output[0...] | PROG=1 plen[0...1] (pcomp) pdata[0...])

(Also, pcomp must be plen bytes long.)

If dd[0] = PASS then output[0...] is output to the destination specified by (filename) in each segment header. Otherwise, the program (pcomp) is run once for each byte of pdata with that byte as input. The output of the program is output to the destination specified by (filename). (pcomp) is allocated $4*2^{ph} + 2^{pm}$ bytes of memory. Both (hcomp) and (pcomp) have the same syntax:

pcomp ::= (opcode)[0...]

hcomp ::= (opcode)[0...]

Opcodes are one to three bytes. Valid opcodes are shown in Table 1 in section 6. A program may have invalid opcodes as long as they are not executed.

A ZPAQ level 1 compliant stream shall conform to this syntax. In addition, a stream where either program does not halt, or executes ERROR or any reserved instruction, or where the program counter goes outside the range of the program is not compliant.

A compliant stream may exceed the memory requirements of a decompressor. We say that a program is *compliant with memory limit M* if it accepts any stream that allocates at most M bytes as described in any block header. The memory requirement is $4*2^{hh} + 2^{hm} + 4*2^{ph} + 2^{pm}$ plus the following by component in (comp), in bytes, where $SIZE = 2^{sizebits}$.

- For each CM, $4*SIZE$
- For each ICM, $64*SIZE + 1024$
- For each MATCH, $8*SIZE$

- For each MIX2, $8 \cdot \text{SIZE}$
- For each MIX, $4 \cdot \text{SIZE} \cdot m$
- For each IMIX2, $64 \cdot \text{SIZE} + 2048$
- For each SSE, $128 \cdot \text{SIZE}$.

3. Decoding

A ZPAQ decoder has the following state information:

- A virtual machine HCOMP as described by (*hcomp*), containing the externally accessible context array $H[0 \dots 2^{hh}-1]$ each in $(0 \dots 2^{32}-1)$.
- An array of n ($1 \dots 255$) components, $\text{COMP}[0 \dots n-1]$ as described by (*comp*), such that $\text{COMP}[i]$ takes context hash $H[i]$ (and also the output of $\text{COMP}[0 \dots i-1]$) as input.
- An array of n predictions, $P[0 \dots n-1]$, where $P[i]$ in $(-2^{31} \dots 2^{31}-1)$ is the output of $\text{COMP}[i]$. $P[i]$ represents a belief by $\text{COMP}[i]$ that the next bit will be a 1 with probability $1/(1 + e^{-P[i]/256})$.
- A partially or fully decoded byte C8 in $(1 \dots 511)$.
- An arithmetic decoder (section 4).
- POST, a post-processor (section 5).

Throughout this specification, we will use the convention that array indexes are modulo the array size. Thus if $hh = 8$, then $H[256] = H[0]$.

Let the function $\text{predict}(\text{COMP}[i])$ assign a prediction to $P[i]$. The function $\text{update}(\text{COMP}[i], Y)$ updates the state of the component with decoded bit Y ($0 \dots 1$) in a way designed to reduce future prediction errors. The function $\text{decode}(p)$ returns Y from the arithmetic coder given a probability p in $[0,1)$, as described in section 4. The function $\text{write}(\text{POST}, C)$ writes byte C ($0 \dots 255$) to the post-processor (section 5). The overall structure of the decompression algorithm is:

$\text{decompress}() =$

```

    For each block do
        For i in  $(0 \dots n-1)$  initialize  $\text{COMP}[i](\text{comp}[i])$ 
        Initialize  $\text{HCOMP}(hh, hm)$ 
        Initialize  $P[0 \dots 255] := 0$ ,  $C8 := 256$ 
        Initialize  $\text{POST}(ph, pm)$ 
        For each segment
            Select output, depending on filename
            Initialize arithmetic decoder (sec. 4)
            While  $\text{decode}(0) = 0$  do
                 $C8 := 1$ 
                While  $C8 < 256$  do
                    For i in  $(0 \dots n-1)$  do

```

```

P[i] := predict(COMP[i], P[0...i-1], H[i], C8)
Y := decode(squash(P[n-1]) + 0.5) / 4096 (Y in 0...1)
For i in (0...n-1) do
    update(COMP[i], P[i], Y)
C8 := C8 * 2 + Y
write(POST, C8 - 256) (sec. 5)
H := run(HCOMP, hcomp, C8 - 256) (sec. 6)
write(POST, 232-1) (end of segment)

```

Define squash(x) =

```

If x > 2047 then return 4095
Else if x < -2048 then return 0
Else return floor(4096 / (1 + e-x/256))

```

Squash() is the approximate inverse of stretch(). It is not exact due to integer roundoff. The exact definition is:

stretch(x) = round(256 * ln((x + 0.5)/(4095.5 - x))), x in (0...4095)

where round(x) = floor(x + 1/2), and floor(x) is the largest integer not greater than x.

If correctly implemented these functions should satisfy the following computations:

```

Σi=0...4095 3i stretch(i) = 1625980894 (mod 232)
Σi=0...4095 3i squash(i-2048) = 1955773538 (mod 232)
squash(2047) = 4094          stretch(4095) = 2307
squash(0) = 2048            stretch(2048) = 0
squash(-2048) = 1           stretch(0) = -2307

```

P[] thus represents stretched probabilities. Squash(P[i]) is in (0...4095) and represents the belief by COMP[i] that the next bit will be a 1 with probability (squash(P[i]) + 0.5) / 4096.

The following functions are defined for each component:

- initialize(COMP[i]) sets the initial state at the start of a block.
- predict(COMP[i], H[i], P[0...i-1], C8) writes a prediction to P[i].
- update(COMP[i], P[0...i-1], Y) modifies the state of COMP[i] to reduce the prediction error for the decoded bit Y.

The (comp) instructions are as follows (with unused parameters omitted):

3.1. CONST c

There is no state to initialize or update. The prediction is P[i] := (c - 128) * 16

3.2. CM sizebits limit

A context model uses a table CM to map a context into a prediction. When updated, it adjusts the table entry to reduce the prediction error. To control the learning rate, it counts predictions in each context in a table CMCOUNT. Initialize:

SIZE := 2^{sizebits}

CM[0...SIZE-1] := 2^{21} (a 22 bit probability in $(0...2^{22}-1)$)

CMCOUNT[0...SIZE-1] := 0 (count, range 0...1023)

predict(COMP[i], H[i], C8) =

CXT := H[i] XOR hmap4(C8)

P[i] := stretch(floor(CM[CXT] / 2^{10}))

update(COMP[i], Y) =

train(i, CM[CXT], CMCOUNT[CXT], limit * 4)

The train() function updates the prediction in CM[CXT] to reduce the prediction error in inverse proportion to CMCOUNT[CXT], then updates the count up to LIMIT. The function is also used by other models:

train(i, T, TCOUNT, LIMIT) =

ERROR := Y * 4095 - squash(P[i])

T := T + ERROR * round(floor($2^{19} / (2 * \text{TCOUNT} + 1)$) / 2^8)

TCOUNT := min(LIMIT, TCOUNT + 1)

hmap4() is a function intended to improve cache locality on 64 byte aligned arrays.

hmap4(C8) =

If (C8 < 16) then return C8

Else if (C8 < 2^{5+0}) then return floor(C8 / 2^0) * 16 + (C8 mod 2^0) + 2^0

Else if (C8 < 2^{5+1}) then return floor(C8 / 2^1) * 16 + (C8 mod 2^1) + 2^1

Else if (C8 < 2^{5+2}) then return floor(C8 / 2^2) * 16 + (C8 mod 2^2) + 2^2

Else if (C8 < 2^{5+3}) then return floor(C8 / 2^3) * 16 + (C8 mod 2^3) + 2^3

hmap4() has the effect of splitting the partially decoded byte into two 4-bit nibbles. After the first nibble is fully decoded, it occupies bits 7...4 of the output with bit 8 set to 1.

0000xxxx -> 00000xxxx

0001xxxx -> 1xxxx0001

001xxxxx -> 1xxxx001x

01xxxxxx -> 1xxxx01xx

1xxxxxxx -> 1xxxx1xxx

3.3. ICM sizebits

An indirect context model uses a hash table HT to map a context to a bit history, and then a direct lookup table CM to map the history to a probability. When a bit is decoded, the history is updated to reflect the new bit, and the history map is adjusted to reduce the prediction error. Initialize:

SIZE := $4 * 2^{\text{sizebits}}$ (size of context map)

HT[0...SIZE-1][0...15] := 0 (checksum and 15 histories, all in 0...255)

CM[0...255] := 2^{21} (history map, probability as a 22 bit value in $0...2^{22}-1$)

CM_COUNT[0...255] := 0 (history map count, range 0...1023)

The next bit is predicted by computing a hash index HI from H[i] and C8 and looking up in the history BH in the hash table HT. The low bits of HI are used as the index, and the next higher 8 bits are used as a checksum to detect (most) hash collisions. If a hash confirmation is not found among 3 adjacent elements, then the lowest priority element is replaced. Then HT is mapped to P[i] through CM.

predict(COMP[i], H[i], C8) =

If C8 = 1 or C8 in (16...31) then HI := find(HT, H[i] + 16 * C8) (first index into HT)

BI := hmap4(C8) (mod 16) (second index in HT, in 1...15)

BH := HT[HI][BI] (bit history in 0...255)

P[i] := stretch(floor(CM[BH] / 2^{10}))

find(HT, CXT) finds the hash index for CXT, replacing an element in HT if needed. It is defined as:

find(HT, CXT) =

CHK := floor(CXT / SIZE) (mod 256) (checksum for hash confirmation)

H0 := CXT mod SIZE (hash index)

H1 := H0 XOR 1 (candidate locations)

H2 := H0 XOR 2

If HT[H0][0] = CHK then return H0

Else if HT[H1][0] = CHK then return H1

Else if HT[H2][0] = CHK then return H2

Else if HT[H0][1] ≤ HT[H1][1] and HT[H0][1] ≤ HT[H2][1] then

HT[H0] := (CHK, 0[1...15]), return H0

Else if HT[H1][1] ≤ HT[H2][1] then

HT[H1] := (CHK, 0[1...15]), return H1

Else CM[H2] := (CHK, 0[1...15]), return H2

HT uses element 0 of each row as a confirmation checksum and element 1 as a priority. Element 1 represents the bit history for a context ending on a 4 bit boundary. This is the history updated most frequently in the row, because the other histories are for contexts that are 1 to 3 bits longer. The histories are represented by states sorted by increasing $n_0 + n_1$, the sum of the 0 and 1 bit counts. Thus, the cache replacement policy is roughly LFU (least frequently used).

When a bit Y is decoded, the bit history is updated in CM, HM[BH] is adjusted to reduce the prediction

error in proportion to $1/\text{HM_COUNT}[\text{BH}]$, and the count is incremented up to limit * 4.

update(COMP[i], Y) =

HT[HI][BI] := NEXT[BH][Y]

train(i, CM[BH], CMCOUNT[BH], 1023) (sec. 3.2)

NEXT[0...255][0...1] is defined by the following table:

```
{ { 1, 2}, { 3, 5}, { 4, 6}, { 7, 10}, { 8, 12}, { 9, 13}, { 11, 14}, // 0
{ 15, 19}, { 16, 23}, { 17, 24}, { 18, 25}, { 20, 27}, { 21, 28}, { 22, 29}, // 7
{ 26, 30}, { 31, 33}, { 32, 35}, { 32, 35}, { 32, 35}, { 32, 35}, { 34, 37}, // 14
{ 34, 37}, { 34, 37}, { 34, 37}, { 34, 37}, { 34, 37}, { 36, 39}, { 36, 39}, // 21
{ 36, 39}, { 36, 39}, { 38, 40}, { 41, 43}, { 42, 45}, { 42, 45}, { 44, 47}, // 28
{ 44, 47}, { 46, 49}, { 46, 49}, { 48, 51}, { 48, 51}, { 50, 52}, { 53, 43}, // 35
{ 54, 57}, { 54, 57}, { 56, 59}, { 56, 59}, { 58, 61}, { 58, 61}, { 60, 63}, // 42
{ 60, 63}, { 62, 65}, { 62, 65}, { 50, 66}, { 67, 55}, { 68, 57}, { 68, 57}, // 49
{ 70, 73}, { 70, 73}, { 72, 75}, { 72, 75}, { 74, 77}, { 74, 77}, { 76, 79}, // 56
{ 76, 79}, { 62, 81}, { 62, 81}, { 64, 82}, { 83, 69}, { 84, 71}, { 84, 71}, // 63
{ 86, 73}, { 86, 73}, { 44, 59}, { 44, 59}, { 58, 61}, { 58, 61}, { 60, 49}, // 70
{ 60, 49}, { 76, 89}, { 76, 89}, { 78, 91}, { 78, 91}, { 80, 92}, { 93, 69}, // 77
{ 94, 87}, { 94, 87}, { 96, 45}, { 96, 45}, { 48, 99}, { 48, 99}, { 88,101}, // 84
{ 88,101}, { 80,102}, {103, 69}, {104, 87}, {104, 87}, {106, 57}, {106, 57}, // 91
{ 62,109}, { 62,109}, { 88,111}, { 88,111}, { 80,112}, {113, 85}, {114, 87}, // 98
{114, 87}, {116, 57}, {116, 57}, { 62,119}, { 62,119}, { 88,121}, { 88,121}, // 105
{ 90,122}, {123, 85}, {124, 97}, {124, 97}, {126, 57}, {126, 57}, { 62,129}, // 112
{ 62,129}, { 98,131}, { 98,131}, { 90,132}, {133, 85}, {134, 97}, {134, 97}, // 119
{136, 57}, {136, 57}, { 62,139}, { 62,139}, { 98,141}, { 98,141}, { 90,142}, // 126
{143, 95}, {144, 97}, {144, 97}, { 68, 57}, { 68, 57}, { 62, 81}, { 62, 81}, // 133
{ 98,147}, { 98,147}, {100,148}, {149, 95}, {150,107}, {150,107}, {108,151}, // 140
{108,151}, {100,152}, {153, 95}, {154,107}, {108,155}, {100,156}, {157, 95}, // 147
{158,107}, {108,159}, {100,160}, {161,105}, {162,107}, {108,163}, {110,164}, // 154
{165,105}, {166,117}, {118,167}, {110,168}, {169,105}, {170,117}, {118,171}, // 161
{110,172}, {173,105}, {174,117}, {118,175}, {110,176}, {177,105}, {178,117}, // 168
{118,179}, {110,180}, {181,115}, {182,117}, {118,183}, {120,184}, {185,115}, // 175
{186,127}, {128,187}, {120,188}, {189,115}, {190,127}, {128,191}, {120,192}, // 182
{193,115}, {194,127}, {128,195}, {120,196}, {197,115}, {198,127}, {128,199}, // 189
{120,200}, {201,115}, {202,127}, {128,203}, {120,204}, {205,115}, {206,127}, // 196
{128,207}, {120,208}, {209,125}, {210,127}, {128,211}, {130,212}, {213,125}, // 203
{214,137}, {138,215}, {130,216}, {217,125}, {218,137}, {138,219}, {130,220}, // 210
{221,125}, {222,137}, {138,223}, {130,224}, {225,125}, {226,137}, {138,227}, // 217
```

```

{130,228},{229,125},{230,137},{138,231},{130,232},{233,125},{234,137}, // 224
{138,235},{130,236},{237,125},{238,137},{138,239},{130,240},{241,125}, // 231
{242,137},{138,243},{130,244},{245,135},{246,137},{138,247},{140,248}, // 238
{249,135},{250, 69},{ 80,251},{140,252},{249,135},{250, 69},{ 80,251}, // 245
{140,252},{ 0, 0},{ 0, 0},{ 0, 0}}; // 252

```

The NEXT table is the same as used in the LPAQ1 file compressor [1], as well as most versions of PAQ8. The code for generating the table can be found in all versions of PAQ8, which also shows the values of n_0 and n_1 for each state.

The meaning of the table is as follows. Each of the 252 states (253...255 are unused) represents a bit history of the form (n_0, n_1, s) where n_0 and n_1 are counts of 0 and 1 bits respectively, and s (of length $|s|$) is the exact sequence of the last $|s|$ bits observed. State 0 is the initial state $(0, 0, "")$. Given state $BH = (n_0, n_1, "s")$, then $NEXT[BH][0] = (n_0+1, \text{discount}(n_1), "s0")$ and $NEXT[BH][1] = (\text{discount}(n_0), n_1+1, "s1")$ subject to the following restrictions, which are symmetric with regard to n_0 and n_1 :

If $n_0 = 0$ then $n_1 \leq 41$	If $n_1 = 0$ then $n_0 \leq 41$.
If $n_0 = 1$ then $n_1 \leq 40$	If $n_1 = 1$ then $n_0 \leq 40$.
If $n_0 = 2$ then $n_1 \leq 12$	If $n_1 = 2$ then $n_0 \leq 12$.
If $n_0 = 3$ then $n_1 \leq 5$	If $n_1 = 3$ then $n_0 \leq 5$.
If $n_0 = 4$ then $n_1 \leq 4$	If $n_1 = 4$ then $n_0 \leq 4$.
If $n \geq 15$ then $ s = 0$ ($s = ""$).	
Else if $n \geq 5$ then $ s = 1$ ($s = "0"$ or $"1"$).	

If the next state exceeds the bounds on $|s|$ then the oldest bits are discarded.

The $\text{discount}()$ function reduces the count that was not incremented. This has the effect of favoring recent history over older. $\text{discount}()$ is defined by the following mapping:

```

(0...6) -> (0...6) (no change)
(7) -> 6
(8...10) -> 7
(11...13) -> 8
(14...17) -> 9
(18...22) -> 10
(23...29) -> 11
(30...38) -> 12
(39...41) -> 13

```

If the counts of the new state are out of range, then the next state is the nearest state which approximately preserves the ratio of n_0 to n_1 . For the case of $n_1 \geq n_0$, the counts are updated as follows. For the case of $n_1 < n_0$, the update is symmetric with regard to n_0 and n_1 .

```

If  $n_0 = 0$  then  $n_1 := \min(n_1, 41)$ 
Else if  $n_0 = 1$  then  $n_1 := \min(n_1, 40)$ 

```

Else while n_0 and n_1 are out of range do

$n_1 := \text{floor}(n_1 * ((n_0 - 1) + \text{floor}(n_0 / 2)) / n_0)$ (approximately $n_1 - n_1/n_0$)

$n_0 := n_0 - 1$

3.4. MATCH sizebits

A match model finds the most recent context match in an output buffer and predicts the next bit as a function of the length of the match. The match is maintained until a bit mismatch is found. On each byte boundary, if there is no current match then it looks up the context in the index to find a new one. On each byte boundary it updates the output buffer and the index. Initialize:

$\text{SIZE} := 2^{\text{sizebits}}$

$\text{INDEX}[0 \dots \text{SIZE}-1] := 0$

$\text{OFFSET} := 0$ (distance back to match)

$\text{LEN} := 0$ (length of match in bytes, up to 255)

$\text{BUF}[0 \dots \text{SIZE} \cdot 4 - 1] := 0$, $\text{BUF}[0] := 1$ (decoded data buffer in $\text{BUF}[0 \dots \text{POS}-1]$ (0...255))

$\text{POS} := 0$ (number of decoded bytes)

$\text{BP} := 0$ (number of decoded bits after last full byte, 0...7)

$\text{predict}(\text{COMP}[i]) =$

If $\text{LEN} = 0$ then $P[i] := 0$

Else

$\text{BIT} := \text{floor}(\text{BUF}[\text{POS} - \text{OFFSET}] / 2^{7-\text{BP}}) \pmod{2}$ (predicted bit)

If $\text{BIT} = 1$ then $P[i] := \text{stretch}(4096 - \text{floor}(256 / \text{LEN}))$

Else $P[i] := \text{stretch}(\text{floor}(256 / \text{LEN}))$

$\text{update}(\text{COMP}[i], Y) =$

If $\text{BIT} \neq Y$ then $\text{LEN} := 0$

$\text{BUF}[\text{POS}] := \text{BUF}[\text{POS}] * 2 + Y \pmod{256}$

$\text{BP} := \text{BP} + 1$

If $\text{BP} = 8$ then (a byte was fully decoded)

$\text{POS} := \text{POS} + 1$

$\text{BP} := 0$

If $\text{LEN} = 0$ then (look for a match)

$\text{OFFSET} := \text{POS} - \text{INDEX}[\text{H}[\text{h}]]$

If $\text{OFFSET} \neq 0 \pmod{\text{BUFSIZE}}$ then

While $\text{LEN} < 255$

```

        and BUF[POS - LEN - 1] = BUF[POS - LEN - OFFSET - 1]
    LEN := LEN + 1
Else if (LEN < 255) then LEN := LEN + 1
INDEX[H[i]] = POS

```

3.5. AVG j k wt

There is no state to initialize or update.

```

predict(P[0...i-1]) =
    P[i] := floor(int32(P[j] * wt + P[k] * (256-wt)) / 256).

```

Integer overflow is unlikely to occur, but is included for strict compatibility. It defines normal overflow for 32 bit signed arithmetic:

$$\text{int32}(x) = ((x + 2^{31}) \bmod 2^{32}) - 2^{31}$$

3.6. MIX sizebits j m rate mask

A MIX adaptively combines m predictions by weighted averaging, where the weights are selected by a context. After a bit is decoded, the weights are adjusted to favor the most accurate models. Initialize:

```

SIZE := 2sizebits
WT[0...SIZE-1][0...m-1] := floor(216/m)

```

The output prediction is a weighted sum of inputs $P[j...j+m-1]$

```

predict(COMP[i], H[i], P[0...i-1], C8) =
    CXT := H[i] + (C8 AND mask)
    P[i] := round((∑k in (j...j+m-1) int32(round(WT[CXT][k] / 256) * P[k])) / 256)

```

After decoding, the weights are adjusted to favor the most accurate input models for the given context.

```

update(COMP[i], P[0...i-1], Y) =
    ERROR := (Y * 4095 - squash(P[i])) * rate
    For k in (j...j+m-1) do
        WT[CXT][k] := int32(WT[CXT][k] + round(ERROR * P[k]) / 216)

```

3.7. MIX2 sizebits j k rate mask

A MIX2 is a MIX with $m = 2$ inputs, $P[j]$ and $P[k]$ instead of $P[j...j+m-1]$.

3.8. IMIX2 sizebits j k wt rate

An indirect mixer maps a context to a bit history (like ICM), which is then used as the context for a MAP2. The MAP2 takes $P[j]$ and $P[k]$ as inputs. The weights are initialized to $\text{wt}/256$ for $P[j]$ and $(256-\text{wt})/256$ for $P[k]$. (Typically, one input is a CONST with a small initial weight). The other parameters set the bit history table size, bit mask for context $\text{BUF}[\text{POS}]$, and the MAP2 learning rate. Initialize:

```

SIZE := 4 * 2sizebits (size of hash table)
HT[0...SIZE-1][0...15] := 0 (checksum and 15 histories, all in 0...255)
WT[0...255][0] := 256 * wt (input for P[j])
WT[0...255][1] := 256 * (256 - wt) (input for P[k])
predict(COMP[i], H[i], P[0...i-1], C8) =
  If C8 = 0 or C8 in (16...31) then HI := find(HT, H[i] + 16 * C8) (first index into HT)
  BI := hmap4(C8) (mod 16) (second index in HT, in 1...15)
  BH := HT[HI][BI] (bit history in 0...255)
  P[i] := round((
    int32(round(WT[BH][0] / 256) * P[j]) +
    int32(round(WT[BH][1] / 256) * P[k])) / 256)
update(COMP[i], P[0...i-1], Y) =
  HT[HI][BI] := NEXT[BH][Y]
  ERROR := (Y * 4095 - squash(P[i])) * rate
  WT[CXT][0] := int32(WT[CXT][0] + round(ERROR * P[j] / 216))
  WT[CXT][1] := int32(WT[CXT][1] + round(ERROR * P[k] / 216))

```

3.9. SSE sizebits j start limit mask

A secondary symbol estimator (SSE) takes an input prediction $P[j]$ quantized to 32 levels and a context $H[i]$ and outputs a new prediction. The prediction is interpolated between the two nearest quantized value. The closer of those two points is then updated. The table SM is initialized to output the same prediction as the input for all contexts. Each element is associated with a count SMCOUNT in (start...limit*4) that determines the update rate. Initialize:

```

SIZE := 2sizebits
For k in (0...31)
  SM[0...SIZE-1][k] := squash(k*256 - 3968) * 210
  SMCOUNT[0...SIZE-1][k] := start
predict(COMP[i]) =
  CXT := H[h] + (BUF[POS] AND mask)
  PQ := min(7935, max(0, (P[j] + 3968)))
  W := PQ (mod 256) (interpolation weight)
  PQ := floor(PQ / 256) (quantized to 0...30)
  P[i] := stretch(floor((SM[CXT][PQ]*(256-W) + SM[CXT][PQ+1]*W) / 218))
  If W ≥ 128 then PQ := PQ + 1

```

When bit Y is decoded, the prediction is adjusted to reduce the prediction error in inverse proportion to its count, and the count is incremented to a maximum of limit*4.

```

update(COMP, Y) =
    train(i, SM[CXT][PQ], SMCOUNT[CXT][PQ], limit*4)

```

4. Arithmetic Decoder

The arithmetic decoder receives bit predictions ($PR = \text{squash}(P[n-1]) + 0.5/4096$), and the compressed input stream and outputs uncompressed bits, Y . The end of segment is decoded with $PR = 0$. All other bits are decoded with PR equal to odd multiple of $1/8192$ between 0 and 1.

The decoder state is initialized:

```

LOW := 1(in  $1 \dots 2^{32}-1$ )
HIGH :=  $2^{32}-1$  (in  $0 \dots 2^{32}-1$ ,  $HIGH > LOW$ )
CURR := 0
Do 4 times: CURR := CURR * 256 + next_byte(ecd)

```

decode(PR) returns a bit as follows:

```

PR := PR * 8192 (an integer in  $0 \dots 8191$ )
MID := LOW + floor((HIGH - LOW) * PR / 8192)
If CURR ≤ MID then Y := 1, HIGH := MID
else Y := 0, LOW := MID + 1
While floor(LOW /  $2^{24}$ ) = floor(HIGH /  $2^{24}$ ) do
    LOW := LOW * 256 (mod  $2^{32}$ )
    If LOW = 0 then LOW := 1
    HIGH := HIGH * 256 + 255 (mod  $2^{32}$ )
    CURR := CURR * 256 + next_byte(ecd) (mod  $2^{32}$ )
Return Y

```

next_byte() reads one byte of the compressed data, *ecd*. When decoding ends, next_byte() will have read the 4 trailing 0 bytes so that $LOW > 0$, $CURR = 0$, $HIGH = 2^{31}-1$. At all other times, $LOW \leq CURR \leq HIGH$ and $LOW < HIGH$. The next byte to read would be EOS.

5. Post Processing

Recall that a block decoded as described in sections 3 and 4 has the following syntax:

```

dd ::= (PASS=0 output[0...] | PROG=1 plen[0...1] (pcomp) pdata[0...] )

```

This data is written to a post-processor in the decoding algorithm in section 3 by calling write(POST, C) for each byte C in *dd*. POST has the following state:

```

PCOMP, a virtual machine, initialized PCOMP(ph, pm)
PBUF, an input buffer string, initialized to ""

```

write(POST, C) =

```

    If PBUF = "" then append C to PBUF

```

Else if $PBUF[0] = \text{PASS}$ then output C

Else if $|PBUF| < 3$ or $|PBUF| < plen + 3$ then append C to $PBUF$ (where $plen = PBUF[1..2]$)

Else run($PCOMP, pcomp, C$) (run program $pcomp$ (in $PBUF[3..plen+2]$) with input C)

When the first byte of dd is $PROG$, the output of $run()$ is the output of the decompressor.

6. ZPAQL

There are 2 ZPAQL virtual machines: HCOMP in the bit prediction model, and PCOMP in the post-processor. A machine COMP is initialized (at the beginning of a block):

COMP($hbits, mbits$) =

PC := 0 (program counter)

A, B, C, D := 0 (general purpose registers in $0..2^{32}-1$)

F := 0 (condition flag in $0..1$)

H[$0..2^{hbits}-1$] (memory, each element in $(0..2^{32}-1)$, initialized to 0). In HCOMP, $hbits = 8$ and

H[i] is the input to COMP[i].

M[$0..2^{mbits}-1$] (memory, each element in $(0..255)$ initialized to 0).

R[$0..255$] (memory, each element in $(0..2^{32}-1)$, initialized to 0).

A program is executed by calling $run(\text{COMP}, prog, input)$, where $prog$ is a string of opcodes as in table 1, and $input$ is an input in $(0..2^{32}-1)$.

$run(\text{COMP}, prog, input) =$

PC := 0

A := input

$prog := (prog \ 0 \ 0)$ (append two 0 bytes (ERROR opcodes))

Do forever

If PC not in $(0..|prog|-3)$ then exit with an error

If $(prog[PC] = 7 \pmod{8})$ then $OPCODE := prog[PC..PC+1]$, PC := PC + 2

Else $OPCODE := prog[PC]$, PC := PC + 1

If $OPCODE = \text{ERROR}$ or is undefined then exit with an error

Else if $OPCODE = \text{HALT}$ then return

Else execute(OPCODE)

Opcode	0	1	2	3	4	5	6	7
0 * 8	ERROR	A++	A--	A!	A=0			A=R N
1	B<>A	B++	B--	B!	B=0			B=R N
2	C<>A	C++	C--	C!	C=0			C=R N
3	D<>A	D++	D--	D!	D=0			D=R N
4	*B<>A	*B++	*B--	*B!	*B=0			JT N
5	*C<>A	*C++	*C--	*C!	*C=0			JF N
6	*D<>A	*D++	*D--	*D!	*D=0			R=A N
7	HALT	OUT		HASH	HASHD			JMP N
8	A=A	A=B	A=C	A=D	A=*B	A=*C	A=*D	A= N
9	B=A	B=B	B=C	B=D	B=*B	B=*C	B=*D	B= N
10	C=A	C=B	C=C	C=D	C=*B	C=*C	C=*D	C= N
11	D=A	D=B	D=C	D=D	D=*B	D=*C	D=*D	D= N
12	*B=A	*B=B	*B=C	*B=D	*B=*B	*B=*C	*B=*D	*B= N
13	*C=A	*C=B	*C=C	*C=D	*C=*B	*C=*C	*C=*D	*C= N
14	*D=A	*D=B	*D=C	*D=D	*D=*B	*D=*C	*D=*D	*D= N
15								
16	A+=A	A+=B	A+=C	A+=D	A+=*B	A+=*C	A+=*D	A+= N
17	A-=A	A-=B	A-=C	A-=D	A-=*B	A-=*C	A-=*D	A-= N
18	A*=A	A*=B	A*=C	A*=D	A*=*B	A*=*C	A*=*D	A*= N
19	A/=A	A/=B	A/=C	A/=D	A/=*B	A/=*C	A/=*D	A/= N
20	A%=A	A%=B	A%=C	A%=D	A%*=B	A%*=C	A%*=D	A%= N
21	A&=A	A&=B	A&=C	A&=D	A&=*B	A&=*C	A&=*D	A&= N
22	A&~A	A&~B	A&~C	A&~D	A&~*B	A&~*C	A&~*D	A&~ N
23	A =A	A =B	A =C	A =D	A =*B	A =*C	A =*D	A = N
24	A^=A	A^=B	A^=C	A^=D	A^=*B	A^=*C	A^=*D	A^= N
25	A<<=A	A<<=B	A<<=C	A<<=D	A<<=*B	A<<=*C	A<<=*D	A<<= N
26	A>>=A	A>>=B	A>>=C	A>>=D	A>>=*B	A>>=*C	A>>=*D	A>>= N
27	A==A	A==B	A==C	A==D	A==*B	A==*C	A==*D	A== N
28	A<A	A<B	A<C	A<D	A<*B	A<*C	A<*D	A< N
29	A>A	A>B	A>C	A>D	A>*B	A>*C	A>*D	A> N
30								
31								LJ N M

Table 1. ZPAQL opcodes

Note that the state of COMP is retained between runs except for A and PC. Opcodes are given in Table 1. The numeric value is 8 times the row number plus the column number. Opcodes in column 7 are two bytes where the second byte is N in (0...255). Opcode 255 (LJ) is 3 bytes.

The meaning of execute(OPCODE) is as follows. Most opcodes have the form "X op Y" where X and Y are one of A, B, C, D, *B, *C, *D, or N. A, B, C, and D are 32 bit registers with values in $(0...2^{32}-1)$. N is a number in (0...255), the second byte of a 2 byte opcode. *B means M[B]. *C means M[C]. *D means H[D]. Operations on *B and *C are modulo 256. Operations on A, B, C, D, and *D are modulo 2^{32} . As usual, indexes into M and H are modulo 2^{mbits} and 2^{hbits} respectively. Operations are as follows:

- ERROR causes the decompressor to fail (for debugging). It is equivalent to any undefined instruction, except that it is not reserved for future use.
- X++ means add 1 to X. (Note that *B++ increments *B, not B).
- X-- means subtract 1 from X.
- X! means $X := -1-X$ (complement all bits).
- X=0 means set $X := 0$. (This is a 1 byte opcode. It is equivalent to the 2 byte opcode X=N when N is 0).
- $X \diamond A$ means swap X with A. If X is *B or *C then only the low 8 bits of A are changed.
- $X=R\ N$ means $X := R[N]$.
- $R=A\ N$ means $R[N] := A$.
- JT N (jump if true) means if $F = 1$ then add $((N+128) \bmod 256) - 128$ to PC. This is a conditional jump in the range (-128...127) relative to the next instruction, e.g. JT 0 has no effect, 1...127 jumps forward and 128...255 jumps backwards from the next instruction.
- JF N (jump if false) means if $F = 0$ then add $((N+128) \bmod 256) - 128$ to PC.
- JMP N means add $((N+128) \bmod 256) - 128$ to PC (regardless of F).
- LJ N M (long jump) means $PC := N + 256 * M$, where N and M are in (0...255). This is the only 3 byte instruction.
- HALT terminates execution and returns to the calling algorithm.
- OUT means to output A. In PCOMP, A (mod 256) is written to output. In HCOMP it has no effect.
- HASH means $A := (A + *B + 512) * 773$ (a useful byte hashing function for HCOMP).
- HASHD means $*D := (*D + A + 512) * 773$.
- X=Y assigns $X := Y$.
- X+=Y adds $X := X + Y$.
- X-=Y subtracts $X := X - Y$.
- X*=Y multiplies $X := X * Y$.
- X/=Y divides: if $Y > 0$ then $X := X / Y$ else $X := 0$.
- X%=Y: if $Y > 0$ then $X := X \bmod Y$, else $X := 0$.

- $X \&= Y$ computes $X := X \text{ AND } Y$, which clears any bit in the binary representation of X if the corresponding bit of Y is 0.
- $X \&\sim Y$ computes $X := X \text{ AND NOT } Y$, which clears any bit in X that is set in Y .
- $X |= Y$ computes $X := X \text{ OR } Y$, which sets any bit in X that is set in Y .
- $X \wedge= Y$ computes $X := X \text{ XOR } Y$, which complements any bit in X that is set in Y .
- $X \ll= Y$ (left shift): $X := X * 2^{Y \bmod 32}$
- $X \gg= Y$ (right shift): $X := \text{floor}(X / 2^{Y \bmod 32})$
- $X == Y$ (equals): If $X = Y$ then $F := 1$ else $F := 0$.
- $X < Y$ (less than): If $X < Y$ then $F := 1$ else $F := 0$.
- $X > Y$ (greater than): If $X > Y$ then $F := 1$ else $F := 0$.

7. Compliance

A program that accepts any data that conforms to the requirements in sections 2 through 6 in this document is ZPAQ level 1 compliant. There is no requirement for a compliant program to behave in any particular way for any non conforming data. There is no requirement that a compressor that produces ZPAQ level 1 data shall support all of the features described. However, it is the responsibility of the compressor to produce compliant data.

A decompressor might not have enough memory to decompress a compliant stream. A decompressor is said to be compliant up to its memory limit if it will accept all streams that require less memory.

A block header begins with "zPQ" followed by LEVEL=1 to indicate the compression level supported. Future versions will use (2...127) in increasing order. Each level L shall support reading all levels in the range (1... L). Levels (128...255) are reserved for private use and are not part of this or any future standard. Likewise, HPROG=1, PROG=1, and RESERVED=0 shall not use values in the range (128...255) in future versions for the same reason. HPROG and PROG are intended to indicate the language used in the (*hcomp*) and (*pcomp*) sections, respectively.

LEVEL=0 is experimental. Different versions of level 0 programs are not compatible with each other.

A segment header has a filename and a comment string. The ZPAQ standard makes no requirements with respect to these strings. They may be empty. The filename string is intended to support archives with multiple files. The intended meaning is that the segment should be written to the named file, unless overridden by the user. An archiver may split a file into more than one segment or block. If the filename is empty, then the intended meaning is that the segment should be output to the same file as the previous segment.

The comment field has no meaning. It may contain arbitrary data, for example, to be displayed when listing the contents of an archive. For archivers that restore timestamps to extracted files, it is recommended that the time stamp be stored at the beginning of the comment in the format "YYYY/MM/DD HH:MM:SS" with optional additional data following.

8. Implementation Notes

This document does not specify a compression algorithm. However, it will generally be the case that the models (on the input side of the arithmetic decoder) will be identical for compression and

decompression. For the arithmetic coder, range splitting and normalization would be identical, except that when the high bits of the range are shifted out, they are written to the output. Immediately after coding EOS, the compressor should write an end of segment marker (0 0 0 0 254). It is not necessary to flush the encoder.

It is the responsibility of the compressor to ensure that preprocessing is exactly reversed by post-processing. The recommended way to do this is to test during compression by running both transforms and comparing with the original.

Memory requirements for a typical decompressor implementation can be calculated almost entirely from information in the block headers. Compression typically requires at least as much memory as decompression.

The design is optimized for arrays aligned on 64 byte cache line boundaries (in particular, ICM and IMIX2).

9. Intellectual Property

I (Matt Mahoney) am not aware of any patents protecting any of the techniques needed to fully implement a compression or decompression algorithm or product according to this specification. I have not filed for patents on any of the techniques described here and will not do so.

This document may be copied and distributed freely as long as the contents are not modified.

References

1. M. Mahoney, PAQ data compression programs. 2000-2008.
<http://cs.fit.edu/~mmahoney/compression/>