

FJXL and FPNGE

Very Fast SIMD lossless image encoders

Luca Versari (veluca@google.com)

Overview

In November 2021, [QOI](#) took the Internet by storm.

Main selling point: simpler and faster than PNG.

I decided to find out whether you can do better with existing formats.


The result of this effort are FJXL and FPNGE: new fast encoders 100% conformant to the JPEG XL and PNG specification, respectively.

Results

small photos: <https://goo.gl/cmHkl>


large photos: JPEG XL CTC images, excluding non-photographic images

Lossless compression speed in MP/s (more is better)




AMD Ryzen 7 5800X 8-Core Processor, 1 thread


Compression speed, wrt lodepng, higher is better ([log scale](#))



Lossless compression density in bpp (smaller is better)



Compression density, wrt lodepng, lower is better (aggregate, zoomed)



Density vs speed, compared to lodepng (aggregate)



FJXL

FJXL - JPEG XL lossless subset overview

- (optional) Palettization is applied
- Image is divided into 256x256 tiles
- Color transform (YCoCg) applied (8 -> 9 bit range)
- Channels are separated into planes
- Prediction is applied (ClampedGradient) (9 -> 10 bit range)
- LZ77 can be applied (in this encoder, just run-length encoding)
- “Hybrid UInt encoding” is applied on raw symbols, split into symbol + bits
- Prefix coding is applied on symbols

Input is 8-bit, but the added bits require switching to 16-bit arithmetic. On AVX2, that's 16 integers per vector.

FJXL - palette detection

- Hash table with 65k possible entries
- Any collision -> no palette
- Palette is sorted by luma to make the prediction still work well
- At most 256 palette entries are used

On non-palette-friendly images, this fails quickly (birthday paradox says after ~256 distinct pixels).

On palette images, encoding 1 channel rather than 4 more than compensates the cost of detection.

FJXL - prediction

ClampedGradient predictor:

- predicts a pixel with an estimate of the gradient: $T + L - TL$
- ... but without extrapolation, i.e. clamped to the $[\min(T, L, TL); \max(T, L, TL)]$ range

TL	T
L	

On the encoder side, simple to (auto)vectorize.

FJXL - RLE SIMD-fication

Very simple SIMD approach to run length encoding:

- if the current vector of values is identical to the last value from the previous vector, increment run count and skip producing output
- if not, emit a LZ77 copy length + distance symbol (if run length > 0) and encode the current vector raw

FJXL - Hybrid UInt Encoding

Given a number (written in base 2)

$$1b_{p-1}b_{p-2}\dots b_1b_0$$

we split it into a symbol that represents p (or a special symbol that represents 0) + p raw bits “ $b_{p-1}b_{p-2}\dots b_1b_0$ ”

This requires a fast-log2 operation; `__builtin_clz()` does the job, can be emulated for AVX2 vectors with some table lookups in a 16-entry table (`vpshufb`)

FJXL - Prefix coding

Usual prefix coding. Optimization: sample a few parts of the image (at random) to produce an image-adapted table.

We only have <16 symbols. We can SIMDfy prefix coding by doing table lookups with `vpshufb`.

We still need to concatenate all the Huffman and raw bits into a single bit stream. This could be done with 32 calls to a `PutBits` function, but SIMDfication provides a significant speed up: can be done with a sequence of bitwise operations to reduce to just 4 `PutBits`.

FPNGE

FPNGE - PNG vs JXL differences

- No division in tiles
- No color transforms
- One of 5 predictors (“filters”) is chosen per row
- Channels are interleaved, not split into planes
 - Makes RLE somewhat less effective
 - Requires masking to disable unused channels, or specialized code paths
- All operations are byte-wise, with wraparound
 - Can use 8-bit integers, i.e. 32 integers per vector
- No Hybrid Uint Encoding - Huffman raw alphabet has 256 symbols
- Two different checksums of image data (Adler32 and CRC32)

FPNGE - PNG filters

- Filter 0: noop
- Filter 1: subtract a
- Filter 2: subtract b
- Filter 3: subtract $(a+b)/2$
- Filter 4: subtract the Paeth predictor

```
p = a+b-c;  
pa = |p - a|;  
pb = |p - b|;  
pc = |p - c|;  
if pa <= pb && pa <= pc PAETH = a;  
else if pb <= pc PAETH = b;  
else PAETH = c;
```

c	b
a	

FPNGE - Paeth SIMDfication

Intermediate quantities for Paeth do **not** use wraparound: problematic for SIMD.

Alternative, equivalent formulation with just 8-bit intermediate quantities:

```
bc = b - c;  
ca = c - a;  
pa = c < b ? bc : -bc;  
pb = a < c ? ca : -ca;  
pc = (a < c) == (c < b) ? (bc >= ca ? bc - ca : ca - bc) : 255;  
PAETH = pa <= pb && pa <= pc ? a : pb <= pc ? b : c;
```

FPNGE - Filter choice

- libpng: filter that minimizes the sum of absolute values of residuals for the row
- fpnge: filter that minimizes the bit cost of the row (doing full mock-encodes)
- fpnge4: always use the Paeth predictor
- fpnge2: always use the Top predictor

FPNGE - Huffman coding

Fast Huffman coding in FJXL works because we have at most 16 raw symbols, which fit in a single vector pair.

For PNG, we can pick a custom Huffman table so that:

- Symbols [0, 16) (0000xxxx in binary) have their own Huffman code
- Symbols [240, 256) (1111xxxx in binary) have their own Huffman code
- All other symbols (yyyyxxxx in binary) have LUT[yyyy]xxxx as their Huffman code, i.e. lowest 4 bits are copied as-is.

The Huffman table can then fit in 3 vector pairs, which still allows fast lookups with `vpshufb + vpblendvb`.

Code

Code

FJXL: https://github.com/libjxl/libjxl/tree/main/experimental/fast_lossless

FPNGE: <https://github.com/veluca93/fpng>

Both are single-file encoders (+ main file).

FJXL has AVX2, ARM NEON and plain-C++ implementations.

FPNGE is AVX2 only.

Questions?