

JPEG2000 Performance Result (so far...)

A HiRISE-centric update by

Myche McAuley

2005-08-31

JPEG2000 Testing

- Characteristics tested for:
 - Lossless compression, does it really work?
 - Attainable compression ratios?
 - Affect of tile sizes on compression ratio?
 - Timing characteristics of (de)compression?
 - Not critical, but important as JPEG2000 format used for Web artifacts
 - Data type support?
 - Byte?
 - Short?
 - Float? (still TBD)
 - Large image file support (> 500MB)?

JPEG2000 Reference Software

- Used both Java and C reference implementations:
 - Java: JJ2000
 - <http://jj2000.epfl.ch/>
 - Caveat: compresses 8- and 16-bit data, but doesn't seem to support decompression to 16-bit
 - Didn't test extensively due to this restriction, but will investigate further
 - C: JasPer
 - <http://www.ece.uvic.ca/~mdadams/jasper/>
 - Compresses and decompresses 8- and 16-bit data
 - Both implementations currently require netpbm formatted files for input (PGM here)

JPEG2000 Testing (cont.)

- Test images must be somewhat indicative of expected MRO HiRISE RDRs:
 - Use MEX HRSC data.
 - Two HRSC images used:
 - h0068_0000_s22.img (~600MB)
 - h0068_0009_s22.img (~55MB)
 - 1MB == 1024*1024 bytes

JPEG2000 Test Images

- Additional stresses for the JPEG2000:
 - Pretend the HRSC header, binary header table and line prefixes are part of the image data.
 - Basically, compress the whole kaboodle and see what happens
 - Seeing how the algorithm performs in the face of this extra entropy should be interesting, especially since what we're after is lossless compression.

h0068_0000_s22.img

- 2618 x 119757
- Short data type
- Total size:
 - 627047673 bytes
- Histogram range:
 - With binary “noise”
 - Full 16-bit dynamic range
 - Image data only
 - [67,308], $\sigma = 52.5$



h0068_0009_s22.img

- 2618 x 11013
- Short data type
- Total size:
 - 57664088 bytes
- Histogram range:
 - With binary “noise”
 - Full 16-bit dynamic range
 - Image data only
 - [62,188], $\sigma = 9.2$



JPEG2000 Tests - computing environments

- Used two platforms
 - Linux
 - AMD64, 2.2GHz CPU with 1GB RAM
 - Mac OS X
 - Dual G5 Mac, 2.5GHz, with 1.5GB RAM
- Only crucial for timing performance, but interesting nonetheless

JPEG2000 Lossless Compression ratios

Using JasPer (C code) reference software

tilesize	h0068_0000_s22		h0068_0009_s22	
	16-bit	8-bit	16-bit	8-bit
1024	5.83	2.89	6.00	2.25
512	5.82	2.88	5.99	2.25
256	5.78	2.87	5.94	2.24
128	5.64	2.81	5.80	2.20
64	<u>5.16</u>	<u>2.59</u>	5.28	2.06
32	<u>4.02</u>	<u>2.06</u>	4.09	1.70

Underlined numbers indicate decompression failure for these tile sizes

JPEG2000 (de)compression processing times

(times in seconds)

Using JasPer (C code) reference software
times are (compression | decompression)

On Mac tilesize	h0068_0000_s22		h0068_0009_s22	
	16-bit	8-bit	16-bit	8-bit
1024	253 217	216 198	21 22	21 19
512	200 187	181 169	18 16	20 17
256	140 129	124 105	15 21	14 10
128	174 172	154 159	18 12	17 12
64	<u>258 ???</u>	<u>227 ???</u>	24 25	24 19
32	<u>488 ???</u>	<u>431 ???</u>	44 42	42 43

On Linux				
1024	238 193	193 169	22 18	20 16
512	210 230	152 142	14 17	14 14
256	157 247	107 109	10 15	11 10
128	162 313	101 127	11 15	10 10
64	<u>208 ???</u>	<u>117 ???</u>	13 21	12 13
32	<u>253 ???</u>	<u>187 ???</u>	18 29	18 27

Underlined numbers indicate decompression failure for these tile sizes

JPEG2000 Test Results

- Lossless compression achieved?
 - Yes, quite. Lossless compression ratios $> 2:1$ are impressive
- Tile size selection can have deleterious impact:
 - Smaller tiles better for error containment and web serving, but decrease compression ratio and decompression performance (insignificant issue). Pathological cases caused decompression to fail (not so insignificant).
 - Seems that a tile size of around 256 is pretty good
- Large file size support?
 - Yes, $> 500\text{MB}$. 2.3GB tests pending.

Pending Tests

- Very large file support
 - > 2GB (beyond 32-bit addressable range)
- Geometrically corrected and projected images (see next slide)
 - Large “blank” border areas
- Images with much larger dynamic range
 - Expectation that this will work, but just with lower compression ratios
- Floating point
 - Appears to be included in Part 10 of the spec.

Images still to be tested

- As HiRISE archive will only use JPEG2000 for RDRs, it would follow that these might be map projected (or otherwise geometrically corrected) giving an image that has blank boundaries
 - This is expected to compress very well.
 - Performance impacts TBD



Other References

- <http://www.oreillynet.com/lpt/a/4370>
- <http://aroundcny.com/technofile/texts/tec032104.html>
- http://www.ee.unsw.edu.au/~taubman/seminars_files/IEEE_IEA_J2K.pdf
- http://www.dsp.toronto.edu/~dsp/JPEG2000/JPEG2000_51to100.pdf