GEOSPATIAL IMAGE FORMATS SHOOTOUT – PART TWO

CHRIS HANSON CONTINUES WITH PART TWO OF HIS INDEPENDENT, IN-DEPTH INVESTIGATION OF IMAGE COMPRESSION FORMAT PERFORMANCE ISSUES

Welcome to part two of the Geospatial Image Formats Shootout. In part one (GEO:connexion International July/August 2007) we reviewed most of the common geospatial image formats and compared their merits. For lossless file formats, we benchmarked their final sizes, and for lossy formats we measured how much control you have over the resulting file size. Now we will place the front-runners under a microscope to examine what their defects look like. We will measure the magnitude of those defects with minute precision. And in the end we will pit them all against each other to see which is faster and under what conditions.

Test Data and Hardware

For these tests, we will revisit the same three datasets we used in Part one. To recap, we have one panchromatic image, one high-detail colour orthophoto and one lower-detail colour Landsat image. These images are all "small" on the scale of typical geospatial image databases, so to make the speed tests more true to real-world usage we will add a larger colour orthomosaic. CAPS-USGS-CS2002 is a 53Gb (758 GeoTIFF files) source dataset consisting of the 2002 vintage USGS Urban area hi-res orthophotos for Colorado Springs, Colorado (supplied courtesy of Colorado Aerial Photography Services). The larger data size will magnify speed differences making benchmarking more definitive.

According to their manufacturers, both programs are 32-bit applications and do not take any advantage of the new 64-bit CPU architectures from AMD and Intel. The manufacturers also indicate that neither program makes use of multiple CPUs or CPU cores.

We will be testing on two different systems, QUADZILLA and

GPStech07 Showcase

GPStech Showcase provides an opportunity for GIS professionals to see how high accuracy GPS solutions can improve GIS data quality and productivity.
Following our successful series of GPStech Seminars earlier in the year **KOREC** have partnered with **Trimble®**, **Ordnance Survey** and **ESRI**, providing access to leading industry specialists and the latest thinking in Mobile GIS.

Why Attend?

- Gain hands-on experience in the use of Trimble GPS with real-time demonstrations
- See the latest FastMap software releases from KOREC
- Presentations and demonstrations from ESRI on Mobile GIS Applications
- Technology updates on the latest Trimble GPS solutions from handheld GPS to RTK GPS using the OS Net™ service
- Learn how other organisations are using GPS solutions for GIS Applications
- Excellent networking opportunity

When & Where?

GPStech Showcase takes place in October, with venues in Edinburgh, York, Manchester, Bristol, London and Southampton.
Numbers are limited, so please register now free of charge to guarantee your place.

For more details and registration please visit:

www.korecgroup.com/gpstech

Event Partners

Trimble · Ordnance Survey Licensed Partner · ESRI UK VISIONARY THINKING

KOREC, Blundellsands House, 34 - 44 Mersey View, Waterloo, Liverpool, L22 6QB T: 0845 603 1214 W: www.korecgroup.com

Trimble Authorised Distribution Partner

KOREC™

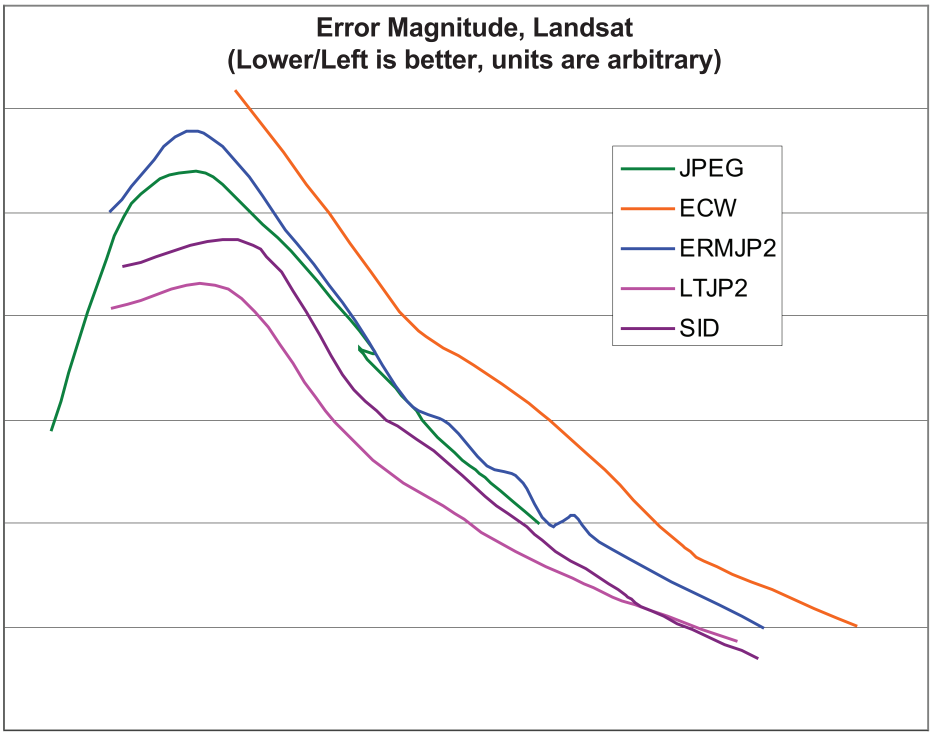

CHIPZILLA. QUADZILLA, purchased as a top-end computation system last year, is a dual dual-core AMD Opteron 64-bit 2GHz system with 8Gb of total RAM and a dual 320Gb SATA RAID0 disk subsystem. CHIPZILLA is a brand new quad-core Intel Core2 series model 6600 at 2.3GHz with 2Gb of RAM and a single 160Gb SATA disk. Both systems were built by JNCS.com specifically for high-performance computation. At purchase time last year, QUADZILLA cost $USD 4700 (and currently is valued at $2785). The newer CHIPZILLA currently costs $1399. These are representative of typical workstation-class machines you might find in a modern GIS department purchase order.

Appearance of Compression Artifacts

Our first test is not a quantitative benchmark, but rather a qualitative examination of the nature and visual appearance of the types of degradation each lossy compression method inflicts upon the data. Lossy means that the result isn't exactly the same as the original data – but how, precisely does it change? Here we will illustrate an exaggerated depiction of what the error looks like for each type. This will be a more drastic error than you will typically see, but will show you the "look" of the image degradation. There are two distinct types of degradation prevalent in the formats in question.

The older JPEG format is based on a Discrete Cosine Transform (DCT). DCT's artifacting manifests as "blocks" (typically 8 pixels by 8 pixels) of different pattern. As the compression rate goes up, the ability of the pattern within the block to approximate the original data is reduced. As well, each block tends to have a harder time matching up against its neighbours, leading to a mosaic-tile appearance. A simple summary of JPEG/DCT is that higher compression levels lead to a reduction of the original signal (the original intended pattern disappearing) as well as an increase in the level of noise (appearance of unintended patterns).

The newer Wavelet formats (ECW, MrSID, JP2) are all very similar in their degradation. Wavelets typically exhibit less introduction of undesired noise, instead offering a smoother decline in the amount of original signal. Noise is usually confined to "ringing" patterns adjacent to sharp features in the original image – smooth gradients are usually well represented.

To see the difference, refer to the large banner image for this article. Uncompressed is on the left, JPEG is in the middle and Wavelet is on the right. JPEG and Wavelet are represented at comparable compression ratios.

On paper, Wavelet's degradation would seem to be preferable; however some feel that this loss-of-texture is itself a worse side effect than JPEG's faux-texture-preserving increase in noise. Either way, lossy formats can't deliver you something for nothing. At higher compression ratios some valuable signal will be lost, and some undesired noise will be introduced. Which of the two is superior for a given purpose is really entirely application-dependent.

Magnitude of Degradation

For each of the lossy formats, we will now judge the quality versus size for each of the small datasets. We will not compare speed on these, nor will architecture differences be measured. We are simply determining which format looks best at a given output size. To judge this quality quantitatively, we will measure RMS error (Root Mean Square). This adds up all of the pixel differences in the image, weighting drastic errors more than minor errors (lower RMS values are better). This test will be run at several compression ratios, using AlphaPixel's PixelDelta, a test tool specially written for this purpose.

In these graphs, error magnitude is expressed as average-RMS-per-pixel multiplied by output file size in bytes. Lower numbers are better. This seemingly-complex formula will cancel out the unavoidable variances in the exact compression ratio and express exactly how good the image representation is at a comparable size. (Example: If doubling the file size results in exactly half as much error, the data point will graph the same, allowing for apples-to-apples measurement.) Without this,

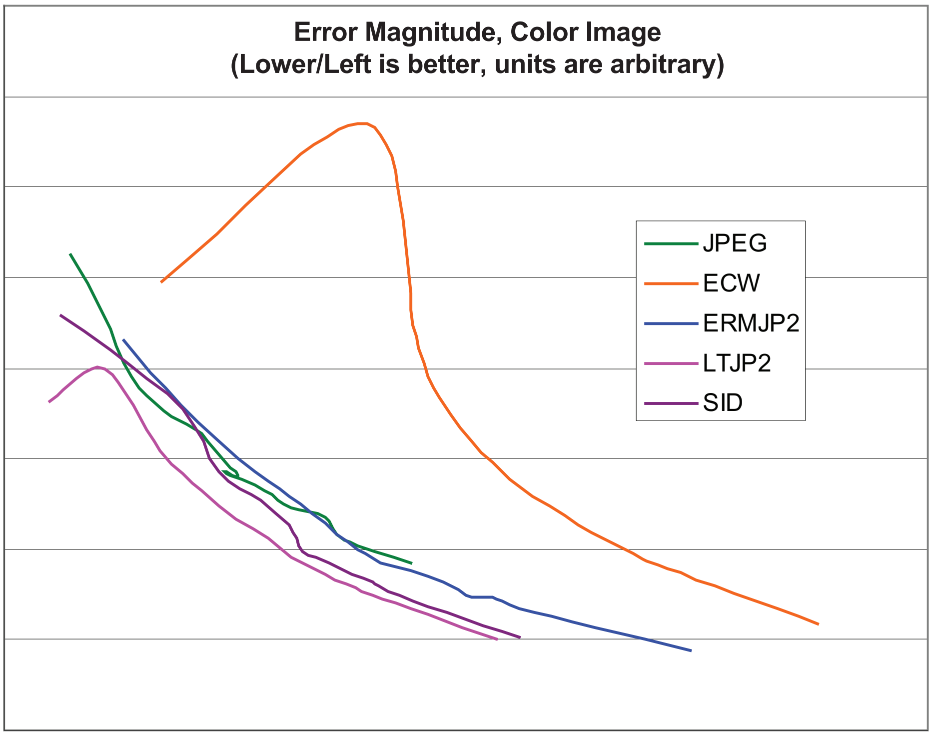

Whatever
Whenever
Wherever…

Your imagery. Delivered.

Express Server enables you to serve more users in less time using widely adopted standards and familiar tools—all without adding hardware or buying more bandwidth.

To learn how your company can benefit from instant online viewing of massive geospatial imagery with Express Server, visit:

www.lizardtech.com/products/exp
Or contact us at 206.652.5211 or geosales@lizardtech.com

Earth photography courtesy of NASA

ExpressServer

LIZARDTECH™
a celartem Company

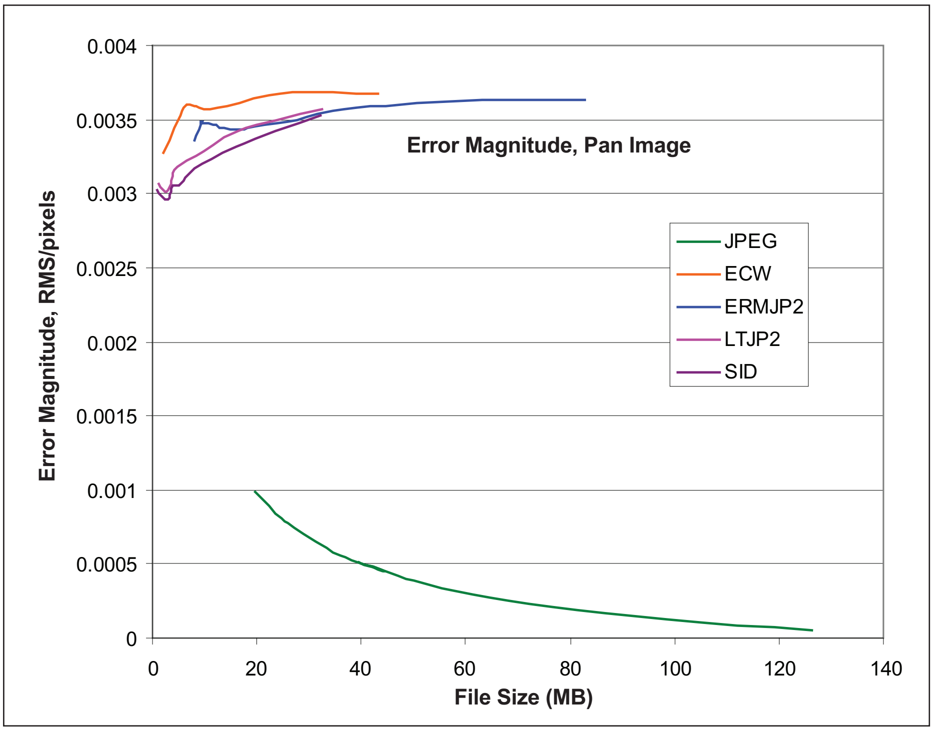

the results often clump too close together to be distinct. JPEG files were tested at quality levels between 1 and 12 (from Adobe Photoshop CS2), wavelet formats were tested at increments of 10 between 10:1 and 100:1, with extra data points at 5:1 and 200:1.

The Landsat image has uncontroversial results. LizardTech's JP2 file is the best representation, followed by MrSID, ER Mapper's JP2 and ECW. JPEG wanders crazily through the middle of the pack, usually between MrSID and ERMJP2 but with some strange excursions.

The Colour image confirms the same findings with only a few statistical oddities: LTJP2, MrSID, meandering JPEG, ERMJP2 and ECW bringing up the rear.

The results from the Pan image were so inscrutable that they are graphed here without the second multiplication factor in order to make sense of them. On the bottom X axis is resulting file size in bytes, the left Y axis represents average per-pixel RMS error. You can see that JPEG produced much larger files overall, but surprisingly in the range where the file sizes are comparable, JPEG did significantly better than the others! Visual inspection of the images confirms the quantitative result. This may be caused by the large amount of high-frequency clutter in this image.

Speed of Compression

Finally, we come to the large dataset speed shootout. We will endeavour to produce 10:1 compressed output from all four wavelet contenders, and see which candidate gets the desired result the fastest. Because the input data is in numerous (758) files, both programs will be required to potentially mosaic and balance the dataset before compressing. In reality, the images do not require any balancing. (Note: To compress large datasets like this to JP2, please refer to LizardTech's KnowledgeBase article F070602 for GeoExpress 6.1.4 and newer.)

To put that into real world units, here is how fast (in hours) our fastest system produces a roughly 10:1 compressed file:

ECW	ERMJP2	LTJP2	SID
1.68h	2.30h	2.49h	8.53h

In nearly mirrored-order of the quality standings, ECW leads the speed test followed by ERMJP2, and the two LizardTech encoders. So who is the overall winner? No clear leader can be crowned. ECW, while fast, produces the lowest-fidelity images. MrSID is the slowest by far. The two JP2 candidates are neck and neck, with LizardTech taking the lead in quality and ER Mapper pulling ahead in speed. JPEG surprisingly stays in the race, but wins no trophy due to its other deficiencies.

As expected with 32-bit programs, the 64-bit architecture did not improve performance, even when running a 64-bit OS. Between the Intel and AMD architectures, both CPUs are similar in performance, yet Quadzilla took 1.5-1.7 times longer. The bottleneck can be traced to the speed of the memory in the two systems. Quadzilla was built more for RAM quantity using cheaper DDR PC3200 200 MHz memory while Chipzilla sports top-speed DDR2 PC2-5300 333 MHz RAM, almost 1.6 times faster. This factor may greatly influence purchase configuration for new systems intended for image compression!

Note: Both LizardTech and ER Mapper suggest that there are options that can be used to tune performance of their products, especially on systems with large amount of memory. For our tests, we used the out-of-the-box settings (with the exception of the necessary F070602 configuration noted above).

While no one winner emerged from this shootout, it is clear that geospatial imagery users have excellent choices available to them in both products. The state of the art will continue to evolve, but the JP2 standard is obviously going to be the geospatial format of choice into the foreseeable future.

In closing, I'd like to thank 3D Nature, Frank Weed II, Colorado Aerial Photography Service, JNCS.com, and Microsoft for assisting with making this article happen. Additional resources for this article are available for download at AlphaPixel.com.

Chris 'Xenon' Hanson *is the co-founder of 3D Nature and AlphaPixel, innovators in the 3D Landscape Visualization and Remote Sensing Image Processing fields respectively. He has been working with computers fo 27 years, computer graphics for 20 years and GIS for 15 years. Contact Chris at Xenon@AlphaPixel.com.*