Efficient Transfer of Images over Networks

J.W. Percival

Space Astronomy Laboratory, University of Wisconsin, Madison, WI 53706

R.L. White

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218

Abstract. Effective remote observing requires sending large images over long distances. The usual approach to the transfer problem is to require high bandwidth transmission links, which are expensive to install and operate. An alternative approach is to use existing low-bandwidth connections, such as phone lines or the Internet, in a highly efficient manner by compressing the images. The combined use of existing low-cost infrastructure and standard networking software means that remote observing can be made practical even for small observatories with limited network resources.

We have implemented such a scheme based on the H-transform compression method developed by White (1992) for astronomical images, which are often resistant to compression because they are noisy. The H-transform can be used for either lossy or lossless compression, and compression factors of at least 10 can be achieved with no noticeable losses in the astrometric or photometric properties of the compressed images. The H-transform allows us to organize the information in an image so that the "useful" information can be sent first, followed by the noise, which makes up the bulk of the transmission. The receiver can invert a partially received set of H-coefficients, creating an image that improves with time. The H-transform is particularly well-suited to this style of incremental reconstruction, because the spatially localized nature of the basis functions of the H-transform prevents the appearance of artifacts such as ringing around point sources and edges.

Our implementation uses the WIYN Telescope Control System's TCP-based communications protocol. We sent an 800x800 16-bit astronomical image over a 2400 baud connection, which would normally take about 71 minutes; after only 60 *seconds*, the partially received H-transform produced an image that did not differ appreciably from the original. This poster will present a quantification of the efficiencies, as well as examples of images reconstructed from partial data.

1. Overview

We want to send images from the WIYN Observatory to Remote Observing sites for target acquisition, data quality monitoring, and quick-look analysis. We want to avoid expensive, private, high-speed data links, and we want to transmit the image in a way that allows the Remote Observer to see something quickly, with improvement as time goes on. This allows the Remote Observer to cancel the transmission as soon as the image is good enough.

We use a type of wavelet transform called the H-transform (White 1992) that organizes the image data for serial transmission.

We compare three methods:

- Sending the raw image in raster order (i.e. row by row)
- Sending the raw image in bitplane order (i.e. MSB first)
- Sending the H-transform in bitplane order (i.e. MSB first)

Sending images or H-transforms in bitplane order (MSB first) is good because the higher bitplanes are sparse and compress well. We use 1-dimensional or 2-dimensional quadtree encoding for each bitplane, with a final run-length encoding of the quadtree codes.

For a test image, we used an 800x800 pixel, 16-bit image of the Coma cluster digitized from the Palomar Sky Survey.

2. Transmission Software

We have developed a message protocol for the WIYN Observatory Control System that will be used for:

- Commands from the user interfaces
- Telemetry (engineering data) from the real-time system
- Data (images, spectra, catalogs, etc.)

This protocol uses TCP/IP sockets as a reliable byte stream, with a special message protocol for exchanging messages between applications.

We have written an image server application that sends images over the network in a machine-independent way, and an X-windows display client that allows the Remote Observer to select the transmission method and display style (linear, log, etc.), updates the display as more data are received, and allows the Remote Observer to cancel the transmission when the received image is good enough, or to wait until the received image has been reconstructed perfectly.

The H-transform can have negative pixels, so a special protocol must be used to avoid sending all the sign extension bits in a negative number. We store the sign bits, and transmit absolute values. After each bitplane is sent, a special "negation plane" is sent to negate any ultimately negative number whose absolute value has just had its most significant one-bit sent. Negating the H-transform coefficients "as you go" keeps the H-transform inversion well behaved on the remote end.

3. Comparison of Results

Figures 1, 2, 3, and 4 show the image after receiving the highest four non-zero bitplanes of the original raw image. After receiving one bitplane, the image is simply a bitmap with two intensities, black or white. After receiving two bitplanes, four intensities are shown, and so on. The full intensity resolution of the original is recovered only after receiving all the raw bitplanes.

Figures 5, 6, 7, and 8 show the image after receiving the highest seven non-zero bitplanes of the H-transform. (The highest seven bitplanes of the H-transform are compared to the highest four bitplanes of the raw image because these represent roughly comparable amounts of data.) Note that a variety of intensity variations is visible in even the earliest of the H-transform reconstructions.

Our quality measure Q is the "reduced χ^2":

$$Q = \frac{1}{\sigma_0^2 N} \sum (I_1 - I_0)^2 \tag{1}$$

where I_1 represents the partially received image, I_0 represents the original image, σ_0 is the standard deviation of the noise in the original image, and the sum is over all N pixels in the image. $Q \approx 1$ occurs when the images are globally consistent to within the noise (although the differences may be highly spatially correlated especially when the image is sent in raster order).

Figure 9 shows Q as a function of the number of transmitted 4-byte words (including protocol overhead). The raster mode converges most slowly, as expected. $Q < 1$ occurs after 326000 words (98.7%) have been transmitted. Sending the image in bitplane order causes convergence after 46861 words (20.1%). Sending the H-transform in bitplane order causes convergence in only 4564 words (1.59%).

Figures 10, 11, and 12 show the intensity histograms of the original image, and the images in Figure 4 and Figure 8 respectively. Figure 4 shows good "visual" convergence after receiving plane 10 of the raw image, but Figure 11 shows that the data are still heavily quantized. Figure 12 shows a much less quantized intensity histogram after receiving plane 10 of the H-transform, more closely approximating that of the original image.

References

White, R.L. 1992, High Performance Compression of Astronomical Images, available via ftp from stsci.edu

Figure 1. Raw image, bitplane order, planes 13-15. 0.313% of the data have been received. Plane 13 is the highest non-zero plane, so the image at this point is only 1-bit deep.

Figure 2. Raw image, bitplane order, planes 12-15. 1.36% of the data have been received. Two non-zero planes have been received, giving 4 intensity levels.

Figure 3. Raw image, bitplane order, planes 11-15. 2.81% of the data have been received. Three non-zero planes have been received, giving 8 intensity levels.

Figure 4. Raw image, bitplane order, planes 10-15. 11.3% of the data have been received. Four non-zero planes have been received, giving 16 intensity levels.

Figure 5. H-transform, bitplane order, planes 13-31. 0.746% of the data have been received. Note the various spatial and intensity scales that are already present in the image.

Figure 6. H-transform, bitplane order, planes 12-31. 1.59% of the data have been received.

Figure 7. H-transform, bitplane order, planes 11-31. 3.07% of the data have been received.

Figure 8. H-transform, bitplane order, planes 10-31. 7.9% of the data have been received.

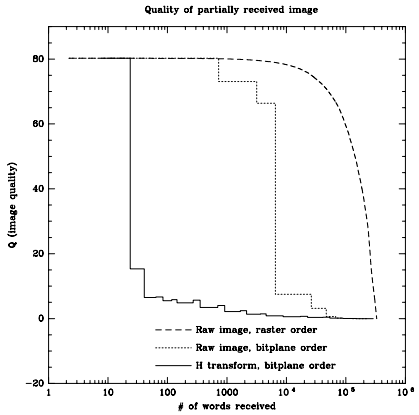


Figure 9. Quantitative improvement: reduced χ^2 measure as a function of the number of transmitted 4-bytes words (including protocol overhead). Note that the H-transform, bitplane method reaches the asymptotic part of its quality curve when fewer than 400 *bytes* of data have been received!

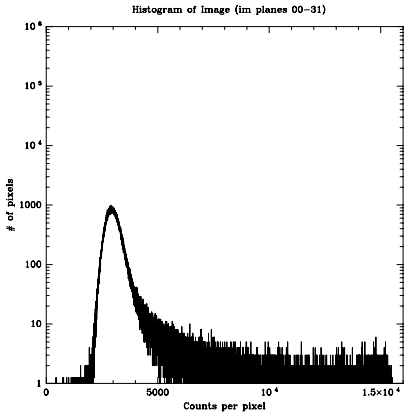


Figure 10. Intensity histogram of original image.

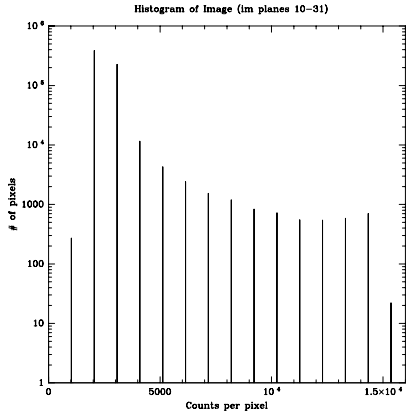


Figure 11. Intensity histogram of image after planes 10-15 of the raw image have been received. Moving through the raw image in bitplane order quantizes the received image according to the number of received bitplanes. The quantization shown here makes it impossible to determine the sky brightness accurately.

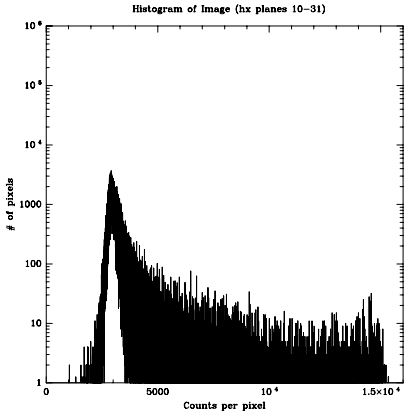


Figure 12. Intensity histogram of image after planes 10-31 of the H-transform have been received. Sending the wavelet coefficients in bitplane order greatly reduces intensity quantization in the received image. Note the improved accuracy possible in estimating the sky brightness.