Subband Encoding of High Dynamic Range Imagery

Greg Ward, Sunnybrook Technologies
Maryann Simmons, Walt Disney Feature Animation

1. Summary

All existing and proposed high dynamic range (HDR) formats yield prohibitively large images that can only be viewed and manipulated by specialized software, which has impeded wider adoption of this technology. In this sketch, we present a simple approach to HDR storage in a backwards-compatible, "lossy" image format. A tone-mapped version of the HDR original is accompanied by restorative information carried in a subband of a standard 24-bit RGB format. This subband contains a compressed *ratio image*, which when multiplied by the tone-mapped foreground, recovers the HDR original. The tone-mapped image data may also be compressed, permitting the composite to be delivered in a standard JPEG wrapper. To naïve software, the image looks like any other, and displays as a tone-mapped version of the original. To HDR-enabled software, the foreground image is merely a tone-mapping suggestion, as the original pixel data are available by decoding the information in the subband. Errors are visible in only a very small percentage of the pixels after decoding, and the technique requires only a modest amount of additional space for the subband data, independent of image size.

2. Method

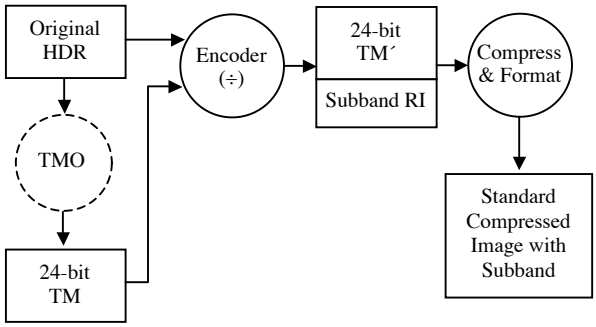

Figure 1. High level HDR subband encoding pipeline.

Figure 1 shows an overview of our HDR encoding pipeline. We start with two versions of our image: a scene referred HDR image, and an output referred tone-mapped version. The generation of the tone-mapped version is a separable problem, and our technique is designed to work with multiple operators. The encoding stage takes these two inputs and produces a composite, consisting of a potentially modified version of the original tone-mapped image, and a subband ratio image that contains enough information to reproduce a close facsimile of the HDR original. The next stage compresses this information, offering the tone-mapped version as the JPEG base image, and storing the subband as metadata in a standard JFIF wrapper.

Figure 2 shows the two possible decoding paths. The high road decompresses both the tone-mapped foreground image and the subband, delivering them to the decoder to recover the HDR pixels. Naïve applications follow the low road, ignoring the subband in the metadata and reproducing only the tone-mapped foreground image.

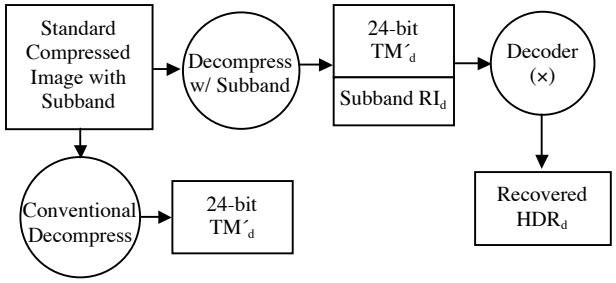

Figure 2. Alternate decoding paths for compressed composite.

In the simplest incarnation of our method, the encoder divides the HDR pixels by the tone-mapped luminances, producing an 8-bit, log-encoded ratio image that is stored as metadata in a standard JPEG image. Decoding follows decompression with a multiplication step to recover the original HDR pixels. Unfortunately, the simplest approach fails for pixels that are mapped to black or white by the tone-mapping operator, and the limited size of metadata in JPEG makes subband compression a challenge. These are two important issues we address with our technique. A third issue is wide-gamut color, and for this we employ a desaturation scheme that reduces tone-mapping induced color shifts while preserving the entire visible gamut.

We demonstrate our encoding on a series of HDR synthetic and natural images. Compression ratios vary, but typical images squeeze down to 3 bits/pixel with visible differences in fewer than 0.2% of the pixels by standard perceptual measures (Daly's VDP).

Figure 3. Memorial Church tone-mapped with Reinhard's global operator next to ratio image needed for HDR recovery. (Image courtesy Paul Debevec.)

3. References

DEBEVEC, P., and MALIK, J. 1997. Recovering High Dynamic Range Radiance Maps from Photographs. In *Proceedings of SIGGRAPH 1997*.

WARD LARSON, G. 1998. Overcoming Gamut and Dynamic Range Limitations in Digital Images. In *Proc. of IS&T/SID 6th Color Imaging Conference*.