Design and Evaluation of Backward Compatible High Dynamic Range Video Compression

Rafał Mantiuk, Alexander Efremov, Karol Myszkowski, Hans-Peter Seidel

MPI–I–2006–4–001 April 2006

Authors' Addresses

Rafał Mantiuk
Max-Planck-Institut für Informatik
Stuhlsatzenhausweg 85
66123 Saarbrücken
`mantiuk@mpi-inf.mpg.de`

Alexander Efremov
Max-Planck-Institut für Informatik
Stuhlsatzenhausweg 85
66123 Saarbrücken
`aefremov@mpi-inf.mpg.de`

Karol Myszkowski
Max-Planck-Institut für Informatik
Stuhlsatzenhausweg 85
66123 Saarbrücken
`karol@mpi-inf.mpg.de`

Hans-Peter Seidel
Max-Planck-Institut für Informatik
Stuhlsatzenhausweg 85
66123 Saarbrücken
`hpseidel@mpi-inf.mpg.de`

Abstract

In this report we describe the details of the backward compatible high dynamic range (HDR) video compression algorithm. The algorithm is designed to facilitate a smooth transition from standard low dynamic range (LDR) video to high fidelity high dynamic range content. The HDR and the corresponding LDR video frames are decorrelated and then compressed into a single MPEG stream, which can be played on both existing DVD players and HDR-enabled devices.

Keywords

video compression, high dynamic range, backward compatibility, MPEG, tone mapping, visual perception, prefiltering, visual masking, visible differences predictor

Contents

1 Introduction

In this report we describe the details of the backward compatible high dynamic range (HDR) video compression algorithm. The algorithm is designed to facilitate a smooth transition from standard low dynamic range (LDR) video to high fidelity high dynamic range content. The HDR and the corresponding LDR video frames are decorrelated and then compressed into a single MPEG stream, which can be played on both existing DVD players and HDR-enabled devices.

This report contains a more detailed description of the algorithm that was proposed in [80]. The report also includes more results, an extended section on implementation details and a comparison with the wavelet compander (Section 6.2). To present the entire description of the encoding system in a single document, this report contains selected sections of the original paper [80].

The encoding algorithm is described in Section 2, and the decoding algorithm in Section 3. We give some details on our implementation in Section 4. Several aspects of the algorithm are tested and its performance is compared with other methods in Section 5. Finally, we compare our method with JPEG HDR and with a wavelet compander in Sections 6.1 and 6.2.

2 Video Encoding Algorithm

An overview of the backward-compatible HDR MPEG encoding is shown in Figure 2.1. Input to our encoder is two sequences of video frames: the LDR frames in the sRGB color space (3 bytes per pixel) and the HDR frames in the absolute CIE XYZ (2° standard observer) color space (three floating point numbers per pixel). The LDR frames usually contain a tone mapped or gamut mapped version of the original HDR sequence, which is intended for LDR devices. The encoder produces three compressed streams: a *LDR stream*, which is fully compatible with MPEG; a *residual stream*, which contains differences between the LDR and the HDR frames; and an *auxiliary stream*, which contains auxiliary data needed for the reconstruction of the HDR frames. In this section we describe each step of the video encoding.

2.1 Step 1: Encode LDR frames

The best performance of our video encoder can be achieved when a HDR stream does not duplicate the information encoded in a LDR stream, i.e. both streams are decorrelated. Most of the processing described in this section is intended to achieve this goal.

For encoding of LDR frames we use an MPEG video encoder compatible with the ISO/IEC 14496-2 standard, but other video encoders can be used as well. The resulting video stream is encapsulated in a media container format, such as Audio Video Interleave (AVI) (used in our implementation) or QuickTime, in such a way that it can be recognized and played back by existing software. This ensures backward compatibility of our video files.

It is important to notice that MPEG is a lossy compression method and therefore the LDR frames that are decoded are not exactly the same as the original LDR frames but contain compression artifacts. This is why we have to use the LDR decoded frames as a reference for the latter steps of the HDR encoding (see *MPEG decode* block in Figure 2.1). Such decoding step

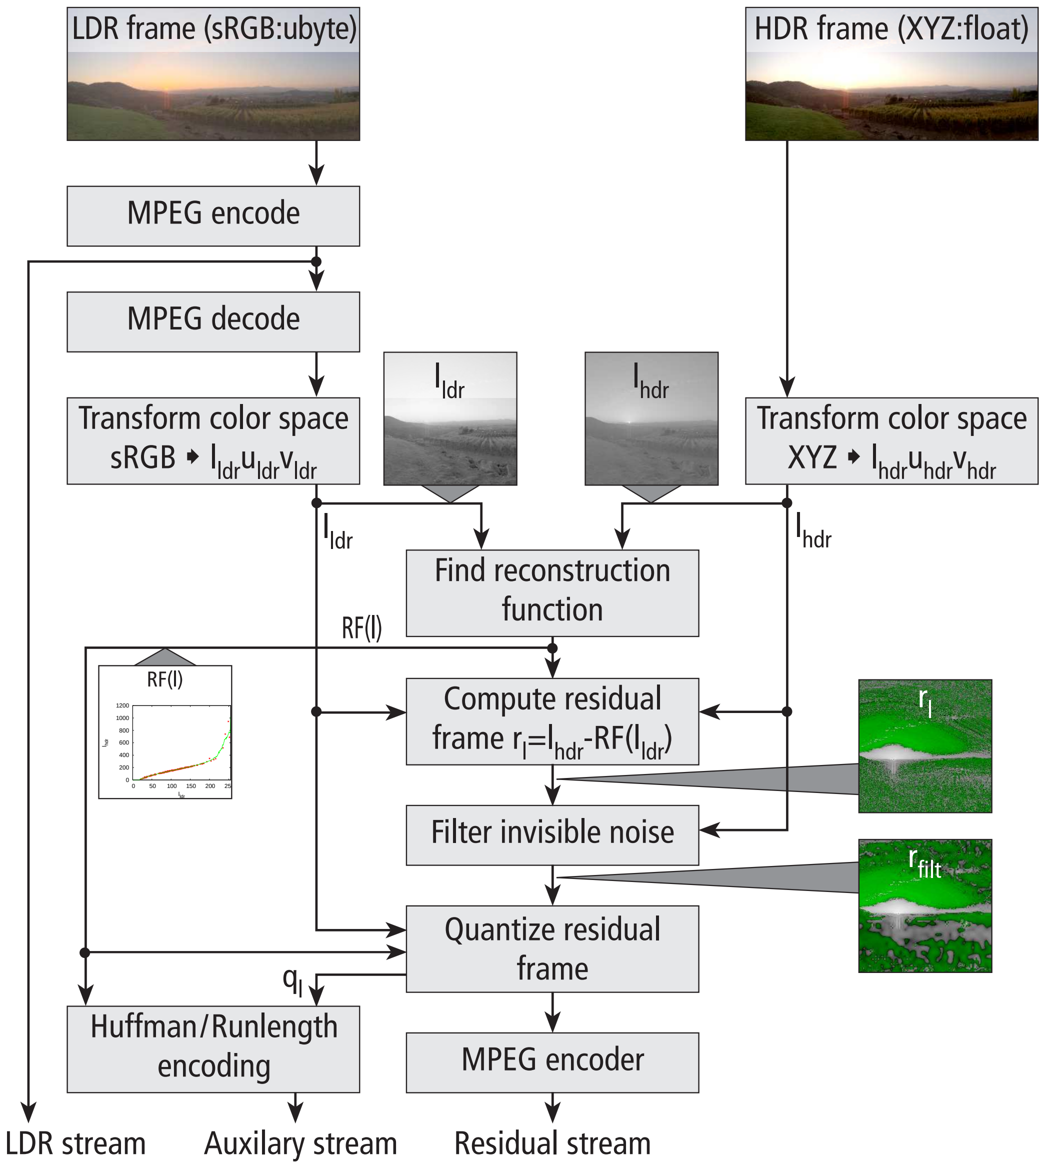

Figure 2.1: A data flow of the backward-compatible HDR MPEG encoding. See text for details.

is in fact done at no costs since the MPEG encoders have to decode frames internally to estimate motion vectors.

2.2 Step 2: Transform color spaces

Before we can make any comparison between LDR and HDR frames to assess their correlation, we need to convert them to two color spaces that would be both "compatible" and perceptually uniform. "compatible" means that color channels of both LDR and HDR color space represent approximately the same information. Perceptual uniformity is needed to estimate color differences according to perceivable, rather than arithmetic, differences.

Although an obvious choice for LDR frames would be a variant of YC_rC_b, which is a common color space for image and video compression, we reject it because of its limited color gamut. The HDR frames should preserve the full visible color gamut, even though it cannot be displayed on the existing displays. We have experimented with several color spaces, including a variant of RGB with an extended gamut (more saturated primaries), but finally we achieved the best results with a combination of the CIE 1976 Uniform Chromacity Scales (u', v') and the nonlinearity used in the *sRGB* color space.

We convert LDR pixels from the *sRGB* color space to the $l_{ldr}u_{ldr}v_{ldr}$ space. First, we compute relative CIE XYZ color coordinates using the sRGB conversion formulas [56]:

$$\begin{aligned} R' &= R_{8bit}/255 \\ R &= \begin{cases} R'/12.92; & \text{if } R' \leq 0.04045 \\ ((R'+0.055)/1.055)^{2.4} & \text{otherwise} \end{cases} \end{aligned} \tag{2.1}$$

for each color channel, R_{8bit}, G_{8bit}, and B_{8bit}, and:

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} 0.4124 & 0.3576 & 0.1805 \\ 0.2126 & 0.7152 & 0.0722 \\ 0.0193 & 0.1192 & 0.9505 \end{bmatrix} \cdot \begin{bmatrix} R \\ G \\ B \end{bmatrix} \tag{2.2}$$

The matrix in the equation above assumes the white point D65. Next, we compute luma[1] for LDR pixels using the sRGB gamma corrected color values:

$$l_{ldr} = 0.2126 \cdot R_{8bit} + 0.7152 \cdot G_{8bit} + 0.0722 \cdot G_{8bit} \tag{2.3}$$

Similarly as in [151], we represent the values for the LDR color channels as u' and v' chromacities as recommended by the CIE 1976 Uniform Chromacity

[1]Luma is a new word proposed by the NTSC in 1953 to prevent confusion between the Y' component of a color signal and the traditional meaning of luminance. While luminance is the weighted sum of the linear RGB components of a color video signal, proportional to intensity, luma is the weighted sum of the non-linear $R'G'B'$ components after gamma correction has been applied, and thus is not the same as either intensity or luminance. Source: http://www.wikipedia.org/

Scales (UCS) diagram and defined by the equations:

$$\begin{aligned} u' &= \tfrac{4X}{X+15Y+3Z} \\ v' &= \tfrac{9Y}{X+15Y+3Z} \end{aligned} \tag{2.4}$$

Then we encode u' and v' using 8-bits:

$$\begin{aligned} u_{ldr} &= u' \cdot 410 \\ v_{ldr} &= v' \cdot 410 \end{aligned} \tag{2.5}$$

Note that we use u' and v' chromatices rather than u^* and v^* of the $L^*u^*v^*$ color space. Although u^* and v^* give better perceptual uniformity and predict the loss of color sensitivity at low light, they are strongly correlated with luminance. Such correlation is highly undesired in image or video compression. Besides, u^* and v^* could reach high values for high luminance, which would be difficult to encode using an 8-bit MPEG encoder.

After converting LDR pixel values, we transform HDR pixels to a perceptually uniform space $l_{hdr}u_{hdr}v_{hdr}$, that would be comparable with the $l_{ldr}u_{ldr}v_{ldr}$ space introduced above. To represent color, we use the same formulas as for LDR pixels, given in Equations 2.4 and 2.5. More difficult is the case of HDR luminance since an ordinary gamma correction cannot be used for a luminance values ranging from 10^{-5} to 10^{10} cd/m^2, which can be found in real world scenes. Instead, we use a perceptually uniform luminance representation that has been derived from the contrast detection measurements for human observers. This space was shown to have similar properties as a gamma correction for LDR, but can encode the full visible range of luminance using 11-12 bits and ensures that quantization errors are below the threshold of visibility. A similar encoding was used in the context of HDR video encoding [81]. The details and the derivation of this space can be found in [84]. To convert HDR luminance, y, into 12-bit HDR luma, l_{hdr}, we use the formula:

$$l_{hdr}(y) = \begin{cases} a \cdot y & \text{if } y < y_l \\ b \cdot y^c + d & \text{if } y_l \leq y < y_h \\ e \cdot \log(y) + f & \text{if } y \geq y_h \end{cases} \tag{2.6}$$

and for the inverse conversion, from 12-bit luma to luminance, we apply:

$$y(l_{hdr}) = \begin{cases} a' \cdot l_{hdr} & \text{if } l_{hdr} < l_l \\ b'(l_{hdr} + d')^{c'} & \text{if } l_l \leq l_{hdr} < l_h \\ e' \cdot \exp(f' \cdot l_{hdr}) & \text{if } l_{hdr} \geq l_h \end{cases} \tag{2.7}$$

The constants are given in the table below:

a	b	c	d	e	f	y_l	y_h
17.554	826.81	0.10013	-884.17	209.16	-731.28	5.6046	10469
a'	b'	c'	d'	e'	f'	l_l	l_h
0.056968	7.3014e-30	9.9872	884.17	32.994	0.0047811	98.381	1204.7

2.3 Step 3: Find reconstruction function

Both LDR and HDR frames contain similar information and are therefore strongly correlated. This is illustrated in Figure 2.2, which shows how the luma values of a LDR frame relate to the luma values of a HDR frame. The relation is different for each tone mapping algorithm, but in general it follows an approximately linear function with more variance at high values. Uncorrelated pixels at the right end of the l_{ldr} axis are the result of luminance clamping that is applied in many tone mapping algorithms. Local tone mapping usually results in higher variance and therefore a more "noisy" shape of this relation, while global tone mapping results in a direct one-to-one relationship unless some pixel values are clamped.

The goal of most compression methods is to decorrelate data, so that the same information is not encoded twice. To decorrelate LDR and HDR frames, we find a *reconstruction function*, which predicts the value of a HDR pixel based on the value of the corresponding LDR pixel. Having such a function we need only to encode the differences between values predicted by the reconstruction function and the actual values from a HDR frame. Such differences are usually close to zero and therefore can be efficiently encoded in *residual frames*. The reconstruction function needs to be defined only for 256 values (bins) for 8-bit per channel LDR encoding. The function does not need to be continuous since its major role is to make the values of the residual frame as small as possible. Some examples of reconstruction functions for different tone mapping algorithms are plotted in Figure 2.2 as continuous green lines.

A mapping from LDR values to HDR values is, in the general case, a one-to-many relationship – there are many HDR pixels values that fall in one of 256 bins of the reconstruction function (LDR pixel values). The questions is how to find a value for each bin that would lead to the best compression performance. We experimented with an arithmetic mean, a median and a midrange [2]. While the midrange gave the worst compression ratio, the arithmetic mean and the median exhibited similar performance. We have

[2] Midrange is defined as an arithmetic mean of the maximum and minimum value in a set.

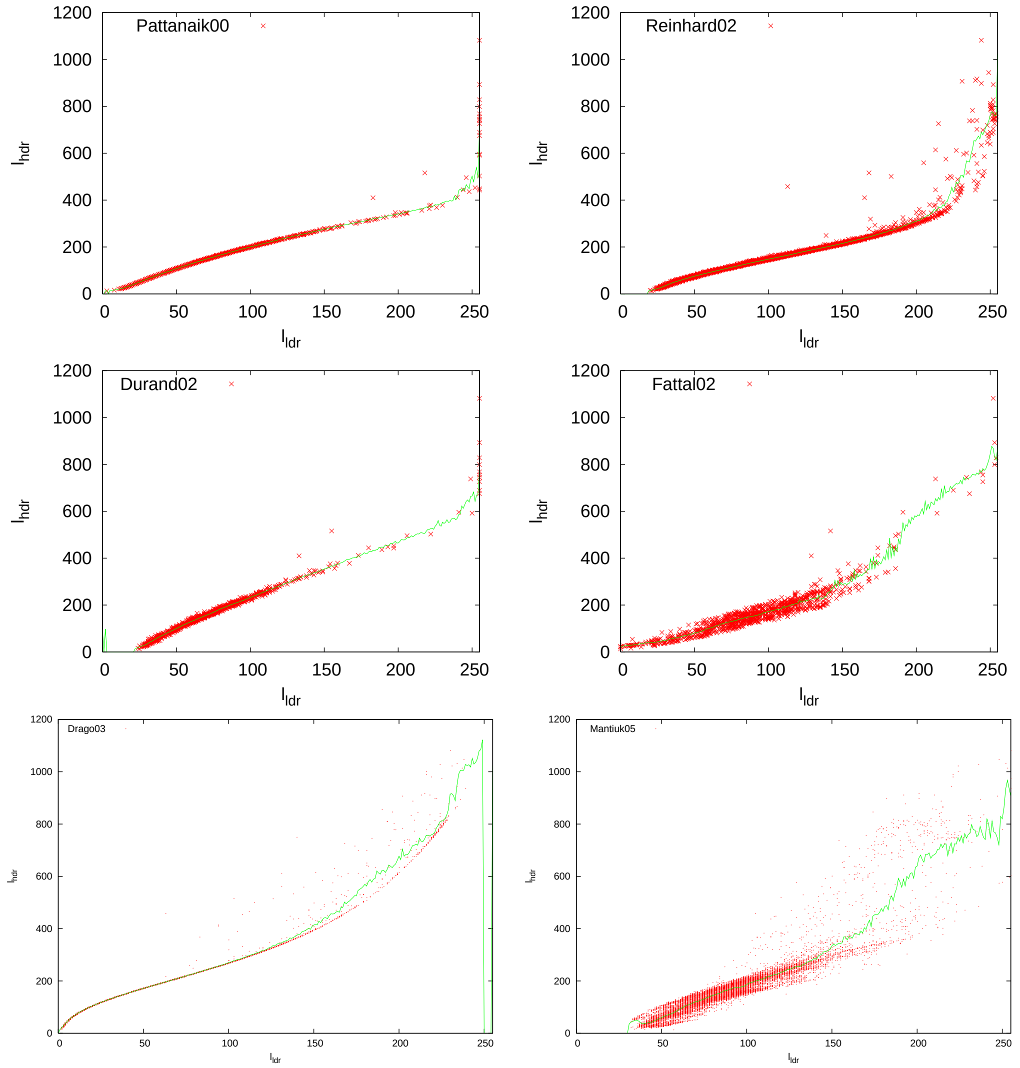

Figure 2.2: The relation between LDR (l_{ldr}) and HDR (l_{hdr}) luma values for various tone mapping algorithms (marked in red) and the corresponding reconstruction functions (marked in green). Tone mapping algorithms (left to right, top to bottom): [98], [109], [36], [39], [33] and [83, Contrast Equalization]. The relations are plotted for the *Memorial Church* image.

decided to use an arithmetic mean because of its lower computational cost.

To summarize, we define the reconstruction function as the arithmetic mean of all pixels falling in a corresponding bin Ω_l:

$$RF(l) = \frac{1}{Card(\Omega_l)} \sum_{i \in \Omega_l} l_{hdr}(i) \text{ where } \Omega_l = \{i = 1..N : \; l_{ldr}(i) = l\} \qquad (2.8)$$

$l = 0..255$ is an index of a bin, N is the number of pixels in a frame, $l_{hdr}(i)$ and $l_{hdr}(i)$ are luma values of the i-th LDR and HDR pixel respectively.

We executed a set of tests on video sequences to decide how often a reconstruction function should be updated: each frame, each intra-encoded frame (I-frame), or if the update should depend on a difference between consecutive frames. We achieved the best compression ratio when the reconstruction function was updated each frame, while updating it each I-frame resulted in severe artifacts.

The relation between LDR and HDR frames is complex only for luminance, and color channels can be quite accurately predicted with simple relations $u_{hdr}(i) = u_{ldr}(i)$ and $v_{hdr}(i) = v_{ldr}(i)$. Although this may not be true for some sophisticated gamut mapping cases, we did not find it necessary to compute a reconstruction function for chroma channels in the case of all the tone mapping operators we tested.

Since the reconstruction function tends to be slowly changing with an increasing slope, we apply an adaptive Huffman algorithm on the differences between the values in consecutive bins to significantly reduce the size of the stored data. The size of the auxiliary data stream, which stores a reconstruction function, is below 1% of the total stream size, therefore its storage overhead is almost insignificant.

We briefly compare our approach with the JPEG HDR compression [148]. A more detailed comparison will be given in Section 6.1. The JPEG HDR compression encodes a ratio between HDR and LDR luminance values, rather than a difference between HDR values and the reconstruction function. However, it can be easily shown that such a ratio is meant to achieve the same goal, which is to decorrelate HDR and LDR pixels. Since a ratio of HDR and LDR luminance corresponds to a difference in the logarithmic domain, and our luminance to luma mapping from Equation 2.6 has roughly logarithmic properties, the ratio encoding of JPEG HDR corresponds to a linear reconstruction function $l_{hdr} = a \cdot l_{ldr}$. As we experimented with such simple reconstruction functions, we found that they give inferior results compared to better fitted ones, like those computed from Equation 2.8. In Figure 2.3 we plot the reconstruction function used in JPEG HDR. Obviously, it does not follow the data well and some luma information is therefore encoded twice in an LDR and a HDR (subband) stream, which leads to worse compression performance.

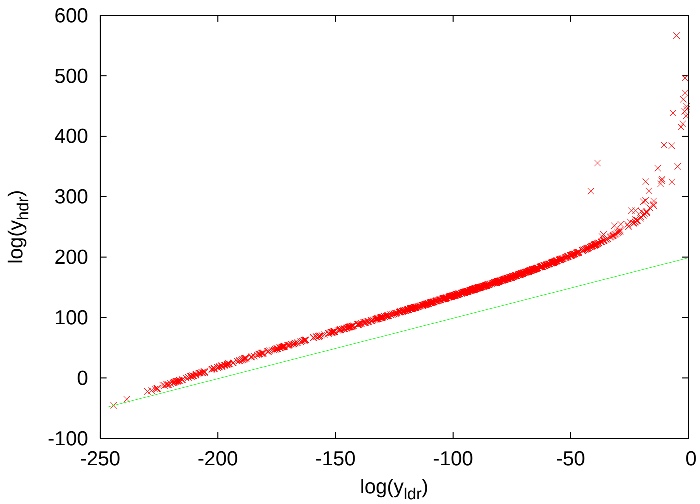

Figure 2.3: A potential reconstruction function for the approach employed in JPEG HDR [148] compression (marked in green) and a relation between LDR and HDR pixel values (marked in red). The ratio used in JPEG HDR is equivalent to a linear reconstruction function in the logarithmic domain. Such a function does not decorrelate HDR and LDR luma well and therefore reduces compression savings.

2.4 Step 4: Compute residual frame

A residual frame is a difference between HDR pixel values and the values predicted by a reconstruction function from LDR pixel values. Each pixel in a residual frame is therefore computed as:

$$r_l(i) = l_{hdr}(i) - RF(l_{ldr}(i)) \tag{2.9}$$

where RF is the reconstruction function from Equation 2.8. A residual frame for chroma is computed in a similar way, but the reconstruction function is assumed to be an identity function:

$$r_u(i) = u_{hdr}(i) - u_{ldr}(i) \tag{2.10}$$

The same formula is used for v chroma channel.

2.5 Step 5: Filter invisible noise

A properly chosen reconstruction function can reduce the amount of data that encodes HDR frames significantly. But despite this huge saving, the residual frames still contain a significant amount of invisible noise, which worsens compression ratio. We employ a model of the Human Visual System to remove invisible noise from the residual frames and thus to improve efficiency of the compression. We give details on this filtering in Section 2.8.

2.6 Step 6: Quantize residual frame

Although the magnitudes of the differences encoded in residual frames are usually small, they can theoretically take values from −4095 to 4095 (for 12-bit HDR luma encoding). Such values cannot be encoded using 8-bit MPEG encoder. Although MPEG standards provide an extension for encoding luma values on 12 bits, such an extension is rarely implemented, especially in hardware. Instead, we would like to reduce the magnitude of residual values so that they can be encoded using a standard 8-bit MPEG encoder.

We have experimented with a non-linear quantization, where large absolute values of residuals were heavily quantized, while small values were preserved with maximum accuracy. Since very few pixels contain a large magnitude of residual, most pixels are not affected by the strong quantization. Such a solution, although giving the best SNR, resulted in poor visual quality for some images. This was because the very few pixels that were heavily quantized attracted attention due to large quantization errors. Therefore the final judgement of quality was mostly based on those few distorted pixels.

A simple clamping of residual values to 8-bit range produced visually better results, but at the cost of losing some details in bright or dark regions. Additionally, to reduce clamping at the cost of a stronger quantization, the residual values can be divided by a constant quantization factor. Such a factor would decide on the trade-off between errors due to clamping and errors due to quantization. Furthermore, we observed that very few bins of a reconstruction function contain residual values that exceed 8-bit range. Therefore the quantization factor can be set separately for each bin, based on the maximum magnitude of the residual that belongs to that bin. Therefore, the residual values after quantization can be computed as:

$$\hat{r_l}(i) = [r_l(i)/q(m)]^{-127 \div 127}, \text{ where } m = k \Leftrightarrow i \subset \Omega_k \tag{2.11}$$

and quantization factor, $q(m)$, is selected separately for each bin Ω_k:

$$q(m) = \max(q_{min}, \frac{\max_{i \in \Omega_l}(|r_l(i)|)}{127}) \tag{2.12}$$

q_{min} is a minimum quantization factor, which is usually set to 1 or 2. $[\cdot]^{-127 \div 127}$ is an operator that rounds the values to the closest integer and then clamps

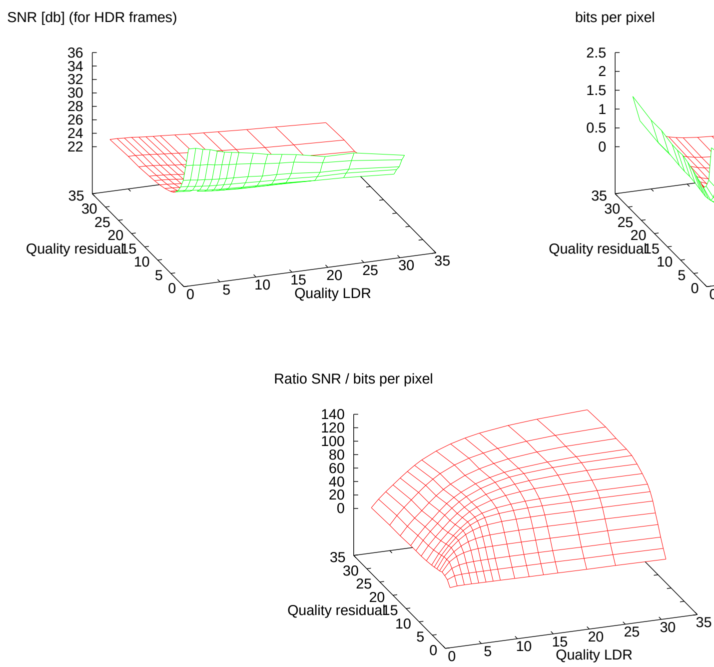

Figure 2.4: Quality of HDR frames and total stream size for possible combinations of quality settings for HDR and residual stream. Quality LDR and residual is given as a quantization scale factor (*qscale*).

them if they are smaller than -127 or larger than 127. The l subscript in r_l denotes a luma channel.

The quantization factors $q(m)$, where $m = 0..255$, need to be stored in an MPEG stream to later restore non-quantized residual values on the decoding stage. We store quantization factors together with the reconstruction function in the auxiliary data stream. Since quantization factors are usually equal q_{min} except a few bins, we found that a run-length encoding followed by the Huffman encoding can effectively reduce the size of this data.

2.7 Step 8: MPEG encode

When the residual values lie within an 8-bit range, they can be encoded using an ordinary MPEG compression. We encode the quantized residual values, $\hat{r_l}$, and chroma residual values r_u and r_v after rounding them to the nearest integer value.

As we encode two video streams: a backward compatible LDR stream and a residual stream, we also have to set quality settings for each stream separately. This is neither intuitive nor convenient and an end user should be provided with a single quality control. Furthermore, some quality settings for LDR and residual streams may give better compression performance than the others. To find the best strategy for controlling quality settings, we executed tests for all combinations of the settings and measured the quality of the resulting video as the Signal to Noise Ratio (SNR) (more on quality metrics in Section 5), the total size of the stream as a number of bits per pixel, and the ratio of both. The results are summarized in Figure 2.4. The maximum quality for HDR frames can be achieved only if a residual stream is encoded with the best quality. The quality settings of a LDR stream also affects the quality of HDR frames but has lower impact. The plot showing a ratio of SNR and bits per pixel suggests that a good set of quality settings lie on the line *Quality residual* = *Quality LDR*. Therefore for the rest of the tests we always set the quality of both streams to the same value.

Note that the operations we have applied so far to get the residual values are approximately linear, assuming that the reconstruction function is close to a linear function and that the effect of the adaptive quantization from Equation 2.11 is minimal. This way the visual information of a residual frame stays in the same frequency bands as the original HDR frame, and the DCT quantization of the residual has a similar effect as for the original HDR pixel values. Therefore, a standard DCT quantization matrix can be used for the encoding residual frames.

Since both LDR and HDR frames contain the same scenes, the optical flow should be also the same for both of them. This should allow as to reuse motion vectors computed for LDR frames for residual frames. Those vectors also do not need to be stored twice. However, in our implementation we compute motion vectors separately for LDR and residual frames.

2.8 Filtering of Invisible Noise

Residual frames do not compress well mainly because they contain a large amount of high frequencies. These high frequencies come from three sources: noise in the source HDR images, rounding errors from the tone mapping algorithm, and the DCT quantization errors due to MPEG encoding of LDR frames (refer to Figure 2.1). However, much of this high frequency information does not need to be preserved in the residual stream since it is not visible to the human eye. To remove such invisible noise and thus improve compression efficiency, we introduce a filtering algorithm based on a simplified model of the human visual system (HVS). Although models of the HVS have been

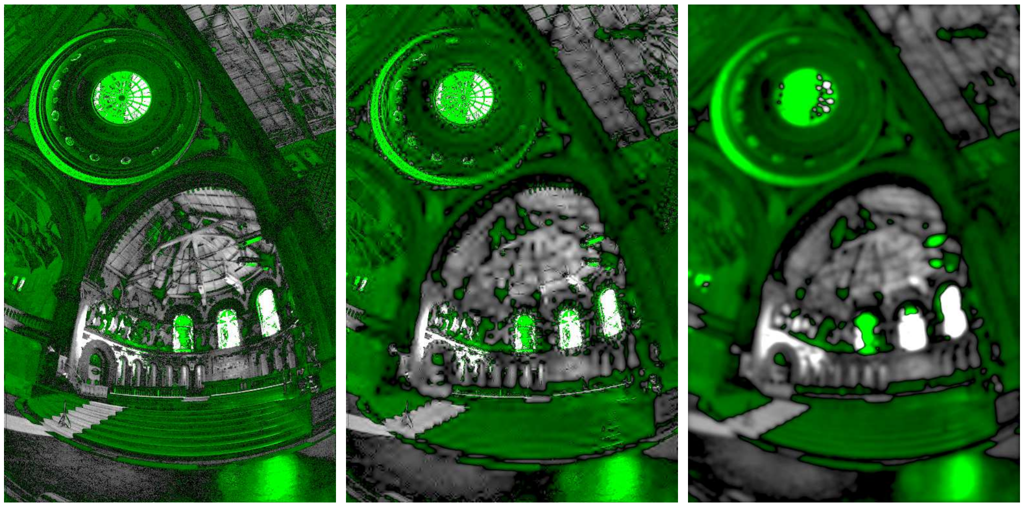

Figure 2.5: Residual frame before (left) and after (center) filtering invisible noise. Details, such as window frame, are lost when low-pass filtering (or downsampling) is used (right). Green color denotes negative values. *Memorial Church* image courtesy of Paul Debevec.

used before in CG to control rendering [42, 11, 106], the proposed filtering algorithm has been specially designed to handle HDR data and it has been optimized for speed, so that it can efficiently process video sequences. It is also different from a typical denoising algorithms, e.g. [9], since it operates on imperceivable, rather than perceivable noise. It can be used as a standard tool which guarantees that all the visual information that cannot be discerned due to imperfections of the human eye and early vision processing will be filtered out from the image.

The standard MPEG encoding already incorporates many aspects of human vision in order to improve compression efficiency. The gamma corrected color space (or transfer function) accounts for luminance masking (sometimes wrongly refereed as the Weber-Fechner law [84]). The limited spatial contrast sensitivity of the HVS is utilized by the DCT quantization matrix. Two different quantization matrices are used for inter- and intra-frames to take advantage of lower sensitivity to high temporal frequencies. However, contrast masking (or visual masking) is very poorly predicted by the mechanism of MPEG encoding. Since contrast masking is primarily responsible for masking invisible high frequency noise, we focus on modeling this aspect of the HVS to filter residual frames.

There are several methods that incorporate more accurate visual models in standard image or video encoding algorithms. An optimal DCT quantization matrix can be computed for each block or region to better predict HVS performance, which would account for both spatial contrast sensitivity and visual masking [3, 155]. This however requires storing multiple quantization matrices together with the encoded data, which is often impractical. Storing quantization matrices is not necessary in the *prequantization* scheme, in which the DCT coefficients are set to zero if they are below a predicted threshold before an image is quantized using a standard DCT quantization matrix [114]. JPEG-2000 standard takes advantages of the wavelet domain and incorporates many aspects of the HVS, which are difficult or impossible to model in the DCT domain [157].

We decided to use yet another approach for improved compression of the residual frames, which involves prefilering of frames before they are passed to the MPEG encoder. Prefiltering methods have been shown to improve video compression [12]. They do not depend on the compression algorithm and therefore do not need any changes to the encoder. In fact, incorporating physically based model of visual system in MPEG encoding is problematic because inter-frame encoded frames contain only differences between consecutive frames. Such differences do not need to correspond to temporal contrast, since optical flow is modelled only approximately by motion compensation. At the prefiltering step we can precisely model contrast masking effect in the wavelet domain, which is quite difficult and inaccurate in the DCT domain. Furthermore, prefiltering is especially effective for residual frames. This is because residual frames contain mostly low magnitude contrast, while prefiltering involves thresholding of wavelet coefficients that are below the predicted visibility level. If the wavelet coefficient are low, most of them are set to zero, which significantly improves the compression efficiency. Prefiltering affects only encoding speed while decoding speed is unaffected or improved due to the reduced stream size.

The input to our residual filtering algorithm consists of two frames: a residual frame (Figure 2.5 left) and an original HDR frame, which is a masker for the residual. Both frames should be stored in the perceptually uniform luma / chroma color space. Output of the filtering is a residual frame with high frequencies attenuated in those regions where they are not visible (Figure 2.5 center). The data flow of the algorithm is shown in Figure 2.6. Though we describe processing that is done on a luma channel, the same processing is performed for two chroma channels, which are subsampled to half of their original resolution. This approximately accounts for the differences between luminance and chrominance Contrast Sensitivity Function (CSF).

In the first step we apply the Discrete Wavelet Transform to split a residual frame into several frequency and orientation selective channels. We have

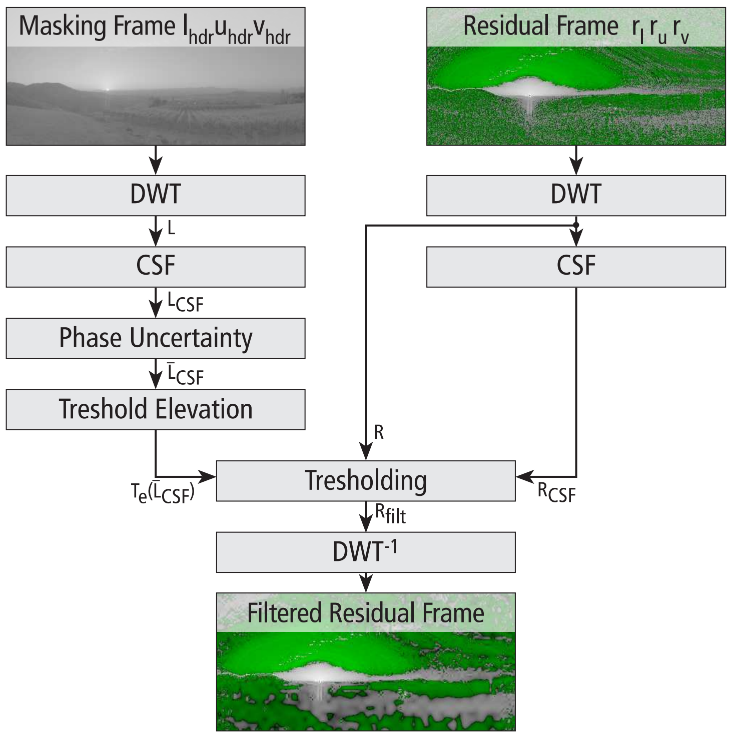

Figure 2.6: A data flow of the residual frame filtering, which removes imperceptible noise for better compression performance.

experimented with the cortex decomposition [153] performed in the Fourier domain, which can better approximate visual channels, but we rejected this approach due to prohibitively long execution times (up to 1 minute per frame). Wavelets, on the other hand, lead to computationally more efficient algorithms and were shown to be useful for modeling many aspects of the HVS [14, 157]. We used CDF 9/7 discrete wavelet basis which is also used for the lossy compression of JPEG-2000. This wavelet basis gives a good trade-off between smoothness and computational efficiency. We use only the three finest scales of the wavelet decomposition since filtering of lower spatial frequencies at coarser scales could lead to noticeable artifacts.

In the next step we account for lower sensitivity of the HVS for high frequencies, which is usually modelled with the Contrast Sensitivity Function (denoted as CSF in Figure 2.6). We weight each band of wavelet coefficients by a constant value in the same way as is done in JPEG-2000. The weighting factors for a viewing distance of 1,700 pixels are given in the table below.

Scale	LH	HL	HH
1	0.275783	0.275783	0.090078
2	0.837755	0.837755	0.701837
3	0.999994	0.999994	0.999988

The visual channels have limited phase sensitivity, ranging from 45° to more than 90°. Because of this, the masking signal affects not only regions where the values of wavelet coefficients are the highest, but may also affect neighboring regions. Phase uncertainty reduces the effect of masking at edges, as opposed to textures which show a high amount of masking. Following the point-wise extended masking in JPEG-2000 [157], we model phase uncertainty with the $L_{0.2}$-norm:

$$\bar{L}_{CSF} = \frac{1}{Card(\Theta)} \left(\sum_{\Theta} |L_{CSF}|^{0.2} \right)^{\frac{1}{0.2}} \tag{2.13}$$

where Θ denotes a neighborhood of a wavelet coefficient (we use a box 13×13 kernel in our implementation).

In the following step we predict how contrast thresholds change in the presence of a masking signal, which is an original HDR frame in our case. To model contrast masking, we employ a threshold elevation function, which we derive from the model proposed by Daly [21] (also used in [106]). We assume a masking slope of 1.0, which was shown to be appropriate for natural images [79]. We modify the original threshold elevation function to make it applicable to the perceptually uniform luma space, which we introduced in Section 2.2. Threshold elevation for this space can be approximated by the function:

$$T_e(\bar{L}_{CSF}) = \begin{cases} 1 & \text{if } \bar{L}_{CSF} \leq a \\ c \cdot (\bar{L}_{CSF}))^b & \text{otherwise} \end{cases} \tag{2.14}$$

where $\bar{L}_{CSF}$ is a wavelet coefficient, $a = 0.093071$, $b = 1.0299$ and $c = 11.535$. The function with original data points is plotted in Figure 2.7.

Next, we compare each CSF weighted coefficient of a residual frame, R_{CSF}, with the corresponding value of the threshold elevation T_e. If the residual is smaller than the visibility threshold predicted by the threshold elevation function from Equation 2.14, we can safely set this coefficient to zero without introducing visually noticeable changes. Formally, it can be written as:

$$R_{filt} = \begin{cases} 0 & \text{if } T_e(\bar{L}_{CSF}) \leq R_{CSF} \\ R & \text{otherwise} \end{cases} \tag{2.15}$$

Finally, we transform the filtered wavelets coefficients, R_{filt} back to the image domain (DWT^{-1} block in Figure 2.6).

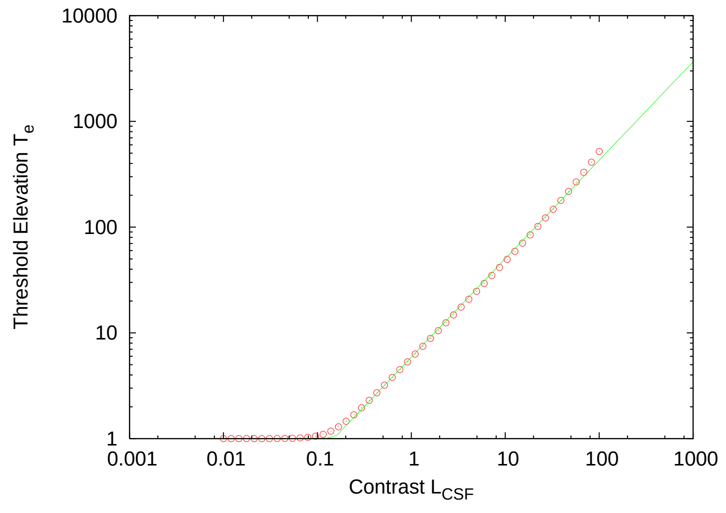

Figure 2.7: The threshold elevation function for contrast represented as wavelet coefficients. Data points were found from the model proposed by Daly [21] after conversion to new units of contrast. The solid line is a model from Equation 2.14, which approximates these data points.

The prefiltering method presented above can substantially reduce the size of a residual stream and is a reasonable trade-off between computational efficiency and accuracy of the visual model. The encoding time is affected by no more than 80% when filtering is used and it can only reduce decoding times. We have resolved to simplify some aspects of the visual model in order to bring the performance to an acceptable level. For example, we do not model the Optical Transfer Function (OTF) since we found that its local effect is negligible (close or below the MTF of a monitor) for typical viewing conditions and the flare effect would require much larger kernels or operations in the Fourier domain, which would slow down filtering significantly. For performance reasons we use wavelets, which do not model visual channels as accurately as other transformations designed especially for that purpose. Since we do not have precise information on the optical flow, we can not model temporal aspects of the CSF. The temporal CSF is partially taken into account by the MPEG encoder. Nevertheless, the proposed prefiltering method takes into account more perceptual factors than most state-of-the-art video compression techniques and can additionally handle HDR scenes.

3 Video Decoding Algorithm

The decoding of HDR video is simpler and computationally less expensive than the encoding. As illustrated in Figure 3.1, input of an HDR decoder consists of three streams: an *LDR stream*, an *auxiliary stream* and a *residual stream*. The decoder can produce either LDR video frames by decoding the LDR stream, or HDR video frames by decoding all three streams and combining the pieces of information they contain.

Two MPEG decoders are involved in the decoding of HDR frames. The first one is responsible for decoding of LDR frames, the other one decodes residual frames. The pixel values from the LDR frame (l_{ldr} in Figure 3.1) together with the decoded reconstruction function ($RF()$) give a prediction of HDR pixel values ($RF(l_{ldr})$). LDR chroma pixel values, u_{ldr} and v_{ldr}, are directly used for the prediction of HDR chroma pixel values, u_{hdr} and v_{hdr}, without the use of the reconstruction function. Those predictions are refined with the content of the residual frame to produce HDR pixel values (**Combine** in Figure 3.1). These pixels are then converted from the $l_{hdr}u_{hdr}v_{hdr}$ color space to the XYZ color space to produce a final HDR frame. An optimized implementation of the decoder can skip the last step and send the $l_{hdr}u_{hdr}v_{hdr}$ color channels directly to a graphics card, where the chroma channels are upsampled and pixel values are converted to the RGB color space.

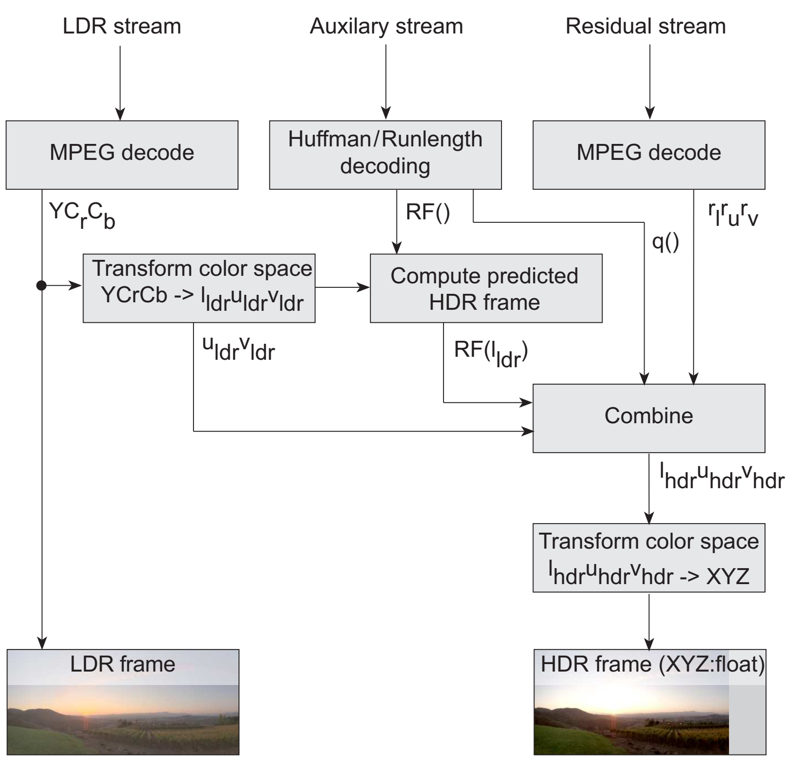

Figure 3.1: The data flow of the decoding algorithm.

4 Implementation Details

The implementation of MPEG-4 Advanced Simple Profile ISO/IEC 14496-2, available from `http://www.xvid.org/`, was used as a base MPEG encoder/decoder. However, our method is not restricted to any particular implementation and any other video or image encoder could be used instead. The backward compatible HDR encoder/decoder has been implemented as a dynamic library to simplify integration with external software. We separately implemented a set of command-line tools for encoding and decoding video streams to and from HDR image files and integrated them with the *pfstools* framework (`http://pfstools.sourceforge.net/`). An LDR stream can be played back using any video player capable of decoding MPEG-4 video. To play back an HDR stream, we have developed a custom HDR video player, which can display video on both LDR and HDR displays [117].

Since HDR video playback involves decoding two MPEG-4 streams, a LDR and a residual stream, achieving an acceptable frame rate is more challenging than in the case of an ordinary LDR video. To boost playback frame rate, we moved some parts of the decoding process to graphics hardware. We found that both color space conversion and up-sampling of color channels are computationally expensive when executed on a CPU while the same operations can be performed in almost no time on a GPU as fragment programs. The remaining parts of the decoding and encoding algorithm were implemented using the *SSE* instruction set whenever possible. Additionally, some color conversion functions were significantly accelerated with the use of fixed point arithmetic and lookup tables. All those optimizations let us achieve real-time software playback of HDR movies (25–50 frames per second for the VGA resolution, depending on a hardware configuration and quality settings of the compression).

The development of the encoding and decoding software was simplified by the reuse of the existing libraries. We would like to acknowledge the following open source projects and the authors who made their software freely available:

- *xvid 1.1.0-beta2* (http://www.xvid.org/) – an implementation of ISO MPEG-4 compliant video codec (GPL);
- *BlitzWave 0.5.0* (http://blitzwave.sourceforge.net/) for wavelet analysis and synthesis, used for filtering of residual frames (GPL);
- *Blitz++ 0.8* (http://www.oonumerics.org/blitz/) used by the BlitzWave 0.5.0 library (GPL);
- *Adaptive Huffman Encoding* by Josh MacDonald (http://www.xcf.berkeley.edu/ ali/K0D/Algorithms/huff/) used for compression of the auxiliary streams;
- *PFS TMO* (http://www.mpii.mpg.de/resources/tmo/) by Grzegorz Krawczyk, as the implementation of tone mapping algorithms, used to render LDR frames from HDR sequence.

5 Results

To test the performance of our backward compatible HDR MPEG compression, we have executed an extensive set of over 1,500 tests on images and video sequences. A well performing video compression should produce a video stream of the smallest size (measured in our tests as the number of bits per pixel) at the highest quality. Although simple arithmetic metrics, such as Signal to Noise Ratio (SNR), are usually used to measure the quality of compressed images, we follow a common practice in CG [148, 165] and also use advanced metrics that account for the aspects of the HVS. We used the following metrics to evaluate the quality of the decoded images and video sequences:

HDR VDP — Visual Difference Predictor for High Dynamic Range images [79]. This is a fidelity metric that can predict the differences between two images that are likely to be noticed by a human observer. This metric has been especially designed for HDR images and takes into account such effects as light scattering in the optics of the eye, luminance masking for the visible range of luminance, spatial contrast sensitivity, local adaptation and visual masking. The result of the HDR VDP is a probability of detection map, which assigns for each pixel a probability that the difference can be noticed. For easier interpretation of the results we have summarized the prediction of the HDR VDP with a single number, which is a percentage of pixels in an image that exceed 75% probability of detection. The lower percentage denotes a better quality, as fewer pixels are noticeably affected by compression distortions. We used the original implementation of the HDR VDP provided by the authors.

UQI — Universal Image Quality Index [143]. This is a quality metric, which is designed by modeling any image distortion as a combination of three factors: loss of correlation, luminance distortion, and contrast distortion. The index, although it does not employ any model of the HVS, shows consistency with a subjective quality measurement and performs better than the mean squared error. The quality index can range from -1 (the worst quality) to 1 (the best quality). We have implemented this metric according to the original paper [143]. To adapt this metric to HDR images, we provide for input luma values computed with Equation 2.6.

SNR — Signal to Noise Ratio. This is the simplest but also the most commonly used metric, which does not model any aspects of the HVS and may not be consistent with a subjective quality measurement. We used the standard formulas to compute the SNR for the luma values computed with Equation 2.6. The larger value of SNR usually results in higher quality.

In the following sections we analyze several aspects of our encoding scheme based on the collected test results.

5.1 Influence of Tone Mapping Operator

Although there is no restrictions on tone mapping / gamut mapping or stylizing used to obtain LDR frames, the choice of such processing will obviously affect the efficiency of compression. We tested our encoder with five tone mapping operators (TMOs) from the *pfstmo* package[1]. The labels introduced in bold in the list below are latter used to refer to the TMOs in the text and on plots.

- **Pattanaik00** — Time-Dependent Visual Adaptation for Realistic Image Display [98];
- **Durand02** — Fast Bilateral Filtering for the Display of High-Dynamic-Range Images [36];
- **Reinhard02** — Photographic Tone Reproduction for Digital Images [109];
- **Fattal02** — Gradient Domain High Dynamic Range Compression [39];
- **Drago03** — Adaptive Logarithmic Mapping for Displaying High Contrast Scenes [33].

We used the default parameters for all TMOs.

Figure 5.1 shows how the efficiency of compression is affected by a TMO. The results for most TMOs are in fact similar, with the exception of *Fattal02*, which results in significantly larger video streams for the same quality settings. This is mainly because the operator introduces the largest changes of local contrast in LDR frames, which results in the high variance of residual values. The result is consistent with our earlier considerations in Section 2.3,

[1]More details on the *pfstmo* at: http://www.mpii.mpg.de/resources/tmo/

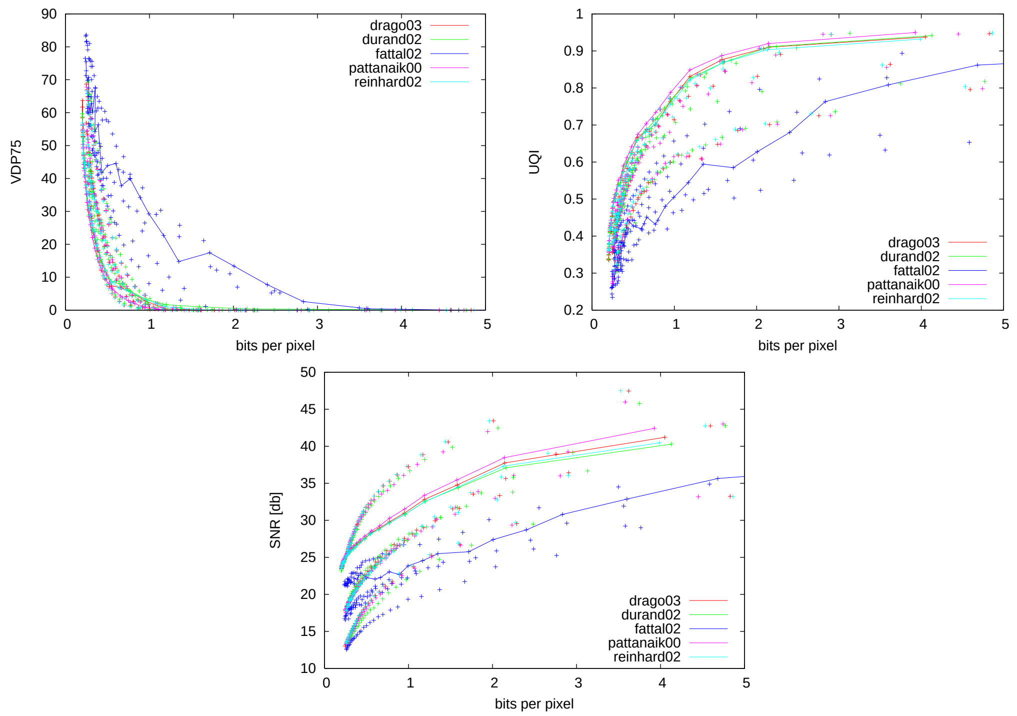

Figure 5.1: Comparison of compression performance for different tone mapping operators. See Section 5 for the description of the quality metrics. "+" denotes measurement points for a set of representative images; lines connect measurement points for one of these images.

which suggested that global TMOs are better approximated by the reconstruction function and therefore result in smaller magnitudes of the residual. If *Fattal02* is used to generate LDR video, the size of the LDR stream is also affected since high frequencies, which are poorly compressed by the MPEG encoding, are enhanced (we expect similar problems with the tone mapping approach proposed by Li et al. [74]). Nevertheless, *Fattal02* gave in our opinion the most attractive LDR images. Therefore the selection of a proper TMO for compression is often a combined aesthetic and economic choice.

5.2 The Effect of Invisible Noise Filtering

We validate the algorithm for filtering of invisible noise, described in Section 2.8, for a range of MPEG quality settings. Figure 5.2 illustrates how the size of a residual stream is reduced when the filtering is used. Note that the largest savings are possible for the best quality settings. This is because the

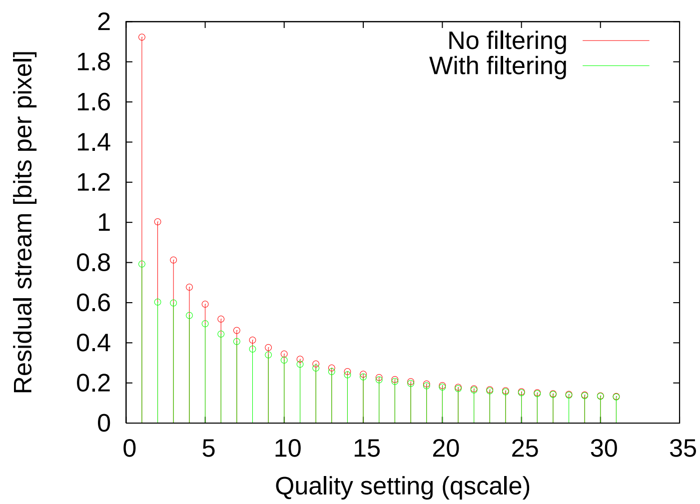

Figure 5.2: The size of a residual stream with and without invisible noise filtering switched on with respect to the quality settings. The largest savings are achieved for the best quality settings.

strength of the filtering is determined by the visibility thresholds, which do not depend on quality settings. The filtering has a minimal impact on the stream size for low quality settings since the distortions introduced by the aggressive DCT quantization are much above the visibility thresholds used in the filtering. Figure 5.4 shows how both the total stream size and quality are affected when the residual frames are filtered. Although the filtering in fact introduces changes that are detected by the HDR VDP (probably due to a mismatch in the visual models used by the filtering and the HDR VDP), the loss of quality is fully compensated by the bit-rate savings (see Figure 5.4). Moreover, we observed that the subjective quality of filtered video is better as predicted by the HDR VDP. This is because the blurry artifacts due to the wavelet based filtering are less objectionable than high frequency distortions of DCT quantization (for example blocky artifacts, as shown in Figure 5.3). Although HDR VDP can predict visible changes, it can not estimate what is the impact of these changes on the perceived quality.

5.3 Comparison with Lossy HDR Compression Methods

The performance of the proposed method (labeled as **HDR MPEG**) has been compared with two others lossy HDR compression methods:

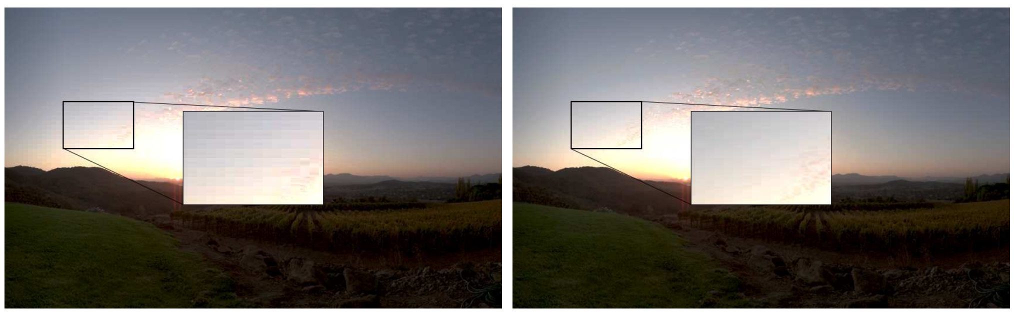

Figure 5.3: Quality comparison for an image compressed without filtering (left) and with invisible noise filtering (right). Both images were compressed to the stream of approximately the same size. The barely visible blocky artifacts in "filtered" image are well noticeable in the image compressed without filtering. No filtering: *qscale* = 6, bpp = 1.37, HDR VDP 75%= 3.12%; With Filtering: *qscale* = 2, bpp = 1.23, HDR VDP 75%= 1.11%.

- **HDRV** — Perception-motivated HDR Video Encoding [81]. This is the first lossy HDR video compression method, which, however, does not offer backward compatibility. The method encodes HDR pixels using 11 bits for luminance and twice 8 bits for chrominance. Since the resulting video stream does not contain any information on LDR frames, it can be expected that this compression method gives better results than backward compatible methods. We used the original implementation provided by the authors.

- **JPEG HDR** — Subband encoding of high dynamic range imagery [148]. This is a backward compatible HDR image encoding, which is conceptually the closest to our method. A detailed comparison of both our approach and JPEG HDR is given in Section 6.1. The method involves sub-sampling of a subband layer (a residual frame in the case of our encoding), which can lead to the loss of high frequencies. To prevent this loss, the method suggests three approaches: **pre-correction** of LDR layer, to encode within this layer high frequencies that can be lost due to sub-sampling; **post-correction** which tries to restore high frequencies that have been lost rather than modifying LDR image, and **full-sampling**, which means that no sub-sampling is performed. We used the original encoding/decoding library provided by the authors.

To evaluate the performance of intra-frame (image) compression, we ran the tests on eight representative HDR images. We chose the *Reinhard02* TMO to compare our algorithm with other lossy compression methods. This

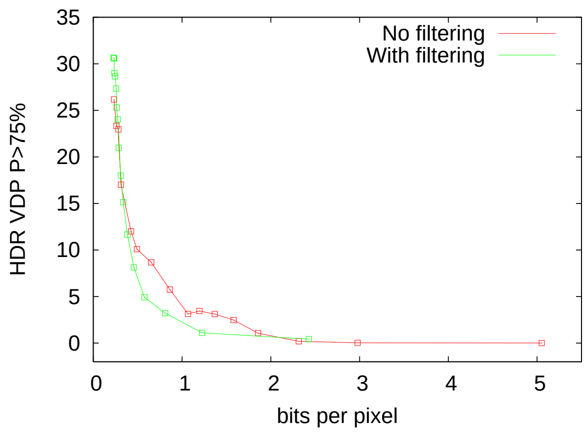

Figure 5.4: Performance of HDR video compression with and without invisible noise filtering (refer to Section 2.8). The bit-rate savings gained on the filtering compensate for the loss of quality.

TMO performed similar to the others and is also used in JPEG HDR. Figure 5.5 shows the averaged results. The HDRV encoding clearly shows the best performance for all three quality metrics. This can be explained by the lack of any information on an LDR stream, which reduces the amount of information that needs to be stored but also makes this encoding incompatible with the LDR MPEG format. For the HDR VDP and the UQI, JPEG HDR performs almost the same as our method for the pre-correction and the post-correction approach, but is worse for the full-sampling. Note that our compression method does not involve sub-sampling and therefore is closer to the full-sampling than two other approaches. JPEG HDR performs worse than our method for the SNR metric. The improved performance of our encoding over JPEG HDR for images is surprising, since the image encoding algorithms, such as JPEG, are known to perform better than intra-frame video encoding. This is due to better arithmetic encoding and a quantization matrix, which is especially optimized for images. Another difference between two methods that affects the performance is that HDR MPEG encodes information on all color channels in the residual stream while the JPEG HDR encodes only luminance in the additional subband layer (see details in Section 6.1)

The performance of inter-frame (video) compression was tested on two video sequences for both HDR MPEG and HDRV, while JPEG HDR was not included in these tests. Since both the VDP and the UQI are not suitable for video mostly because of a huge computational cost, we computed the SNR over all video frames to measure quality. The averaged results for two

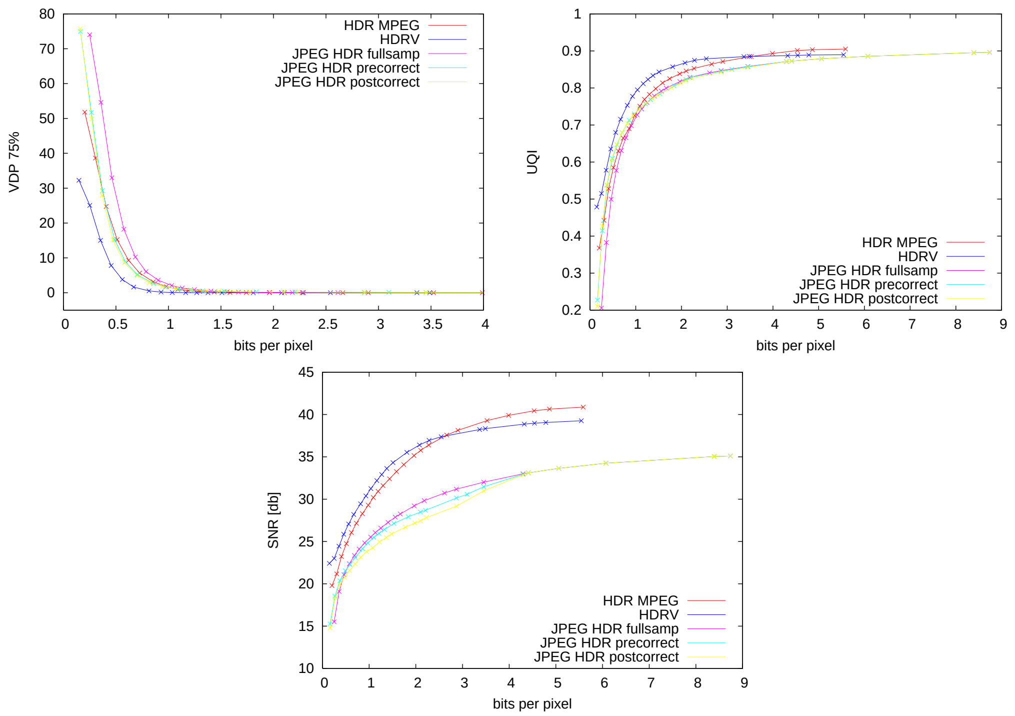

Figure 5.5: Comparison of lossy HDR compression algorithms. Averaged results for a set of images.

video sequences are shown in Figure 5.6. Similarly as for images, HDRV gave better SNR than HDR MPEG for the same number of bits. HDR MPEG, however, could achieve a higher SNR than HDRV for very high bit-rates.

5.4 The Cost of Encoding Residual Stream

The proposed HDR encoding method is designed to be an extension to the existing MPEG formats. Therefore, it is interesting to know how much more data must be stored to include additional HDR information. We plot the size of the total HDR stream (LDR + Residual + Auxiliary stream) against the size of the LDR stream in Figure 5.7. The residual stream does not seem to depend on the quality settings as much as the LDR stream. Therefore its share in the total stream size is the smallest for high quality settings. This can be expected since the residual stream encodes the difference between LDR and HDR frames, including those differences that result from lossy compression of the LDR stream (refer to the MPEG encoding and decoding stages in Figure 2.1). The lower quality LDR stream means that more

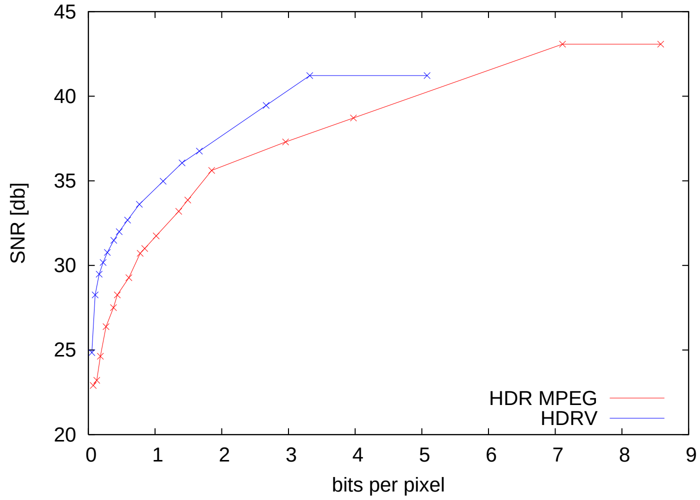

Figure 5.6: Comparison of lossy HDR compression algorithms. Averaged results for two video sequences.

information needs to be stored in the residual stream. Overall, the share of residual stream ranges from a few to 70 percent, depending on the image, quality settings and a TMO. A well chosen TMO and a decent quality settings result in a residual stream that is 25–30% of the LDR stream. The size of an auxiliary stream is negligible.

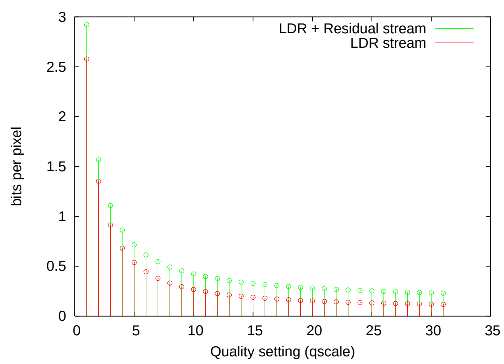

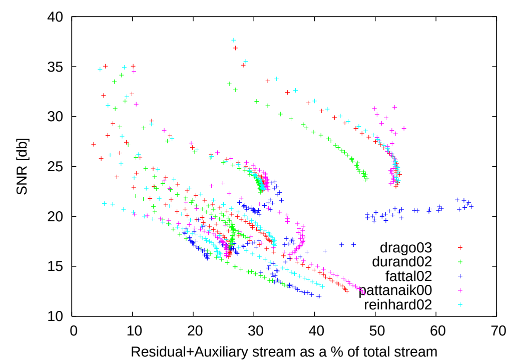

Figure 5.7: The size of a backward compatible HDR stream (LDR+Residual+Auxiliary) compared with the size of the LDR stream alone. Left: with respect to the quality settings (low *qscale* denotes high quality) for a single image and the *Reinhard02* TMO; Right: with respect to the SNR of the resulting images for several images and TMOs.

6 Discussion

6.1 Comparison with JPEG HDR

Although the proposed backward compatible HDR encoding algorithm seems to be conceptually similar to the JPEG HDR compression [148], there are several important differences between both approaches.

The HDR MPEG has been designed for video sequences, however it can be used for images as well. The JPEG HDR, on the other hand, is not suitable for video sequences in its original version. One of the reasons for that is linear scaling of subband values (the counterpart of the residual) to the range from 0-255 before those values are compressed with JPEG algorithm. This scaling depends on the maximum and minimum value in the subband, which may differ from image to image. Therefore for video sequences the scaling factor can change from frame to frame, which would result in temporal flickering and lack of temporal coherence in the subband stream. Such a lack of temporal coherence can significantly impact the performance of MPEG inter-frame compression.

The linear scaling of subband values in the JPEG HDR also makes the quantization of the subband layer difficult to predict and control. The JPEG HDR will quantize subband values with high accuracy for those images that lead to small magnitude of subband values, perhaps wasting some bit-rate on invisible contrast details. For another set images, which result in large magnitude of subband values, the JPEG HDR can quantize too coarsely, leading to contouring artifacts. HDR MPEG quantizes color values consistently for consecutive frames and the quantizer is based on the visibility thresholds of the HVS rather than frame content. This makes the result of the quantization predictable and easy to control with the MPEG quality settings.

JPEG HDR reduces quantization in the subband image by skipping a small percent of the brightest and the darkest pixels in an image. This however can lead to loss of some small but salient details, such as the sun in

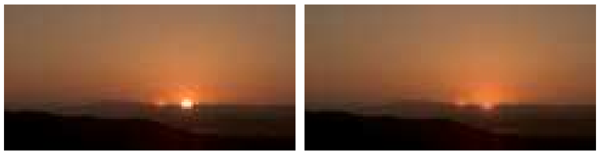

Figure 6.1: Very bright pixels in the original image (left) are lost after compression with JPEG HDR (right) at quality setting 90 and with the *precorrection* (default settings). This is because a small percentage of the brightest and darkest pixels is skipped when computing minimum and maximum value of the subband image.

Figure 6.1).

Unlike the JPEG HDR, the proposed compression method does not impose any restrictions on the choice of a TMO and a gamut mapping algorithm. A TMO for HDR MPEG can saturate both luminance and color, change color vales and enhance local contrast. Such changes may result in lower compression ratio, but both LDR and HDR frames will be preserved in the resulting video stream. JPEG HDR will lose most of color differences between HDR and LDR since it does not store color in the subband layer. Such unrestricted control over appearance of both LDR and HDR streams is very important for our major application - a storage of DVD movies.

As discussed in Section 2.3, the JPEG HDR assumes a non-optimal reconstruction function, which worsens the compression performance. This is visible in Figure 5.1, though the results would be even more favorable if we used JPEG algorithm instead of MPEG intra-frame compression to encode images.

Finally, a sub-sampling of the subband layer in the JPEG HDR may lead to loss of visible details. While the pre-correction may be used to avoid loss of high frequency details, this leads to a modified LDR images, which is not acceptable for our application. The post-correction, although does not modify the source image, does not give as good results as pre-correction. Full-sampling, on the other hand, does not give as good compression ratio as the other two approaches. The counterpart of sub-sampling in the proposed HDR MPEG is the filtering of residual frames (see Section 2.8). The filtering has a similar function as the sub-sampling — to reduce high frequency noise and improve compression, but it does it in a more selective manner. The proposed filtering removes only those high frequency details, which are not visible and therefore can be smoothed out without impairing the visual quality of the resulting video. Sub-sampling, obviously, cannot give such a guarantee.

6.2 Comparison with Wavelet Compander

Li et al. [74] propose a multiscale wavelet architecture, which can compress an HDR image to a lower bit-depth and later expand it to obtain a result that is close to the original HDR image (the so-called compander). The information loss is reduced by amplifying low amplitudes and high frequencies at the compression stage, so that they survive the quantization step to the 8-bit LDR image. Such technique is conceptually similar to *pre-correction* in JPEG HDR. Since the expansion is a fully symmetric inverted process, the amplified signals are properly suppressed to their initial level in the companded HDR image. To further reduce the information loss, the compressed image is iteratively modified to improve the correlation of its subbands with respect to the original HDR image. The authors observe a good visual quality of both the compressed and companded images, but they admit that any guarantee concerning their fidelity to tone mapped (i.e. undergoing just one compression iteration) and original HDR images cannot be given. Moreover, they could clearly see the differences between the corresponding image pairs but they describe them as "visually not disturbing". The obtained PNSR for the companded HDR image is even worse than for ordinary LUT (Look-Up-Table) companding. Given our research goals, the lack of fidelity of tone mapped images is not acceptable, since the original material quality cannot be compromised. Also, the multiscale wavelet framework as proposed by Li et al. severely limits the choice of tone mapping operator. The emphasis on high frequencies at the compression step makes the proposed framework less suitable for standard JPEG and MPEG techniques, which use the quantization matrices that are perceptually tuned to discard a great deal of visually non-important high frequencies. This is confirmed by relatively poor compression rates reported the authors when they attempted to combine JPEG with their companding. We also achieved poor results for the tone mapping operator that enhanced high frequencies (see *Fattal02* in Figure 5.1). Moreover, many existing quantization schemes, which incorporate a visual masking model [155, 94], assume that due to a better visibility of the contouring artifacts in smooth image regions than in textured ones a finer quantization is required in those regions. For such image compression approaches superficially high quality coefficients should be set to preserve the high frequency details as required by the compander. It is not clear, how the compander approach can be adopted for lossy HDR video compression, in which apart from just raised quality concerns the issues of temporal coherence and computation efficiency arise (the authors recommend oversampling, i.e., handling the subband computation in the full image resolution to avoid aliasing and do not report any timings for their compander).

6.3 Conclusions

In this report we propose an algorithm for a backward compatible compression of HDR video. The video streams produced by the algorithm are about 30% larger than the corresponding LDR MPEG streams, which is a modest cost compared to the huge storage requirements of High Definition video. We evaluate the performance of the proposed algorithm and compare it with two other existing lossy HDR encodings: JPEG HDR and HDRV. Although our algorithm is not optimized for images, it performs better than JPEG HDR. Since HDRV compression is not backward compatible and does not preserve LDR frames, it achieves better performance than our method. We discuss and dismiss the HDR encoding methods that require modification of LDR content. Such modifications are not acceptable for a professional DVD movie production. We believe that the proposed methods offers a solution for the smooth transition of from LDR to HDR content for the DVD movie distribution.

Bibliography

[1] A. Adams. *The Print, The Ansel Adams Photography Series 3*. New York Graphic Society, 1981.

[2] A. Agarwala, M. Dontcheva, M. Agrawala, S. Drucker, A. Colburn, B. Curless, D. Salesin, and M. Cohen. Interactive digital photomontage. *ACM Trans. on Graph.*, 23(3):294–302, 2004.

[3] A. Ahumada and H. Peterson. Luminance-model-based DCT quantization for color image compression. In *Human Vision, Visual Processing and Digital Display*, SPIE, volume 3299, pages 191–201, 1993.

[4] M. Ashikhmin. A tone mapping algorithm for high contrast images. In *Rendering Techniques 2002: 13th Eurographics Workshop on Rendering*, pages 145–156, 2002.

[5] M. Ashikhmin. A tone mapping algorithm for high contrast images. In *13th Eurographics Workshop on Rendering*. Eurographics, June 2002.

[6] P. Barten. Subjective image quality of high-definition television pictures. In *Proc. of the Soc. for Inf. Disp.*, volume 31, pages 239–243, 1990.

[7] P. G. Barten. *Contrast sensitivity of the human eye and its effects on image quality*. SPIE – The International Society for Optical Engineering, P.O. Box 10 Bellingham Washington 98227-0010, 1999. ISBN 0-8194-3496-5.

[8] S. Bech, R. Hamberg, M. Nijenhuis, C. Teunissen, H. L. de Jong, P. Houben, and S. Pramanik. The RaPID perceptual image description method (RaPID). In *SPIE Proceedings*, volume 2657, pages 317–328, P.O. Box 10 Bellingham Washington 98227-0010, January/February 1996. SPIE – The International Society for Optical Engineering.

[9] E. P. Bennett and L. McMillan. Video enhancement using per-pixel virtual exposures. *ACM Trans. on Graph.*, 24(3):845–852, 2005.

[10] R. Bogart, F. Kainz, and D. Hess. OpenEXR image file format. In *ACM SIGGRAPH 2003, Sketches & Applications*, 2003.

[11] M. R. Bolin and G. W. Meyer. A perceptually based adaptive sampling algorithm. In *Proceedings of ACM SIGGRAPH 1998*, Computer Graphics Proceedings, Annual Conference Series, pages 299–309, New York, NY, USA, 1998. ACM Press.

[12] P. Border and P. Guillotel. Perceptually adapted MPEG video encoding. In *Human Vision and Electronic Imaging V*, SPIE, volume 3959, pages 168–175, 2000.

[13] I. Borg and P. Groenen. *Modern Multidimensional Scaling: Theory and Applications*. Springer-Verlag, New York, 1997.

[14] A. P. Bradley. A wavelet visible difference predictor. *IEEE Transactions on Image Processing*, 8(5):717–730, 1999.

[15] P. Burt and R. Kolczynski. Enhanced image capture through fusion. In *Proc. of International Conference on Computer Vision (ICCV)*, pages 173–182, 1993.

[16] J. D. Carroll. Models and methods for multidimensional analysis of preference choice or other dominance data. In E. D. Lantermann and H. Feger, editors, *Similarity and Choice*, pages 105–155. Huber, Bern, Switzerland, 1980.

[17] J. Cataliotti and A. Gilchrist. Local and global processes in surface lightness perception. *Perception & Psychphysics*, 57(2):125–135, 1995.

[18] K. Chiu, M. Herf, P. Shirley, S. Swamy, C. Wang, and K. Zimmerman. Spatially nonuniform scaling functions for high contrast images. In *Graphics Interface '93*, pages 245–253, Toronto, Ontario, Canada, May 1993. Canadian Information Processing Society.

[19] CIE. *An Analytical Model for Describing the Influence of Lighting Parameters Upon Visual Performance*, volume 1. Technical Foundations, CIE 19/2.1. International Organization for Standardization, 1981.

[20] T. Cooper. Modifications to retinex to relax reset nonlinearity and implement segmentation constraints. In *Proc. IS&T/ SPIE Symposium on Electronic Imaging: Science and Technology*, San José January 2002.

[21] S. Daly. The Visible Differences Predictor: An algorithm for the assessment of image fidelity. In A. Watson, editor, *Digital Image and Human Vision*, pages 179–206. Cambridge, MA: MIT Press, 1993.

[22] S. Daly. Method and apparatus for determining visually perceptible differences between images. In *US Patent 5,394,483*. 1995.

[23] S. J. Daly and X. Feng. Bit-depth extension using spatiotemporal microdither based on models of the equivalent input noise of the visual system. In *Color Imaging VIII: Processing, Hardcopy, and Applications*, SPIE, volume 5008, pages 455–466, 2003.

[24] S. J. Daly and X. Feng. Decontouring: Prevention and removal of false contour artifacts. In *Human Vision and Electronic Imaging IX*, SPIE, volume 5292, pages 130–149, 2004.

[25] Data Analysis Products Division, MathSoft, Seattle, WA. *S-PLUS 2000 Guide to Statistics*, 1999.

[26] P. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *Proceedings of ACM SIGGRAPH 1997*, Computer Graphics Proceedings, Annual Conference Series, pages 369–378, 1997.

[27] R. Deeley, N. Drasdo, and W. N. Charman. A simple parametric model of the human ocular modulation transfer function. *Ophthalmology and Physiological Optics*, 11:91–93, 1991.

[28] K. Devlin, A. Chalmers, A. Wilkie, and W. Purgathofer. Tone reproduction and physically based spectral rendering. In *Eurographics 2002: State of the Art Reports*, pages 101–123. Eurographics, 2002.

[29] DICOM. Part 14: Grayscale standard display function. In *Digital Imaging and Communications in Medicine (DICOM)*. 2001.

[30] F. Drago, W. Martens, K. Myszkowski, and H.-P. Seidel. Perceptual Evaluation of Tone Mapping Operators with Regard to Similarity and Preference. Technical Report MPI-I-2002-4-002, `http://domino.mpi-sb.mpg.de/internet/reports.nsf/AG4NumberView?OpenView`, Max-Planck-Institut fuer Informatik, Oct. 2002.

[31] F. Drago, W. L. Martens, K. Myszkowski, and N. Chiba. Design of a tone mapping operator for high dynamic range images based upon

psychophysical evaluation and preference mapping. In *The human Vision and Electronic Imaging VIII Conference*. IS&T SPIE Electronic Imaging, 2003.

[32] F. Drago and K. Myszkowski. Validation proposal for global illumination and rendering techniques. *Computers & Graphics*, 25(3):511–518, June 2001. ISSN 0097-8493.

[33] F. Drago, K. Myszkowski, T. Annen, and N. Chiba. Adaptive logarithmic mapping for displaying high contrast scenes. *Computer Graphics Forum, Proceedings of Eurographics 2003*, 22(3):419–426, 2003.

[34] R. Dumont, F. Pellacini, and J. A. Ferwerda. Perceptually-driven decision theory for interactive realistic rendering. *ACM Trans. on Graph.*, 22(2):152–181, 2003.

[35] F. Durand and J. Dorsey. Interactive tone mapping. In *Rendering Techniques 2000: 11th Eurographics Workshop on Rendering*, pages 219–230, 2000.

[36] F. Durand and J. Dorsey. Fast bilateral filtering for the display of high-dynamic-range images. *ACM Trans. on Graph.*, 21(3):257–266, 2002.

[37] F. Durand and J. Dorsey. Fast bilateral filtering for the display of high-dynamic-range images. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pages 257–266, 2002.

[38] M. Fairchild and G. Johnson. Image appearance modeling. In *Human Vision and Electronic Imaging VIII*, SPIE, volume 5007, pages 149–160, 2003.

[39] R. Fattal, D. Lischinski, and M. Werman. Gradient domain high dynamic range compression. *ACM Trans. on Graph.*, 21(3):249–256, 2002.

[40] J. Ferwerda, S. Pattanaik, P. Shirley, and D. Greenberg. A model of visual adaptation for realistic image synthesis. In *Proceedings of SIGGRAPH 96*, Computer Graphics Proceedings, Annual Conference Series, pages 249–258, Aug. 1996.

[41] J. A. Ferwerda, H. Rushmeier, and B. Watson. Frontiers in perceptually based image synthesis: Modeling, rendering, display, validation. *SIGGRAPH 2003, Course*, July 2003.

[42] J. A. Ferwerda, P. Shirley, S. N. Pattanaik, and D. P. Greenberg. A model of visual masking for computer graphics. In *Proceedings of ACM SIGGRAPH 1997*, Computer Graphics Proceedings, Annual Conference Series, pages 143–152, New York, NY, USA, 1997. ACM Press.

[43] J. Frankle and J. McCann. Method and apparatus for lightness imaging. *US Patent No: 4,384,336*, May 17 1983.

[44] B. Funt and F. Ciuera. Control parameters for retinex. In *AIC 2001, Proc 9th Congress of the International Color Association*, June 2001.

[45] B. Funt, F. Ciuera, and J. McCann. Retinex in Matlab. In *Proc. IS&T/SID Eighth Color Imaging Conference*, pages 112–121, 2000.

[46] B. Funt, F. Ciuera, and J. McCann. Retinex in matlab. In *Proc. IS&T/SID Eighth Color Imaging Conference*, pages 112–121, Scottsdale 2000.

[47] M. Georgeson and G. Sullivan. Contrast constancy: Deblurring in human vision by spatial frequency channels. *J.Physiol.*, 252:627–656, 1975.

[48] A. Gilchrist, C. Kossyfidis, F. Bonato, T. Agostini, J. Cataliotti, X. Li, B. Spehar, V. Annan, and E. Economou. An anchoring theory of lightness perception. *Psychological Review*, 106(4):795–834, 1999.

[49] A. A. Gooch, S. C. Olsen, J. Tumblin, and B. Gooch. Color2gray: Salience-preserving color removal. *ACM Trans. on Graph.*, 24(3), 2005.

[50] N. Goodnight, R. Wang, C. Woolley, and G. Humphreys. Interactive time-dependent tone mapping using programmable graphics hardware. In *Rendering Techniques 2003: 14th Eurographics Symposium on Rendering*, pages 26–37, 2003.

[51] D. Heeger and P. Teo. A model of perceptual image fidelity. In *Proc. of IEEE Int'l Conference Image Processing*, pages 343–345, 1995.

[52] D. Hood and M. Finkelstein. Sensitivity to light. In K. Boff, L. Kaufman, and J. Thomas, editors, *Handbook of Perception and Human Performance: 1. Sensory Processes and Perception*, volume 1, New York, 1986. Wiley.

[53] B. Horn. Determining lightness from an image. *Computer Graphics and Image Processing*, 3(1):277–299, 1974.

[54] R. Hunt. *The Reproduction of Colour in Photography, Printing and Television: 5th Edition*. Fountain Press, 1995.

[55] A. Hurlbert. Formal connections between lightness algorithms. *Journal of the Optical Society of America A*, 3(10):1684–1693, 1986.

[56] IEC 61966-2-1:1999. *Multimedia systems and equipment - Colour measurement and management - Part 2-1: Colour management - Default RGB colour space - sRGB*. International Electrotechnical Commission, 1999.

[57] ISO-IEC-14496-2. *Information technology: Coding of audio-visual objects, Part 2: Visual*. International Organization for Standardization, Geneva, Switzerland, 1999.

[58] R. Janssen. *Computational Image Quality*. Spie Press, Bellingham, Washington USA, 2001.

[59] T. Janssen and F. Blommaert. Image quality semantics. *Journal of Imaging Science and Technology*, 41(5):555–560, 1997.

[60] T. Janssen and F. Blommaert. A computational approach to image quality. *Displays*, 21(4):129–142, October 2000.

[61] T. Janssen and F. Blommaert. Predicting the usefulness and naturalness of colour reproductions. *Journal of Imaging Science and Technology*, 44(2):93–104, March/April 2000.

[62] T. Janssen and F. Blommaert. Visual metrics: discriminative power through flexibility. *Perception*, 29(8):965–980, June 2000.

[63] E. W. Jin, X.-F. Feng, and J. Newell. The development of a color visual difference model (CVDM). In *IS&T's 1998 Image Processing, Image Quality, Image Capture, Systems Conf.*, pages 154–158, 1998.

[64] D. J. Jobson, Z. Rahman, and G. A. Woodell. Properties and performance of a center/surround retinex. *IEEE Transactions on Image Processing*, 6(3):451–462, 1997.

[65] D. J. Jobson, Z. Rahman, and G. A. Woodell. A multi-scale retinex for bridging the gap between color images and the human observation of scenes. *IEEE Transactions on Image Processing: Special Issue on Color Processing*, 6(7):965–976, July 1997.

[66] S. Kang, M. Uyttendaele, S. Winder, and R. Szeliski. High dynamic range video. *ACM Trans. on Graph.*, 22(3):319–325, 2003.

[67] F. A. A. Kingdom and P. Whittle. Contrast discrimination at high contrasts reveals the influence of local light adaptation on contrast processing. *Vision Research*, 36(6):817–829, 1996.

[68] G. Krawczyk, K. Myszkowski, and H.-P. Seidel. Perceptual effects in real-time tone mapping. In *SCCG '05: Proceedings of the 21st spring conference on Computer graphics*, pages 195–202, New York, NY, USA, 2005. ACM Press.

[69] J. Kuang, H. Yamagishi, G. Johnson, and M. Fairchild. Testing hdr rendering algorithms. In *IS&T/SID Color Imaging Conference*, pages 315–320, 2004.

[70] E. Land and J. McCann. Lightness and retinex theory. *Journal of the Optical Society of America*, 61(1):1–11, January 1971.

[71] E. H. Land and J. J. McCann. Lightness and the retinex theory. *Journal of the Optical Society of America*, 61(1):1–11, 1971.

[72] P. Ledda, G. Ward, and A. Chalmers. A wide field, high dynamic range, stereographic viewer. In *GRAPHITE 2003*, pages 237–244, 2003.

[73] B. Li, G. Meyer, and R. Klassen. A comparison of two image quality models. In *Human Vision and Electronic Imaging III*, SPIE, volume 3299, pages 98–109, 1998.

[74] Y. Li, L. Sharan, and E. H. Adelson. Compressing and companding high dynamic range images with subband architectures. *ACM Trans. on Graph.*, 24(3):836–844, 2005.

[75] C. Lloyd and R. Beaton. Design of spatial-chromatic human vision model for evaluating full-color display systems. In *Human Vision and Electronic Imaging: Models, Methods, and Appl.*, SPIE, volume 1249, pages 23–37, 1990.

[76] J. Lubin. A visual discrimination model for imaging system design and evaluation. In *Vis. Models for Target Detection*, pages 245–283, 1995.

[77] J. Lubin and A. Pica. A non-uniform quantizer matched to the human visual performance. *Society of Information Display Int. Symposium Technical Digest of Papers*, (22):619–622, 1991.

[78] I. Lucian, S. Felicia, S. Charles, and H. Siefkien. Digital encode and method of encoding high dynamic range video images. In *US Patent 6,867,717*. 2005.

[79] R. Mantiuk, S. Daly, K. Myszkowski, and H.-P. Seidel. Predicting visible differences in high dynamic range images - model and its calibration. In *Human Vision and Electronic Imaging X*, SPIE, volume 5666, pages 204–214, 2005.

[80] R. Mantiuk, A. Efremov, K. Myszkowski, and H.-P. Seidel. Backward compatible high dynamic range mpeg video compression. *ACM Trans. on Graph. (conditionally accepted)*, 25(3), 2006.

[81] R. Mantiuk, G. Krawczyk, K. Myszkowski, and H.-P. Seidel. Perception-motivated high dynamic range video encoding. *ACM Trans. on Graph.*, 23(3):730–738, 2004.

[82] R. Mantiuk, K. Myszkowski, and H.-P. Seidel. Visible difference predicator for high dynamic range images. In *Proceedings of IEEE International Conference on Systems, Man and Cybernetics*, pages 2763–2769, 2004.

[83] R. Mantiuk, K. Myszkowski, and H.-P. Seidel. A perceptual framework for contrast processing of high dynamic range images. In *APGV '05: 2nd Symposium on Appied Perception in Graphics and Visualization*, pages 87–94, 2005.

[84] R. Mantiuk, K. Myszkowski, and H.-P. Seidel. Lossy compression of high dynamic range images and video. In B. E. Rogowitz, T. N. Pappas, and S. J. Daly, editors, *Human Vision and Electronic Imaging XI*, volume 6057 of *Proceedings of SPIE*, page 60570V, San Jose, USA, February 2006. SPIE.

[85] D. Marimont and B. Wandell. Matching colour images: the effects of axial chromatic aberration. *J. Opt. Soc. Am. A*, 11(12):2113–3122, 1994.

[86] K. Matkovic, L. Neumann, and W. Purgathofer. A survey of tone mapping techniques. In *13th Spring Conference on Computer Graphics*, pages 163–170. Comenius University, Bratislava, Slovakia, June 1997.

[87] W. Matusik and H. Pfister. 3D TV: a scalable system for real-time acquisition, transmission, and autostereoscopic display of dynamic scenes. *ACM Trans. on Graph.*, 23(3):814–824, 2004.

[88] A. McNamara. Visual perception in realistic image synthesis. *Computer Graphics Forum*, 20(4):211–224, 2001.

[89] J. Meulman, W. J. Heiser, and J. D. Carroll. PREFMAP-3 user's guide. Technical report, Bell Telephone Laboratories, Murray Hill, NJ, USA, 1986.

[90] A. Michelson. *Studies in Optics*. U. Chicago Press., 1927.

[91] N. Miller, P. Ngai, and D. Miller. The application of computer graphics in lighting design. *Journal of the Illuminating Engineering Society*, 14(1):6–26, 1984.

[92] P. Moon and D. Spencer. Visual data applied to lighting design. *J. Opt. Soc. Am.*, 34(605), 1944.

[93] N. N. Personal communication (witheld for annonimity). 2005.

[94] M. Nadenau. *Integration of Human color vision Models into High Quality Image Compression*. PhD thesis, École Polytechnique Fédéral Lausane, 2000.

[95] S. Nayar and V. Branzoi. Adaptive dynamic range imaging: Optical control of pixel exposures over space and time. In *Proc. of IEEE International Conference on Computer Vision (ICCV 2003)*, pages 1168–1175, 2003.

[96] S. Nayar and T. Mitsunaga. High dynamic range imaging: Spatially varying pixel exposures. In *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition*, pages 472–479, 2000.

[97] M. Nijenhuis, R. Hamberg, C. Teunissen, S. Bech, H. L. de Jong, P.Houben, and S. Pramanik. Sharpness, sharpness related attributes, and their physical correlates. In *SPIE Proceedings*, volume 3025, pages 173–184, P.O. Box 10 Bellingham Washington 98227-0010, February 1997. SPIE – The International Society for Optical Engineering.

[98] S. Pattanaik, J. Tumblin, H. Yee, and D. Greenberg. Time-dependent visual adaptation for realistic image display. In *Proceedings of ACM SIGGRAPH 2000*, Computer Graphics Proceedings, Annual Conference Series, pages 47–54, New York, NY, USA, 2000. ACM Press.

[99] S. N. Pattanaik, J. A. Ferwerda, M. D. Fairchild, and D. P. Greenberg. A multiscale model of adaptation and spatial vision for realistic image

display. In *Siggraph 1998, Computer Graphics Proceedings*, pages 287–298, 1998.

[100] E. Peli. Contrast in complex images. *Journal of the Optical Society of America A*, 7(10):2032–2040, 1990.

[101] F. Pellacini, J. A. Ferwerda, and D. P. Greenberg. Toward a psychophysically-based light reflection model for image synthesis. In *Proceedings of ACM SIGGRAPH 2000*, Computer Graphics Proceedings, Annual Conference Series, pages 55–64. ACM Press / ACM SIGGRAPH / Addison Wesley Longman, 2000.

[102] P. Perez, M. Gangnet, and A. Blake. Poisson image editing. *ACM Trans. on Graph.*, 22(3):313–318, 2003.

[103] K. Perlin and E. M. Hoffert. Hypertexture. In *Computer Graphics (Proceedings of SIGGRAPH 89)*, volume 23, pages 253–262, July 1989.

[104] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. *Numerical Recipes in C*. Cambridge Univ. Press, 1993.

[105] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. *Numerical Recipes in C++*, chapter 2.7, pages 87–92. Cambridge Univ. Press, second edition, 2002.

[106] M. Ramasubramanian, S. N. Pattanaik, and D. P. Greenberg. A perceptually based physical error metric for realistic image synthesis. In *Proceedings of ACM SIGGRAPH 1999*, Computer Graphics Proceedings, Annual Conference Series, pages 73–82, New York, NY, USA, 1999. ACM Press.

[107] R. Reid and R. Shapley. Brightness induction by local contrast and the spatial dependence of assimilation. *Vision Research*, 28(1):115–132, 1988.

[108] E. Reinhard. Parameter estimation for photographic tone reproduction. *Journal of Graphics Tools*, 7(1):45–52, 2002.

[109] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda. Photographic tone reproduction for digital images. *ACM Trans. on Graph.*, 21(3):267–276, 2002.

[110] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda. Photographic tone reproduction for digital images. *ACM Transs on Graph.*, 21(3):267–276, 2002.

[111] E. Reinhard, G. Ward, S. Pattanaik, and P. Debevec. *High Dynamic Range Imaging. Data Acquisition, Manipulation, and Display.* Morgan Kaufmann, 2005.

[112] M. Robertson, S. Borman, and R. Stevenson. Dynamic range improvement through multiple exposures. In *Proceedings of the 1999 International Conference on Image Processing (ICIP-99)*, pages 159–163, Los Alamitos, CA, Oct. 24–28 1999.

[113] B. E. Rogowitz, T. Frese, J. Smith, C. A. Bouman, and E. Kalin. Perceptual image similarity experiments. In *Human Vision and Electronic Imaging III*, SPIE, volume 3299, pages 576–590, 1998.

[114] R. J. Safranek. JPEG compliant encoder using perceptually based quantization. In *Human Vision, Visual Processing, and Digital Display IV*, SPIE, volume 1913, pages 117–126, 1993.

[115] K. Saito. Electronic image pickup device. *Japanese Patent 07-254965*, 1995.

[116] C. Schlick. Quantization techniques for the visualization of high dynamic range pictures. In *Photorealistic Rendering Techniques*, Eurographics, pages 7–20. Springer-Verlag Berlin Heidelberg New York, 1994.

[117] H. Seetzen, W. Heidrich, W. Stuerzlinger, G. Ward, L. Whitehead, M. Trentacoste, A. Ghosh, and A. Vorozcovs. High dynamic range display systems. *ACM Trans. on Graph.*, 23(3):757–765, 2004.

[118] H. Seetzen, L. Whitehead, and G. Ward. High dynamic range display using low and high resolution modulators. In *The Society for Information Display International Symposium*, 2003.

[119] M. Sezan, K. Yip, and S. Daly. Uniform perceptual quantization: Applications to digital radiography. *IEEE Transactions on Systems, Man, and Cybernetics*, 17(4):622–634, 1987.

[120] M. Sezan, K. Yip, and S. Daly. Uniform perceptual quantization: Applications to digital radiography. *IEEE Transactions on Systems, Man, and Cybernetics*, 17(4):622–634, 1987.

[121] R. Shapley and C. Enroth-Cugell. Visual adaptation and retinal gain controls. In *Progress in Retinal Research*, volume 3, pages 263–346. Oxford: Pergamon Press, 1984.

[122] K. Shen and E. Delp. Wavelet based rate scalable video compression. *IEEE Transactions on Circuits and Systems for Video Technology*, 9(1):109–122, 1999.

[123] G. Spencer, P. Shirley, K. Zimmerman, and D. Greenberg. Physically-based glare effects for digital images. In *Proceedings of ACM SIGGRAPH 95*, pages 325–334, 1995.

[124] SpheronVR. http://www.spheron.com/.

[125] SPSS Inc., Chicago, WA. *SPSS Base 10.0 User's Guide*, 1999.

[126] S. Stevens and J. Stevens. Brightness function: parametric effects of adaptation and contrast. *Journal of the Optical Society of America*, 50(11):1139A, Nov. 1960.

[127] H. Stone and J. Siedel. *Sensory Evaluation Practices*. Academic Press, London, 1993.

[128] P. Sumantha and Y. Hector. Adaptive gain control for high dynamic range image display. In *Proceedings of ACM SIGGRAPH 2000*, pages 55–64, 2002.

[129] J. Sun, J. Jia, C.-K. Tang, and H.-Y. Shum. Poisson matting. *ACM Trans. on Graph.*, 23(3):315–321, 2004.

[130] SunnybrookTechnologies. Hdr display. *http://www.sunnybrooktech.com/*.

[131] T. Tanaka and N. Ohnishi. Painting-like image emphasis based on human vision systems. *Computer Graphics Forum*, 16(3):253–260, August 1997. ISSN 1067-7055.

[132] C. Taylor, Z. Pizlo, J. P. Allebach, and C. Bouman. Image quality assessment with a Gabor pyramid model of the Human Visual System. In *Hum. Vis. and Elect. Imaging*, SPIE, volume 3016, pages 58–69, 1997.

[133] W. B. Thomspon, P. Shirley, and J. A. Ferwerda. A spatial post-processing algorithm for images of night scenes. *Journal of Graphics Tools*, 7(1):1–12, 2002.

[134] W. Torgerson. *Theory and methods of scaling*. Wiley, New York, 1958.

[135] L. Tucker and S. Messick. An individual differences model for multidimensional scaling. *Psychometrika*, 38:333–368, 1963.

[136] J. Tumblin. *Three Methods For Detail-Preserving Contrast Reduction For Displayed Images*. PhD thesis, Georgia Institute of Technology, 1999.

[137] J. Tumblin, J. K. Hodgins, and B. K. Guenter. Two methods for display of high contrast images. *ACM Trans. on Graph.*, 18(1):56–94, 1999.

[138] J. Tumblin and H. E. Rushmeier. Tone reproduction for realistic images. *IEEE Computer Graphics and Applications*, 13(6):42–48, Nov. 1993.

[139] J. Tumblin and G. Turk. LCIS: A boundary hierarchy for detail-preserving contrast reduction. In *Siggraph 1999, Computer Graphics Proceedings*, pages 83–90, 1999.

[140] F. Van Nes and M. Bouman. Spatial modulation transfer in the human eye. *Journal of the Optical Society of America*, 57:401–406, 1967.

[141] B. Walter, S. Pattanaik, and D. Greenberg. Using perceptual texture masking for efficient image synthesis. *Computer Graphics Forum*, 21(3):393–399, 2002.

[142] B. Wandell. *Foundations of Vision*. Sinauer Associates, Inc., 1995.

[143] Z. Wang and A. Bovik. A universal image quality index. *IEEE Signal Processing Letters*, 9(3):81–84, 2002.

[144] G. Ward. High dynamic range image encodings. *http://www.anyhere.com/gward/hdrenc/hdr_encodings.html*.

[145] G. Ward. Real pixels. *Graphics Gems II*, pages 80–83, 1991.

[146] G. Ward. Real pixels. In J. Arvo, editor, *Graphics Gems II*, pages 80–83. Academic Press, 1991.

[147] G. Ward. A contrast-based scalefactor for luminance display. *Graphics Gems IV*, pages 415–421, 1994.

[148] G. Ward and M. Simmons. Subband encoding of high dynamic range imagery. In *APGV '04: 1st Symposium on Applied Perception in Graphics and Visualization*, pages 83–90, 2004.

[149] G. Ward and M. Simmons. JPEG-HDR: A backwards-compatible, high dynamic range extension to jpeg. In *Proceedings of the 13th Color Imaging Conference*, pages 283–290, 2005.

[150] G. J. Ward. The radiance lighting simulation and rendering system. In *Proceedings of SIGGRAPH 94*, Computer Graphics Proceedings, Annual Conference Series, pages 459–472, Orlando, Florida, July 1994. ACM SIGGRAPH / ACM Press. ISBN 0-89791-667-0.

[151] G. Ward Larson. LogLuv encoding for full-gamut, high-dynamic range images. *Journal of Graphics Tools*, 3(1):815–30, 1998.

[152] G. Ward Larson, H. Rushmeier, and C. Piatko. A visibility matching tone reproduction operator for high dynamic range scenes. *IEEE Transactions on Visualization and Computer Graphics*, 3(4):291–306, 1997.

[153] A. Watson. The cortex transform: Rapid computation of simulated neural images. *Comp. Vis. Graph. and Image Proc.*, 39:311–327, 1987.

[154] A. B. Watson and J. A. Solomon. A model of visual contrast gain control and pattern masking. *Journal of the Optical Society A*, 14:2378–2390, 1997.

[155] A. B. Watson, J. A. Solomon, A. Ahumada, and A. Gale. DCT basis function visibility: Effects of viewing distances and contrast masking. In *Human Vision, Visual Processing, and Digital Display V*, SPIE, volume 2179, pages 99–108, 1994.

[156] B. Watson, A. Friedman, and A. McGaffey. Measuring and predicting visual fidelity. In *Proceedings of ACM SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, pages 213–220, Los Angeles CA, USA, 2001. ACM Press / ACM SIGGRAPH / Addison Wesley Longman.

[157] Z. Wenjun, S. Daly, and L. Shawmin. Visual optimization tools in JPEG 2000. In *IEEE Intern. Conf. on Image Processing*, pages 37–40, 2000.

[158] S. Westen, R. Lagendijk, and J. Biemond. Perceptual image quality based on a multiple channel HVS model. In *Proc. of IEEE Int'l Conference on Acoustics, Speech and Signal Processing*, pages 2351–2354, 1995.

[159] G. Westheimer. The oscilloscopic view: retinal illuminance and contrast of point and line targets. *Vision Research*, 25(8):1097–1103, 1985.

[160] P. Whittle. Increments and decrements: Luminance discrimination. *Vision Research*, 26(10):1677–1691, 1986.

[161] H. Wilson. A transducer function for threshold and suprathreshold human vision. *Biological Cybernetics*, 38:171–178, 1980.

[162] H. Wilson. Psychophysical models of spatial vision and hyperacuity. In D. Regan, editor, *Vision and Visual Dysfunction: Spatial Vision*, pages 64–86. Pan Macmillan, 1991.

[163] J. Woods and G. Lilienfield. A resolution and frame-rate scalable subband/wavelet video coder. *IEEE Transactions on Circuits and Systems for Video Technology*, 11(9):1035–1044, 2001.

[164] G. Wyszecki and W. Stiles. *Color Science*. John Willey & Sons, 2000.

[165] R. Xu, S. Pattanaik, and C. Hughes. High-dynamic range still-image encoding in JPEG 2000. *IEEE Comp. Graph. and Appl.*, 26(6):57–64, 2005.

[166] H. Yee and S. Pattanaik. Segmentation and adaptive assimilation for detail-preserving display of high-dynamic range images. *The Visual Computer*, 19:457–466, 2003.

[167] C. Zetzsche and G. Hauske. Multiple channel model for the prediction of subjective image quality. In *Human Vision, Visual Processing, and Digital Display*, SPIE, volume 1077, pages 209–216, 1989.

Below you find a list of the most recent technical reports of the Max-Planck-Institut für Informatik. They are available by anonymous ftp from `ftp.mpi-sb.mpg.de` under the directory `pub/papers/reports`. Most of the reports are also accessible via WWW using the URL `http://www.mpi-sb.mpg.de`. If you have any questions concerning ftp or WWW access, please contact `reports@mpi-sb.mpg.de`. Paper copies (which are not necessarily free of charge) can be ordered either by regular mail or by e-mail at the address below.

Max-Planck-Institut für Informatik
Library
attn. Anja Becker
Stuhlsatzenhausweg 85
66123 Saarbrücken
GERMANY
e-mail: `library@mpi-sb.mpg.de`

MPI-I-2006-5-004	F. Suchanek, G. Ifrim, G. Weikum	Combining Linguistic and Statistical Analysis to Extract Relations from Web Documents
MPI-I-2006-5-003	V. Scholz, M. Magnor	Garment Texture Editing in Monocular Video Sequences based on Color-Coded Printing Patterns
MPI-I-2006-5-002	H. Bast, D. Majumdar, R. Schenkel, M. Theobald, G. Weikum	IO-Top-k: Index-access Optimized Top-k Query Processing
MPI-I-2006-5-001	M. Bender, S. Michel, G. Weikum, P. Triantafilou	Overlap-Aware Global df Estimation in Distributed Information Retrieval Systems
MPI-I-2006-4-001	R. Mantiuk	?
MPI-I-2006-2-001	T. Wies, V. Kuncak, K. Zee, A. Podelski, M. Rinard	On Verifying Complex Properties using Symbolic Shape Analysis
MPI-I-2006-1-001	M. Dimitrios	?
MPI-I-2005-5-002	S. Siersdorfer, G. Weikum	Automated Retraining Methods for Document Classification and their Parameter Tuning
MPI-I-2005-4-006	C. Fuchs, M. Goesele, T. Chen, H. Seidel	An Emperical Model for Heterogeneous Translucent Objects
MPI-I-2005-4-005	G. Krawczyk, M. Goesele, H. Seidel	Photometric Calibration of High Dynamic Range Cameras
MPI-I-2005-4-004	C. Theobalt, N. Ahmed, E. De Aguiar, G. Ziegler, H. Lensch, M.A.,. Magnor, H. Seidel	Joint Motion and Reflectance Capture for Creating Relightable 3D Videos
MPI-I-2005-4-003	T. Langer, A.G. Belyaev, H. Seidel	Analysis and Design of Discrete Normals and Curvatures
MPI-I-2005-4-002	O. Schall, A. Belyaev, H. Seidel	Sparse Meshing of Uncertain and Noisy Surface Scattered Data
MPI-I-2005-4-001	M. Fuchs, V. Blanz, H. Lensch, H. Seidel	Reflectance from Images: A Model-Based Approach for Human Faces
MPI-I-2005-2-004	Y. Kazakov	A Framework of Refutational Theorem Proving for Saturation-Based Decision Procedures
MPI-I-2005-2-003	H.d. Nivelle	Using Resolution as a Decision Procedure
MPI-I-2005-2-002	P. Maier, W. Charatonik, L. Georgieva	Bounded Model Checking of Pointer Programs
MPI-I-2005-2-001	J. Hoffmann, C. Gomes, B. Selman	Bottleneck Behavior in CNF Formulas
MPI-I-2005-1-008	C. Gotsman, K. Kaligosi, K. Mehlhorn, D. Michail, E. Pyrga	Cycle Bases of Graphs and Sampled Manifolds
MPI-I-2005-1-008	D. Michail	?
MPI-I-2005-1-007	I. Katriel, M. Kutz	A Faster Algorithm for Computing a Longest Common Increasing Subsequence
MPI-I-2005-1-003	S. Baswana, K. Telikepalli	Improved Algorithms for All-Pairs Approximate Shortest Paths in Weighted Graphs
MPI-I-2005-1-002	I. Katriel, M. Kutz, M. Skutella	Reachability Substitutes for Planar Digraphs
MPI-I-2005-1-001	D. Michail	Rank-Maximal through Maximum Weight Matchings

MPI-I-2004-NWG3-001	M. Magnor	Axisymmetric Reconstruction and 3D Visualization of Bipolar Planetary Nebulae
MPI-I-2004-NWG1-001	B. Blanchet	Automatic Proof of Strong Secrecy for Security Protocols
MPI-I-2004-5-001	S. Siersdorfer, S. Sizov, G. Weikum	Goal-oriented Methods and Meta Methods for Document Classification and their Parameter Tuning
MPI-I-2004-4-006	K. Dmitriev, V. Havran, H. Seidel	Faster Ray Tracing with SIMD Shaft Culling
MPI-I-2004-4-005	I.P. Ivrissimtzis, W.-. Jeong, S. Lee, Y.a. Lee, H.-. Seidel	Neural Meshes: Surface Reconstruction with a Learning Algorithm
MPI-I-2004-4-004	R. Zayer, C. Rssl, H. Seidel	r-Adaptive Parameterization of Surfaces
MPI-I-2004-4-003	Y. Ohtake, A. Belyaev, H. Seidel	3D Scattered Data Interpolation and Approximation with Multilevel Compactly Supported RBFs
MPI-I-2004-4-002	Y. Ohtake, A. Belyaev, H. Seidel	Quadric-Based Mesh Reconstruction from Scattered Data
MPI-I-2004-4-001	J. Haber, C. Schmitt, M. Koster, H. Seidel	Modeling Hair using a Wisp Hair Model
MPI-I-2004-2-007	S. Wagner	Summaries for While Programs with Recursion
MPI-I-2004-2-002	P. Maier	Intuitionistic LTL and a New Characterization of Safety and Liveness
MPI-I-2004-2-001	H. de Nivelle, Y. Kazakov	Resolution Decision Procedures for the Guarded Fragment with Transitive Guards
MPI-I-2004-1-006	L.S. Chandran, N. Sivadasan	On the Hadwiger's Conjecture for Graph Products
MPI-I-2004-1-005	S. Schmitt, L. Fousse	A comparison of polynomial evaluation schemes
MPI-I-2004-1-004	N. Sivadasan, P. Sanders, M. Skutella	Online Scheduling with Bounded Migration
MPI-I-2004-1-003	I. Katriel	On Algorithms for Online Topological Ordering and Sorting
MPI-I-2004-1-002	P. Sanders, S. Pettie	A Simpler Linear Time 2/3 - epsilon Approximation for Maximum Weight Matching
MPI-I-2004-1-001	N. Beldiceanu, I. Katriel, S. Thiel	Filtering algorithms for the Same and UsedBy constraints
MPI-I-2003-NWG2-002	F. Eisenbrand	Fast integer programming in fixed dimension
MPI-I-2003-NWG2-001	L.S. Chandran, C.R. Subramanian	Girth and Treewidth
MPI-I-2003-4-009	N. Zakaria	FaceSketch: An Interface for Sketching and Coloring Cartoon Faces
MPI-I-2003-4-008	C. Roessl, I. Ivrissimtzis, H. Seidel	Tree-based triangle mesh connectivity encoding
MPI-I-2003-4-007	I. Ivrissimtzis, W. Jeong, H. Seidel	Neural Meshes: Statistical Learning Methods in Surface Reconstruction
MPI-I-2003-4-006	C. Roessl, F. Zeilfelder, G. Nrnberger, H. Seidel	Visualization of Volume Data with Quadratic Super Splines
MPI-I-2003-4-005	T. Hangelbroek, G. Nrnberger, C. Roessl, H.S. Seidel, F. Zeilfelder	The Dimension of C^1 Splines of Arbitrary Degree on a Tetrahedral Partition
MPI-I-2003-4-004	P. Bekaert, P. Slusallek, R. Cools, V. Havran, H. Seidel	A custom designed density estimation method for light transport
MPI-I-2003-4-003	R. Zayer, C. Roessl, H. Seidel	Convex Boundary Angle Based Flattening
MPI-I-2003-4-002	C. Theobalt, M. Li, M. Magnor, H. Seidel	A Flexible and Versatile Studio for Synchronized Multi-view Video Recording
MPI-I-2003-4-001	M. Tarini, H.P.A. Lensch, M. Goesele, H. Seidel	3D Acquisition of Mirroring Objects
MPI-I-2003-2-004	A. Podelski, A. Rybalchenko	Software Model Checking of Liveness Properties via Transition Invariants
MPI-I-2003-2-003	Y. Kazakov, H. de Nivelle	Subsumption of concepts in $DL\ \mathcal{FL}_0$ for (cyclic) terminologies with respect to descriptive semantics is PSPACE-complete