Optimizing a Tone Curve for Backward-Compatible High Dynamic Range Image and Video Compression

Zicong Mai, *Student Member, IEEE*, Hassan Mansour, *Member, IEEE*, Rafal Mantiuk, Panos Nasiopoulos, *Member, IEEE*, Rabab Ward, *Fellow, IEEE*, and Wolfgang Heidrich

***Abstract*—For backward compatible high dynamic range (HDR) video compression, the HDR sequence is reconstructed by inverse tone-mapping a compressed low dynamic range (LDR) version of the original HDR content. In this paper, we show that the appropriate choice of a tone-mapping operator (TMO) can significantly improve the reconstructed HDR quality. We develop a statistical model that approximates the distortion resulting from the combined processes of tone-mapping and compression. Using this model, we formulate a numerical optimization problem to find the tone-curve that minimizes the expected mean square error (MSE) in the reconstructed HDR sequence. We also develop a simplified model that reduces the computational complexity of the optimization problem to a closed-form solution. Performance evaluations show that the proposed methods provide superior performance in terms of HDR MSE and SSIM compared to existing tone-mapping schemes. It is also shown that the LDR image quality resulting from the proposed methods matches that produced by perceptually-based TMOs.**

***Index Terms*—high dynamic range imaging, bit-depth scalable, tone-mapping, HDR video compression.**

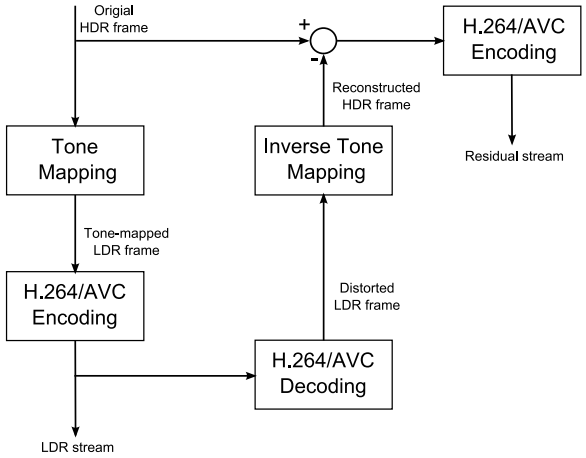

Fig. 1. General structure of the scalable approach used for backward-compatible HDR video encoding. The base layer encodes an 8-bit LDR representation of the HDR input. The enhancement layer encodes the difference (residual) between the inverse tone-mapped base layer and the original HDR source.

I. Introduction

Natural scenes contain far more visible information than can be captured by the majority of digital imagery and video devices. This is because traditional display devices can only support a limited dynamic range (contrast) and color gamut. New display and projection technologies, however, employ narrow-wavelength LED light sources that expand the boundaries of the displayable color gamut. This expansion will be again vastly enlarged with the next generation display technologies that will employ dual modulation [1] or backlight dimming that enhance intra- and inter-frame contrast.

For video compression, these advances in display technologies have motivated the use of extended gamut color spaces. These include xvYCC (x.v.Color) for home theater [2] and the Digital Cinema Initiative color space for digital theater applications. Yet, even these extended color spaces are too limited for the amount of contrast that can be perceived by the human eye. High dynamic range (HDR) video encoding goes beyond the typical color space restrictions and attempts to encode all colors that are visible and distinguishable to the human eye [3], and is not restricted by the color gamut of the display technology used. The main motivation is to create a video format that would be future-proof, independent of a display technology, and limited only by the performance of the human visual system (HVS).

HDR images preserve colorimetric or photometric pixel values (such as CIE XYZ) within the visible color gamut and allows for intra-frame contrast exceeding 5-6 orders of magnitude ($10^6 : 1$), without introducing contouring, banding or posterization artifacts caused by excessive quantization. The photometric or colorimetric values, such as luminance ($\mathrm{cd} \cdot \mathrm{m}^{-2}$) or spectral radiance ($\mathrm{W{\cdot}sr^{-1}{\cdot}m^{-3}}$), span much larger range of values than luma and chroma values (gamma corrected) used in typical video encoding (JPEG, MPEG, etc.). The obvious representation for the colorimetric values are floating point numbers, which, however, are impractical for image and video coding applications. For that reason several HDR color encodings and file formats have been proposed, including the Radiance RGBE (.hdr) [4], OpenEXR (.exr) [5] and LogLuv TIFF (.tiff) [6] file formats. They employ either more efficient floating point coding (OpenEXR and RGBE) or perceptually motivated compressive functions (LogLuv and [3]), which extends a typical 'gamma correction' to the entire range of luminance values.

Although HDR image and video encoding offers truly device-independent representation, the majority of existing digital display devices can only support 8-bit video content.

Zicong Mai, Hassan Mansour, Panos Nasiopoulos, and Rabab Ward are with the Department of Electrical and Computer Engineering, University of British Columbia, Vancouver, BC, Canada (email: zicongm@ece.ubc.ca, hassanm@ece.ubc.ca, panosn@ece.ubc.ca, and rababw@ece.ubc.ca).

Rafal Mantiuk is with the School of Computer Science, Bangor University, North Wales, United Kingdom (email: mantiuk@bangor.ac.uk).

Wolfgang Heidrich is with the Department of Computer Science, University of British Columbia, Vancouver, BC, Canada (email: heidrich@cs.ubc.ca).

Therefore, high dynamic range video formats are unlikely to be broadly accepted without the backward-compatibility with these devices. Such backward-compatibility can be achieved if the HDR video stream contains 1) a backward-compatible 8-bit video layer which could be directly displayed on existing devices, and 2) additional information which along with this 8-bit layer can yield a good quality reconstructed version of the original HDR content. Such a stream can also contain a residual layer to further improve the quality of the HDR reconstruction. Fig. 1 illustrates the general coding structure used to provide a backward compatible HDR video bitstream.

Several proposals have been suggested to allow the above-mentioned HDR backward-compatibility function within the scalable extension of the H.264/AVC video coding standard [7]–[15]. The contrast of the original HDR content is first quantized into the 8-bit range using a tone mapping operator (TMO) to produce an LDR representation. The LDR sequence is then compressed using a standard video encoder (H.264/AVC). A larger dynamic range video can then be reconstructed by decoding the LDR layer and applying the inverse of the tone-mapping operator to reconstruct the HDR representation. The shape of the TMO can be encoded using supplemental enhancement information (SEI) messages [8], [16]. Finally, a HDR residual signal can also be extracted and encoded in the bitstream as an enhancement layer.

In this paper, we address the problem of finding an optimal tone-curve for such a backward-compatible encoding scheme. To compute the tone-curve, we propose a method that minimizes the difference in the video quality between the original and the reconstructed HDR video . This difference results in quality loss and is due to tone-mapping, encoding, decoding and inverse tone-mapping the original video. Minimizing this difference would reduce the size of the HDR residual signal in the enhancement layer. We also achieve the primary goal of tone-mapping which is to produce an LDR image with a visual response as similar as possible to the original HDR image. Although the initial assumptions used in our approach pose a difficult optimization problem, we demonstrate that for typical compression distortions there exists a closed-form solution that approaches the optimum.

The remainder of this paper is organized as follows: an overview of related work is presented in Section II. In Section III, the proposed tone-mapping approach that considers tone-mapping together with compression is discussed in detail. Section IV demonstrates and analyzes the performance of the proposed methods. Finally, we draw our conclusions in Section VI.

II. Related work

Backward compatible HDR video encoding has received significant interest recently. Mantiuk et al. [3] derived a color space of encoding HDR content based on the luminance threshold sensitivity of the human visual system. They concluded that 10-12 bit luma encoding is sufficient to encode the full range of visible and physically plausible luminance levels. Their encoding, however, is not backward-compatible with the existing video decoding hardware. For still image compression, backward compatibility can be achieved by encoding a tone-mapped copy of the HDR image together with a residual [17] or a ratio image [18] that allows the reconstruction of the original HDR image. In [19] and [20] this approach was extended for video sequences. A tone-mapping curve was encoded together with the tone-mapped and residual video sequences. In [19] the residual video sequence was additionally filtered to remove the information that is not visible to the human eye.

Recently, several proposals for bit-depth scalability have been introduced to provide backward-compatible HDR video bitstreams. These proposals incorporate backward-compatible encoding of high fidelity video as an extension to the H.264/AVC video encoding standard [8]–[10], [21]. The extension includes tone-mapping SEI messages, which encode the shape of the tone-mapping curve [8]. The scalable video coding (SVC) extension [22] is used to encode an additional residual stream needed to reconstruct the information lost due to tone-mapping, or to provide bit-depth scalability [9], [10], [20], [23]. The proposed tone mapping is meant to be used in combination with the scalable video coding. In that case the backward compatible (tone-mapped) sequence is generated by our operator, instead of being provided by a user, which results in better compression efficiency.

The primary goal of tone-mapping is to produce the best quality low-dynamic range rendering of an HDR scene that is visually close to the visual high contrast signal [24], [25]. A review of such operators can be found in [26, Ch. 6–8]. Li et al. [27] first considered tone-mapping explicitly so as to optimize image compression. They used forward and inverse wavelet-based tone-mapping (compressing and companding) in an iterative optimization loop to minimize HDR quality loss due to quantizing the 8-bit tone-mapped image. As this method requires encoding the tone-mapped image using high bit-rates, it is thus not suitable for video. Lee et al. [28] extended the gradient domain tone-mapping method [29] to video applications using the temporal information obtained from the video decoding process. In [30] the performance of several tone-mapping operators in terms of quality loss due to forward and inverse tone-mapping was compared. Local tone-mapping operators (spatially variant) were found to be more prone to quality loss than global operators (spatially invariant). In this paper we compare the results of our study with the two tone-mapping operators that performed the best in this study: the photographic TMO [31] and the adaptive logarithmic TMO [32].

III. Problem Statement and Proposed Solution

In this section, we present the challenges of obtaining a good quality reconstructed HDR representation in a backward-compatible HDR video encoding system and describe in detail the approach we propose towards overcoming these challenges.

The performance of a backward-compatible HDR video and image encoding system depends on the coding efficiency of the LDR base layer and the HDR enhancement layer. Performance gains can be achieved by finding a TMO that preserves the

necessary information in the LDR base layer so that after it passes through the inverse TMO process, the resultant HDR reconstructed signal is of high quality. The coding efficiency of the base layer does not depend much on the TMO used as most TMOs attain a similar level of contrast in the LDR representation. Therefore, the performance gain lies in the effectiveness of the inverse TMO in producing a high quality inverse tone-mapped HDR representation. This in turn determines the resulting HDR quality (when no enhancement layer exists).

It can be deduced from above that the performance of the whole system strongly depends on the TMO used to produce the LDR representation which, in turn undergoes compression. Our proposed approach attempts to find the best global (spatially invariant) tone-mapping curve that minimizes the mean square error (MSE)[1] between the original HDR content and the reconstructed version obtained after tone-mapping, compression, decompression, and inverse tone-mapping . This process is illustrated in Fig. 2.

Let l denote the input HDR image/frame, and v the tone-mapped LDR version as shown in Fig. 2a. Let $\tilde{v}$ be the decoded LDR frame, and $\tilde{l}$ the reconstructed HDR frame produced after inverse-tone-mapping. Also let θ be the set of parameters that control the tone-mapping operator, then our goal is to find the tone-mapping parameters that minimize MSE, which we denote as $||l - \tilde{l}||_2^2$ using the norm notation.

The above optimization problem can be solved by exhaustive search, repeatedly tone-mapping, encoding, decoding and then inverse tone-mapping, until the best set of TMO parameters θ^* is found. Even though this approach guarantees an optimal solution, this framework requires an unacceptable computational cost. To overcome this problem, we estimate the distortion due to tone-mapping, encoding, decoding, and inverse tone-mapping with a statistical distortion model, as illustrated in Fig. 2b. Then, we show that under certain assumptions that are valid for natural images, an immediate closed-form solution for this problem can be found.

In the following sections we consider only luminance/luma channels. To tone-map color images, we use the same tone-curve for the red, green and blue color channels. Such approach was shown to well preserve color appearance for moderate contrast compression [34]. Encoding of the enhancement layer (for the residual data) is not considered in this paper. The rationale comes from our effort to achieve the best possible HDR reconstruction and thus the smallest possible residual. As a result, the cost of encoding any additional refinement layer would be minimized.

In the following subsections we describe how we parameterize the tone-mapping function, approximate encoding distortions with a statistical model and then find a closed-form solution for an optimal tone-curve.

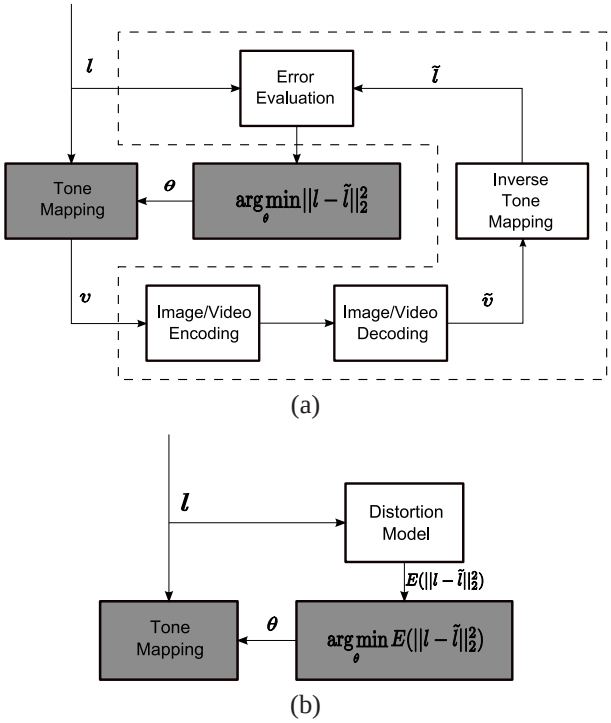

Fig. 2. System overview of the proposed tone-mapping method. (a) demonstrates the ideal scenario where the actual H.264/AVC encoding is employed. (b) shows the practical scenario which is addressed by this paper.

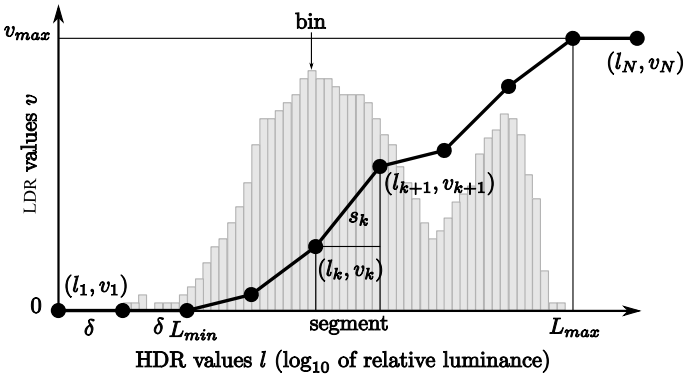

Fig. 3. Parameterization of a tone-mapping curve and the notation. The bar-plot in the background represents an image histogram used to compute $p(l)$.

A. Tone-Mapping Curve

The global tone-mapping curve is a function that maps HDR luminance values to either the display's luminance range [25], or directly to LDR pixel values. In this paper, we consider the latter case. The tone-mapping curve is usually continuous and non-decreasing. The two most common shapes for the tone curves are the sigmoidal ("S-shaped") or a compressive power function with an exponent < 1 (gamma correction).

According to the Weber-Fechner law [35], the sensitivity

[1]We choose the mean square error as the HDR quality metric for its simplicity, despite its shortcomings in reflecting the perceptual quality of images. Moreover, the results shown in Section IV demonstrate that although we minimize the MSE we also achieve image quality gains in terms of SSIM [33].

of the human visual system to light is proportional to the logarithm of luminance. Thus, our tone-mapping method will operate on the logarithmic values of the luminance, which we refer to as HDR values ($l = \log_{10}(L)$ where L is the luminance of the HDR image).

To keep the problem analytically tractable, we parameterize the tone-mapping curve as a piece-wise linear function with the nodes (l_k, v_k), as shown in Fig. 3. Each segment k between two nodes (l_k, v_k) and (l_{k+1}, v_{k+1}) has a constant width in HDR values equal to δ (0.1 in our implementation). The tone-mapping curve can then be uniquely specified by a set of slopes:

$$s_k = \frac{v_{k+1} - v_k}{\delta}, \tag{1}$$

which forms a vector of tone-mapping parameters θ. Using this parameterization, the forward tone-mapping function is defined as:

$$v(l) = (l - l_k) \cdot s_k + v_k, \tag{2}$$

where v is the LDR pixel value, k is the segment corresponding to HDR value l, that is $l_k \leq l < l_{k+1}$. The inverse mapping function is then:

$$\tilde{l}(v; s_k) = \begin{cases} \dfrac{v - v_k}{s_k} + l_k & \text{for} \quad s_k > 0 \\ \displaystyle\sum_{l \in S_0} l \cdot p_L(l) & \text{for} \quad s_k = 0\ , \end{cases} \tag{3}$$

where $s_k \in \{s_1 .. s_N\}$.

When the slope is zero ($s_k = 0$), $\tilde{l}(v; s_k)$ is assigned an expected HDR pixel value for the entire range S_0 in which the slope is equal zero. $p_L(l)$ is the probability of HDR pixel value l.

B. Statistical Distortion Model

As mentioned earlier in section III, accurately computing the distorted HDR values $\tilde{l}$ would be too computationally demanding. Instead, we estimate the error $||\tilde{l} - l||_2^2$ assuming that the compression distortions follow a known probability distribution p_C. Under this assumption, the expected value of the error $||\tilde{l} - l||_2^2$ is:

$$E[||\tilde{l} - l||_2^2] = \sum_{l=l_{min}}^{l_{max}} \sum_{\tilde{v}=0}^{v_{max}} (\tilde{l}(\tilde{v}; s_k) - l)^2 \cdot p_C(v(l) - \tilde{v} | v(l)) \cdot p_L(l), \tag{4}$$

where $p_C(v - \tilde{v}|v)$ is the probability that the encoding error equals $v - \tilde{v}$. Note that Eqs. (2) and (3) show that both v and $\tilde{l}$ are uniquely determined by the values of l and $\tilde{v}$, respectively. Therefore, the conditional probabilities for these two variables and their corresponding summations have been removed from the calculation of the expected value of the error above.

The probability of the HDR pixel value $p_L(l)$ is in practice found from a histogram of HDR values and the summation over l is performed for each bin of that histogram. The number of bins is greater than or equal to the number of tone-curve segments.

Since a tone curve is uniquely defined by a sequence of its slopes $s_1, ..., s_N$, the expected error value in (4) can be expressed as a function $\varepsilon(s_k)$. For a specific tone curve defined by the sequence of slopes, the pixel value v is calculated as in (2). In practice the pixel values v and $\tilde{v}$ are integer valued, such that $v, \tilde{v} \in \{0, 1, ..., v_{max}\}$, whereas, l and $\tilde{l}$ are continuous real variables. The rounding operation makes the encoding error estimate in (4) a non-convex function. Therefore, we impose a convex relaxation on the encoding error function by removing the rounding operator from the calculation of v. Moreover, assuming that the compression error probability is independent of the LDR pixel value v, we can simplify the expression above by removing the dependency of p_C on v. Consequently, the continuously relaxed objective function is written as:

$$\varepsilon(s_k) = \sum_{l=l_{min}}^{l_{max}} \sum_{\tilde{v}=0}^{v_{max}} (\tilde{l}(\tilde{v}; s_k) - l)^2 \cdot p_C(v - \tilde{v}) \cdot p_L(l) \tag{5}$$

The only unknown variable is the probability distribution of the compression error $p_C(v - \tilde{v})$, which can be estimated for any lossy compression scheme. In Appendix A we model such distribution for the H.264/AVC I-frame coding [2]. However, we will show in Section III-D that the distribution of the compression scheme error is not necessary to calculate a good approximation of the encoding error.

C. Optimization Problem

The optimum tone curve can be found by minimizing the function $\varepsilon(s_k)$ with respect to the segment slopes s_k:

$$\arg\min_{s_1 .. s_N} \varepsilon(s_k) \tag{6}$$

subject to:

$$\begin{aligned} & s_{min} \leq s_k \leq s_{max} \quad \text{for} \quad k = 1..N \\ & \sum_{k=1}^{N} s_k \cdot \delta = v_{max}. \end{aligned} \tag{7}$$

The first constraint restricts slopes to the allowable range, while the second ensures that the tone curve spans exactly the range of pixel values from 0 to v_{max}. The minimum slope s_{min} ensures that the tone-mapping function is strictly increasing and thus invertible and $\tilde{l}(\tilde{v}; s_k)$ can be computed. The lack of this assumption introduces discontinuity and local minima, impeding the use of efficient solvers. Since s_{min} is set to a very low value (below $0.5/\delta$), this assumption has no significant effect on the resulting tone-curves, which are rounded to the nearest pixel values. With s_{max} we ensure that we do not try to preserve more information than what is visible to the human eye. Assuming that the luminance detection threshold equals 1% ($\Delta L / L = 0.01$), we can write:

$$\tilde{l}(v + 1; s_k) - \tilde{l}(v; s_k) > \log_{10}(1.01), \tag{8}$$

so that:

$$s_{max} = (\log_{10}(1.01))^{-1}. \tag{9}$$

[2]The derivation is also valid for other compression methods such as H.264/AVC P-frame, B-frame and JPEG coding, which we show in the supplementary materials.

D. Closed-form Solution

The distortion model in (5) gives a good estimate of compression errors, but poses two problems for practical implementation in an HDR compression scheme: 1) it requires the knowledge of the encoding distortion distribution p_C, and 2) the optimization problem can only be solved numerically using slow iterative solvers. In order to reduce the complexity of the optimization problem given in (7), we propose the following assumptions that allow us to cast a simpler optimization problem to which we can find a closed-form solution with almost no noticeable impact on the compression performance.

If we assume local linearity of the tone-curve, so that the slope at the non-distorted pixel value v and that at the distorted pixel value $\tilde{v}$ is the same, we can then substitute $\tilde{l}(\tilde{v}; s_k)$ and l in the distortion model (5) using the inverse mapping function in (3), this gives:

$$\varepsilon(s_k) \approx \sum_{l=l_{min}}^{l_{max}} \sum_{\tilde{v}=0}^{v_{max}} p_C(v-\tilde{v}) \cdot p_L(l) \cdot \left(\frac{v-\tilde{v}}{s_k}\right)^2. \quad (10)$$

After reorganizing we get:

$$\begin{aligned} \varepsilon(s_k) &\approx \sum_{l=l_{min}}^{l_{max}} \frac{p_L(l)}{s_k^2} \cdot \sum_{\tilde{v}=0}^{v_{max}} p_C(v-\tilde{v}) \cdot (v-\tilde{v})^2 \\ &= \sum_{l=l_{min}}^{l_{max}} \frac{p_L(l)}{s_k^2} \cdot Var(v-\tilde{v}). \end{aligned} \quad (11)$$

Since the variance of $(v-\tilde{v})$ does not depend on the slopes s_k, it does not affect the location of the global minimum of $\varepsilon(s_k)$ and thus can be omitted when searching for the minimum.

Our local linearity assumption holds in most cases for two reasons. Firstly, the distortion distribution p_C has high kurtosis (see appendix) so that most of the distorted pixels are likely to lie in the same segment as the non-distorted pixel v. Secondly, even if a distorted pixel $\tilde{v}$ moves to another segment, the slopes of two neighboring segments are usually very close to each other. This assumption has been also confirmed by our results, in which the tone-curves found using the accurate model from (5) and a simplified model from (11) were almost the same (see Section IV).

The most important consequence of using the simplified model from (11) is that the optimal tone-curve does not depend on the image compression error, as long as the compression distortions are not severe enough to invalidate the local linearity assumption. This means that the optimal tone-curve can be found independently of the compression algorithm and its quality settings.

The constrained optimization problem defined in (6) can now be re-written as follows:

$$\begin{aligned} &\arg\min_{s_1..s_N} \sum_{k=1}^{N} \frac{p_k}{s_k^2} \\ &\text{subject to} \quad \sum_{k=1}^{N} s_k = \frac{v_{max}}{\delta}, \end{aligned} \quad (12)$$

where $p_k = \sum_{l=l_k}^{l_{k+1}} p_L(l)$, and l_k and l_{k+1} define the lower and the upper bounds of a segment, respectively.

This problem can be solved analytically by calculating the first order Karush-Kuhn-Tucker (KKT) optimality conditions of the corresponding Lagrangian, which results in the following system of equations:

$$\begin{cases} \frac{-2p_1}{s_1^3} + \lambda = 0 \\ \frac{-2p_2}{s_2^3} + \lambda = 0 \\ \vdots \\ \frac{-2p_N}{s_N^3} + \lambda = 0 \\ \sum_{k=1}^{N} s_k - \frac{v_{max}}{\delta} = 0 \end{cases}, \quad (13)$$

where λ is the Lagrange multiplier. The solution to the above system of equations results in the slopes s_k given by:

$$s_k = \frac{v_{max} \cdot p_k^{1/3}}{\delta \cdot \sum_{k=1}^{N} p_k^{1/3}}. \quad (14)$$

Note that the expression derived in (14) does not consider the upper bound constraint imposed on s_k in (7). Let $\mathcal{I}$ be the set of the index of a segment with a slope that exceeds the upper bound. We overcome the upper bound violation using the following adjustment:

$$s_k = \begin{cases} s_{max} & \text{for} \quad s_k \in \mathcal{I}, \\ \frac{\left(v_{max} - \sum_{i \in \mathcal{I}} s_{max}\delta\right) \cdot p_k^{1/3}}{\delta \cdot \sum_{j \notin \mathcal{I}}^{N} p_j^{1/3}} & \text{for} \quad s_k \notin \mathcal{I}. \end{cases} \quad (15)$$

IV. Experimental Results and Discussion

In this section we first validate the proposed methods: optimization using the statistical model proposed in Section III-B and the closed-form solution based on a simplified model derived in Section III-D. Then, our models are further analyzed based on the generated tone curve and the distortion of the reconstructed HDR content. The performance of our models is also evaluated by comparing it with existing tone-mapping methods. We use H.264/AVC encoding as an example to demonstrate the results. In the experiments below, all tone-mapped images are compressed/decompressed using the intra mode of the H.264/AVC reference software [36] except for IV-E where inter-frame mode is also used. To reconstruct an HDR image from a decoded LDR image, an inverse tone-mapping function is stored as a lookup table with each encoded image.

A. Model validation

In this section, we validate that the statistical model of Section III-B results in a tone-curve that truly reflects the ground-truth results. Ground-truth results are achieved using the ideal scheme illustrated in Fig. 2a, where the actual H.264/AVC encoder and decoder are employed to find the truly optimal piecewise linear tone curve. This ideal scheme is extremely computationally expensive, and its complexity increases exponentially with the number of segments. To make

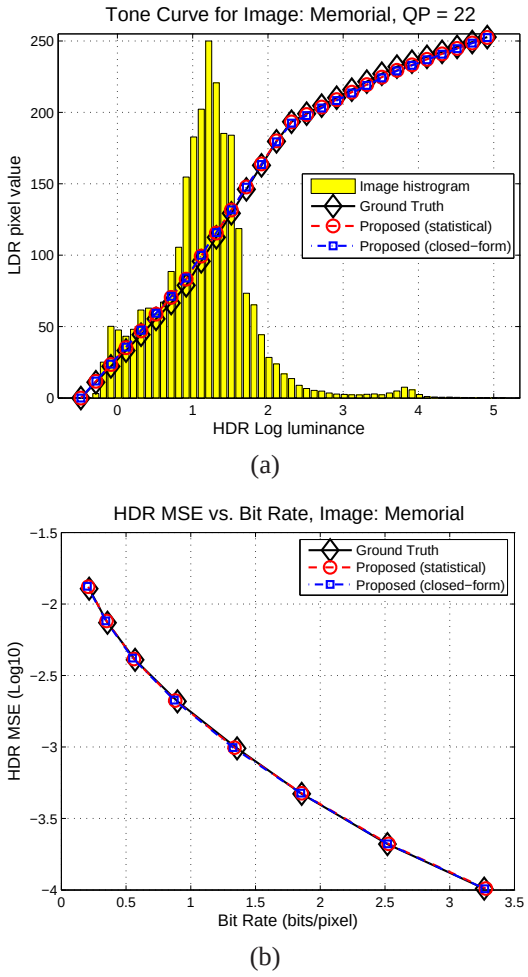

Fig. 4. Validation of the proposed models by comparison with the ground-truth solution. The top figure, (a), shows the tone curves computed using the statistical model, the closed-form solution and the ground-truth optimization for the image "Memorial". The x axis denotes the HDR luminance in the log-10 scale, and y axis is the LDR pixel value. (b) demonstrates the result of HDR MSE (in log10 scale) vs. bit rate (bits/pixel). The lower the MSE value, the better the image quality.

the experiment computationally feasible, we divided the tone curve into four segments of equal width. That is, the dynamic range of each of the segments is identical. Then the ideal scheme is used to find our ground-truth four-segment tone-curve.

Fig. 4a demonstrates the tone curves generated by the statistical (Section III-B), the closed-form (Section III-D) and the ground-truth approaches for the H.264/AVC quantization parameter QP = 22. It can be seen that the tone curves produced by the proposed models are very close to the ground-truth curve. Fig. 4b shows the rate-distortion result in terms of bit rate vs HDR MSE. The results show that for different encoding bit rates, the reconstructed HDR images resulting from the two proposed models have very similar MSE relative to the ground-truth case. This further validates that the performance of the ideal scenario can be closely estimated by our statistical model and that the local linearity assumption we used to derive the closed-form solution is justified.

B. Dependence of the tone curves on QP

Next, we verify that the proposed statistical model can be well approximated by the closed-form solution which produces a tone curve that is independent of QP. The probability distribution of the H.264/AVC compression errors, which is a function of QP, is included in the statistical model proposed in Section III-B. This suggests that the generated tone curve should vary with the value of QP. However, we observed from experiments on a large pool of HDR images encoded at different QPs that the variation in QP has no significant effect on the choice of an optimal tone curve. Fig. 5 illustrates this observation with an example of two images and their corresponding tone curves derived from the statistical model for different QP values (the larger the value of QP, the larger the compression error). The figures show that the tone curves are not significantly affected by the variation of QP.

C. Further analysis of the closed-form solution

The tone curve resulting from the closed-form solution given by (14) can be generalized as follows:

$$s_k = \frac{v_{max} \cdot p_k^{1/t}}{\delta \cdot \sum_{k=1}^{N} p_k^{1/t}}, \tag{16}$$

In our closed-form solution, t is set to be equal to 3. Note that when $t = 1$, (16) is identical to the histogram equalization operation. Therefore, we will investigate the performance of the tone curves obtained from changing the exponent t of (16).

In the experiment, we set t = 1, 2, 3, 4, 5, 10 and 20, and compressed the tone-mapped image using H.264/AVC at different QPs, and evaluated the distortion of the reconstructed HDR image. In addition to HDR MSE, we also used the popular quality metric known as the Structural Similarity Index (SSIM) [33] in order to find which t value gives highest quality for a particular quality metric. Fig. 6 shows the resulting average performance over 40 HDR images. The left row in the figure indicates that our closed-form solution ($t = 3$) is largely better than the histogram equalization method ($t = 1$) and outperforms all other cases for HDR MSE. This can be expected, since our approach explicitly minimizes MSE. However, the same behavior cannot be expected from SSIM: the difference among all cases is minimal for light and medium compression, while the case of ($t = 2$) performs slightly better for strong compression.

From a practical point of view, HDR content is usually prepared for high-quality visual experience where only light or medium compression quantization is allowed. In this sense, the results demonstrated in Fig. 6 indicate that our closed-form solution (t = 3) guarantees good performance.

D. Comparison with existing TMOs

In this subsection, we compare the performance of the proposed models to existing tone-mapping methods. The chosen TMOs are the photographic TMO [31], the adaptive logarithmic TMO [32] and the display adaptive TMO [25]. In [30], a study was conducted to find how different TMOs

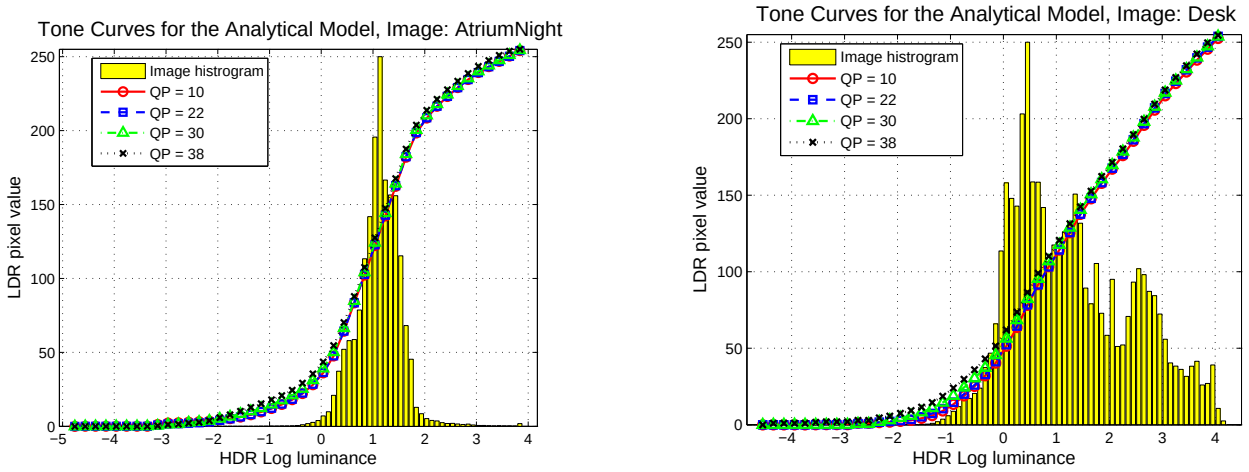

Fig. 5. Tone curves generated using the statistical model with different QP values for the images "AtriumNight" and "Desk". The notation of the axis is the same as Fig. 4a. The smaller the value of QP, the better the compression quality. 87 and 88 segments are used for "AtriumNight" and "Desk" respectively.

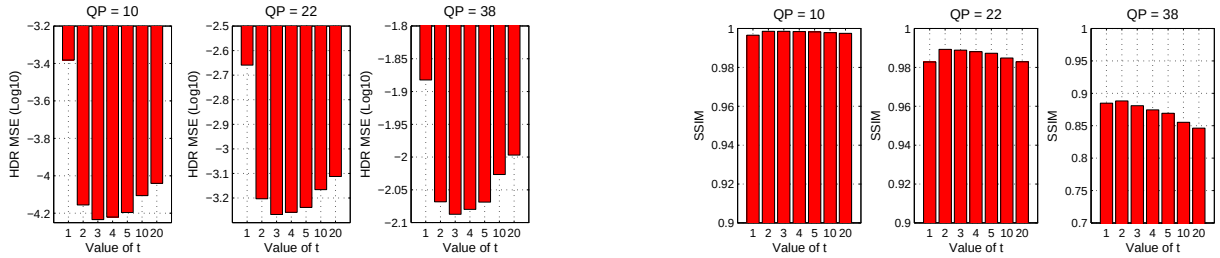

Fig. 6. Distortion measures for the reconstructed HDR images using the generalized solution (see (16)) with different values of t, averaged over 40 images. The tone-mapped images are compressed with different quality (QP = 10, 22 and 38), decoded and used to reconstruct HDR images. The left row, (a), shows the measurement of HDR MSE, where the smaller the value, the better the image quality. (b) compares the SSIM quality. Higher SSIM values mean better quality. For each of the 40 images, the segment width is set to be 0.1.

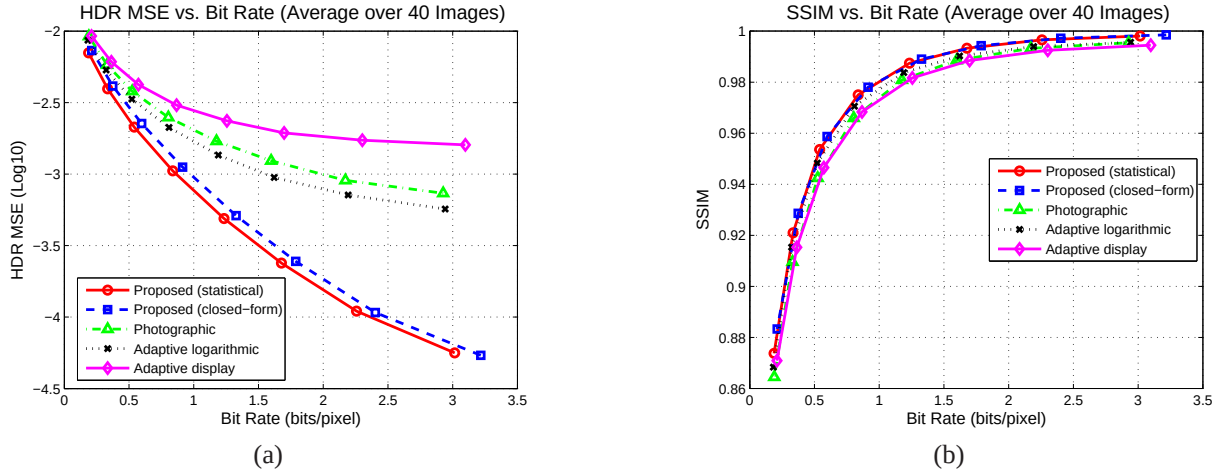

Fig. 7. Comparison with other tone-mapping methods in terms of MSE and SSIM (for the reconstructed HDR image) vs. bit rate, averaged over 40 images. The correlation between image quality and the distortion measures (MSE and SSIM) can be referred to the caption of Figure 6. MSE and SSIM for all methods represent the reconstruction error without the correction of the residual/enhancement layer. In the experiment, the segment width of each image histogram is set to be 0.1.

perform when the LDR is inverse tone-mapped. Of the TMOs that were compared, the photographic TMO and the adaptive logarithmic TMO were found to outperform other popular tone-mapping methods in backward-compatible HDR image

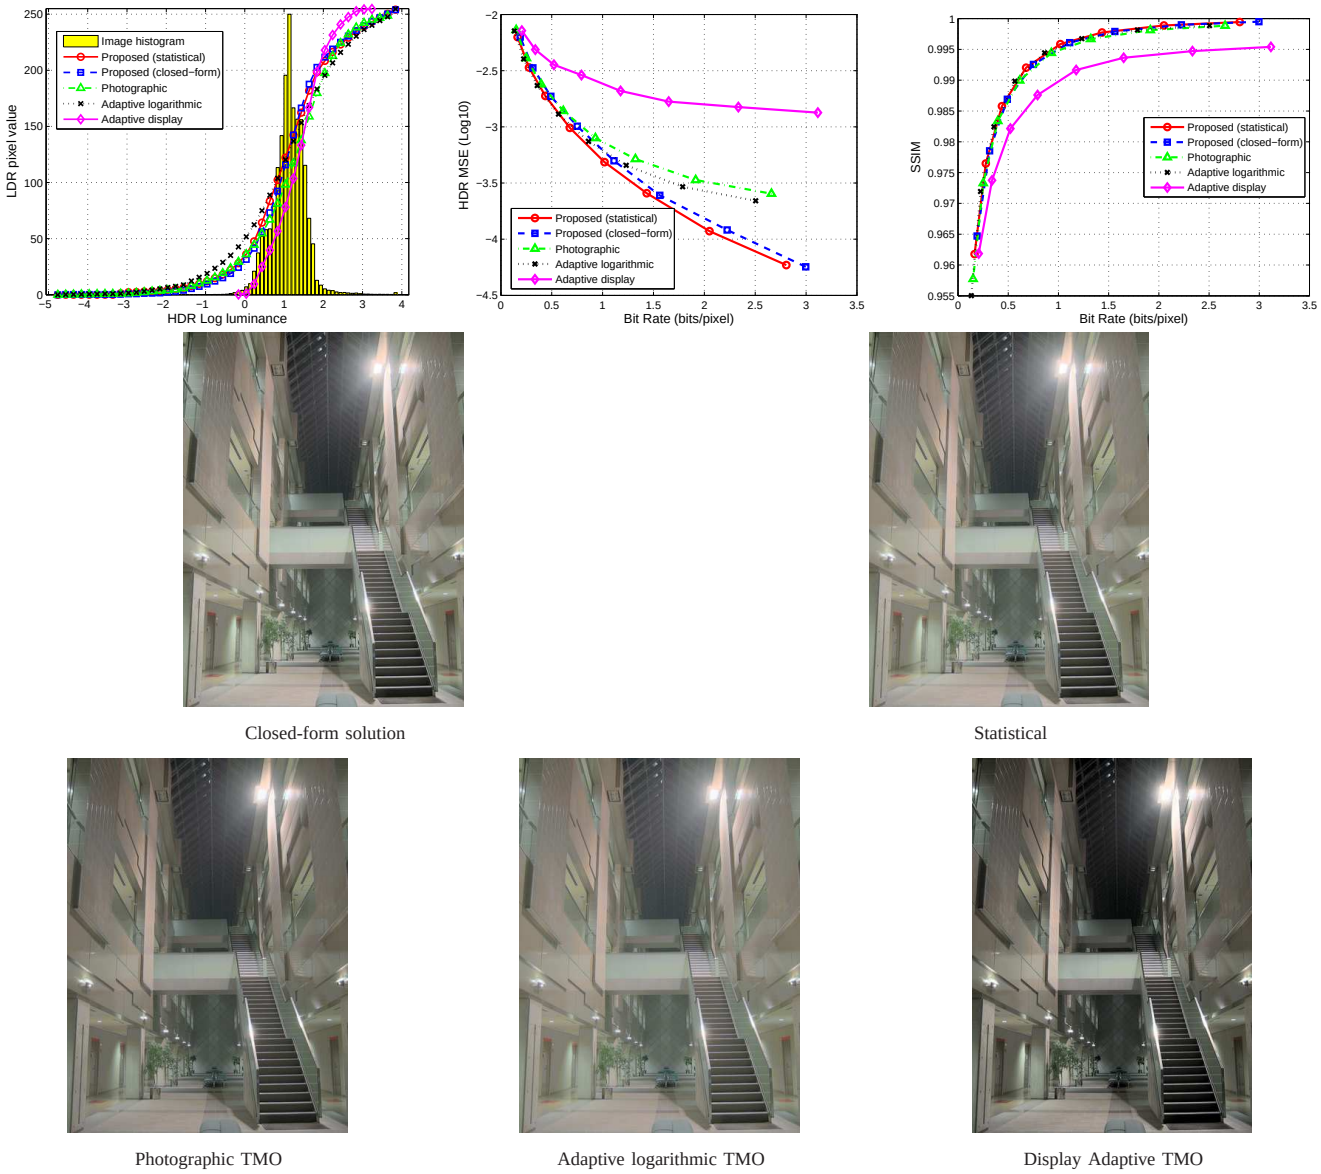

(a) Image: AtriumNight; LDR images shown are compressed with QP = 10

Fig. 8. Rate-distortion curves, tone curves and tone-mapped images for the image "AtriumNight". The first row demonstrates the resulting tone-curves with different TMOs, followed by the results for MSE and SSIM vs. bit rates; the second row shows tone-mapped LDR images using the proposed statistical model and the closed-form solution. The third row shows the tone-mapped images using the existing tone-mapping methods. All the tone-mapped images shown are compressed. The compression quantization parameters used for "AtriumNight" is 10. The number of segments used for the histogram is 87.

compression. However, Display adaptive TMO is a recent tone-mapping algorithm at the time of writing and it employs a similar optimization loop as our technique.

Fig. 7 compares the distortion of the reconstructed HDR image versus the compressed LDR bit rate for different TMOs, averaged over 40 images. This test demonstrates how successful each TMO is at delivering a good quality HDR by inverse tone mapping the corresponding LDR representation. The results show that our proposed methods clearly outperform the other methods in terms of MSE. The difference is not very large for low bit rates (heavy compression), but the performance of our model dramatically improves when the compression reaches the point of the medium and light compression. For HDR MSE = -3 (in log10 scale), which corresponds to QP = 25, we save about 50% of the bit-rate compared to the best performing competitive TMO for the same quality. Fig. 7 also shows that the proposed statistical model results in a better MSE performance compared to the closed-form solution.

Although our models are designed for minimizing MSE, the results also indicate that the proposed TMOs show superior performance for the advanced quality metric, SSIM. In terms of SSIM, Fig. 7b shows that our proposed TMOs result in a better performance and the improvement is sustained for higher bit rates.

Fig. 8, 9 and 10 display the tone curve, rate-distortion curves and tone-mapped LDR images for three images. Ad-

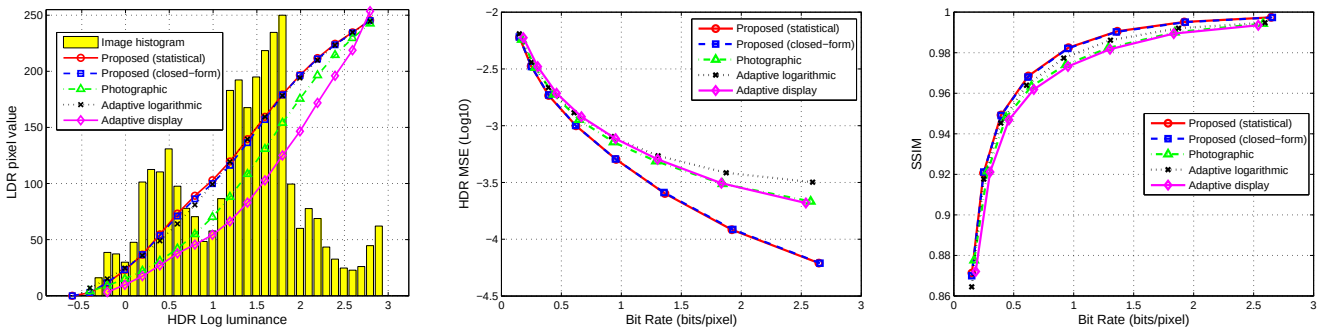

Closed-form solution

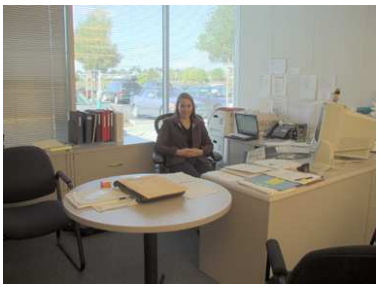

Statistical

Photographic TMO

Adaptive logarithmic TMO

Display Adaptive TMO

(a) Image: Coby; LDR images shown are compressed with QP = 10

Fig. 9. Rate-distortion curves, tone curves and tone-mapped images for the image "Coby". The notation is the same as Fig. 8. The compression quantization parameters used for "Coby" is 22. The number of segments used for the histogram is 36.

ditional results for more images are included in the supplementary material. The LDR images shown in these figures demonstrate that the images tone-mapped using our method also provide good quality. To further demonstrate the quality of the LDR images generated by the proposed models, Fig. 11 shows the distortion maps of the LDR images compared with their original HDR counterparts. The distortion maps were generated using the dynamic range independent image quality metric [37], which is the only available computational metric capable of comparing HDR and tone-mapped images. The metric visualizes the areas where the visible contrast is lost (green color), or distorted (red color). The distortions maps indicate that the proposed method causes less contrast loss than the other tested tone-mapping operators.

The computer graphics community often uses a perceptually-based image difference predictor, HDR-VDP [38], to compare a distorted HDR image to a reference HDR image. We compared the different TMOs in terms of HDR-VDP and found that there is no consistent improvement or degradation in performance of the different TMOs when compression is applied at the LDR layer. Therefore, we do not include performance evaluations relative to HDR VDP in this paper.

The computational complexity of most global tone mapping operators, such as the photographic and the logarithmic operators compared in our study, is linear to the number of pixels O(N). The same holds for our closed form solution, in which the most expensive part is computing an image histogram. The solution based on the statistical model is more computationally expensive because it requires several iteration for the optimization procedure to converge. For comparison, the closed-form solution requires about 0.9 seconds and the statistical approach 20 seconds to complete using non-optimized MATLAB code on a 3 GHz CPU computer.

E. Optimized tone-curves for JVT bit-depth scalable encoding

In this section, we demonstrate the efficiency gains that can be expected when the proposed tone-mapping technique is used in combination with the JVT bit-depth scalable extensions [22]. For that purpose, we tone-map the JVT stan-

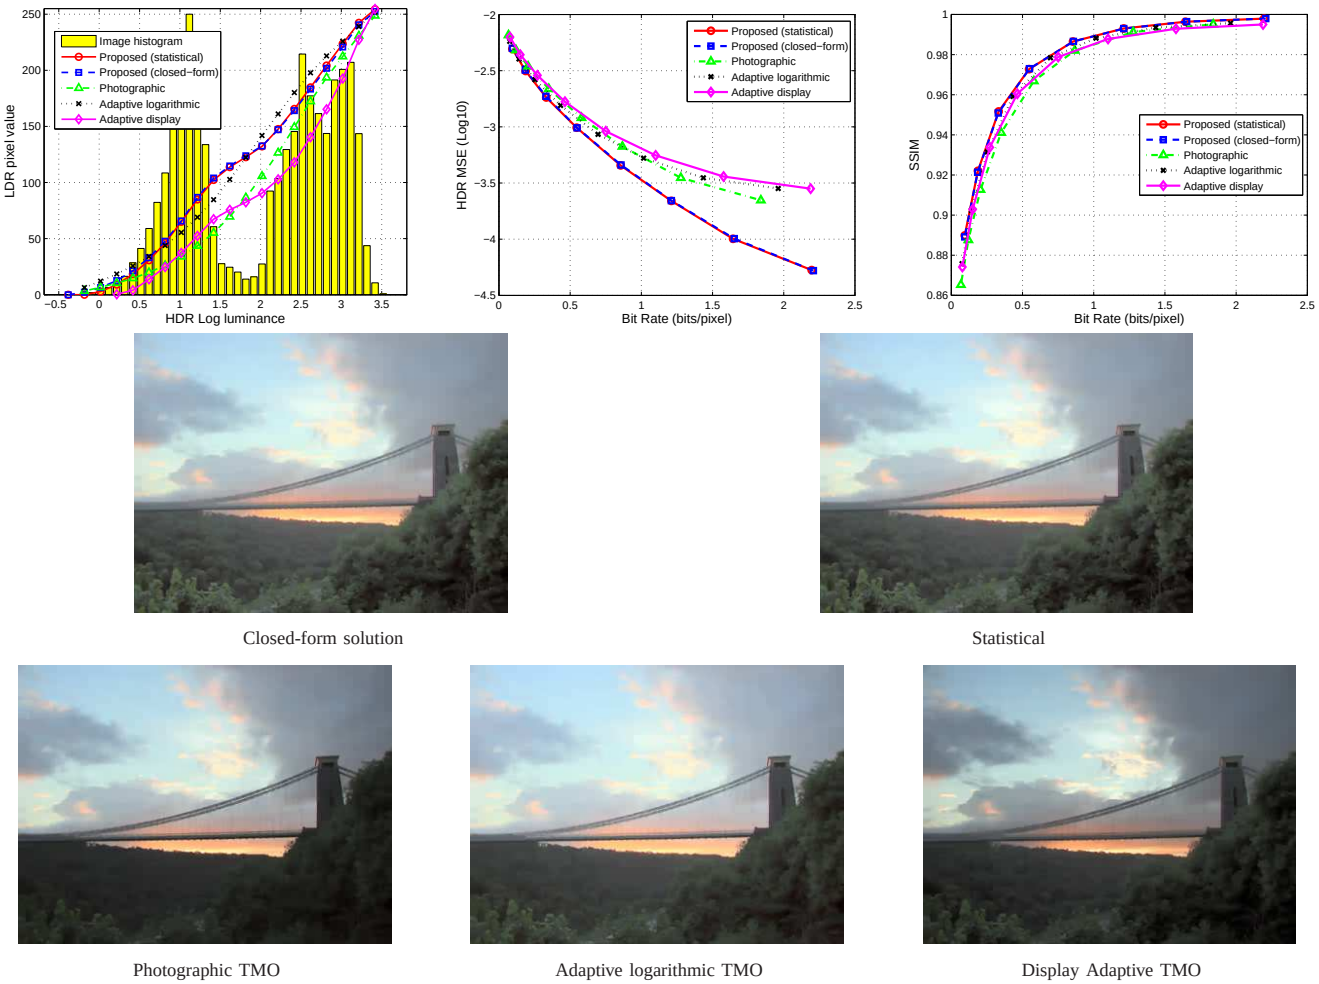

(a) Image: BristolBridge; LDR images shown are compressed with QP = 10

Fig. 10. Rate-distortion curves, tone curves and tone-mapped images for the image "BristolBridge". The notation is the same as Fig. 8. The compression quantization parameters used for "BristolBridge" is 38. The number of segments used for the histogram is 39.

dard sequences [39] using our closed-form method, and then compare the compression performance between the sequences generated by our method and the generic tone-mapping used in the test sequences.

The JVT test sequences [39] are provided as 10-bit extended dynamic range frames and the corresponding 8-bit tone-mapped frames. The 10-bit frames contain gamma corrected footage from a high-end camera that can capture an extended dynamic range. We assume that the gamma correction has a similar effect as the logarithmic function that we apply to linear luminance values to account for the Weber law. Therefore, we use the 10-bit frames directly as input to our algorithm. One important issue to consider is flickering, which our method can cause when used on video sequences. This is because the computed tone-curves solely depend on scene content, which can abruptly change from frame to frame. To prevent such flickering, we apply a low-pass filter to the generated tone-curves, identical to that in [25]. For comparison, we also generate tone-mapped sequences without the temporal filter.

Fig. 12 shows the comparison of the compression performance for our method (closed-form solution) with and without the temporal filtering, and the generic tone-mapping used for the JVT test sequences. The temporal filtering did not change significantly compression performance for these test sequences (compare black and red curves in Fig. 12). This is because the frames did not contain any abrupt scene changes that could cause flickering. For two sequences (Freeway and Waves) our method gave significant improvement over a generic tone-mapping that was not optimized for video compression (compare red and dashed-blue curves). The improvement is especially large for higher bit-rates. For the third sequence (Plane), the compression performance was very similar for both methods. Slightly worse result for our method for medium bit-rates can be explained by the approximations used in the closed-form solution. The tone-mapped sequence provided by the JVT was coincidentally well conditioned for video compression. However large improvements for the two remaining sequences illustrate gains that can be expected from the proposed method when used in combination with the JVT bit-depth scalable coding.

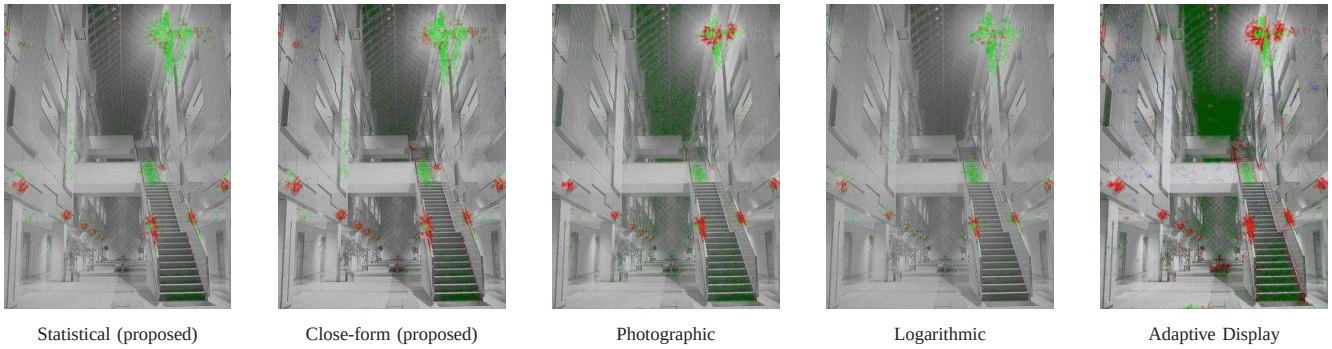

(a) Image: AtriumNight; distortion maps of tone-mapped images (no compression)

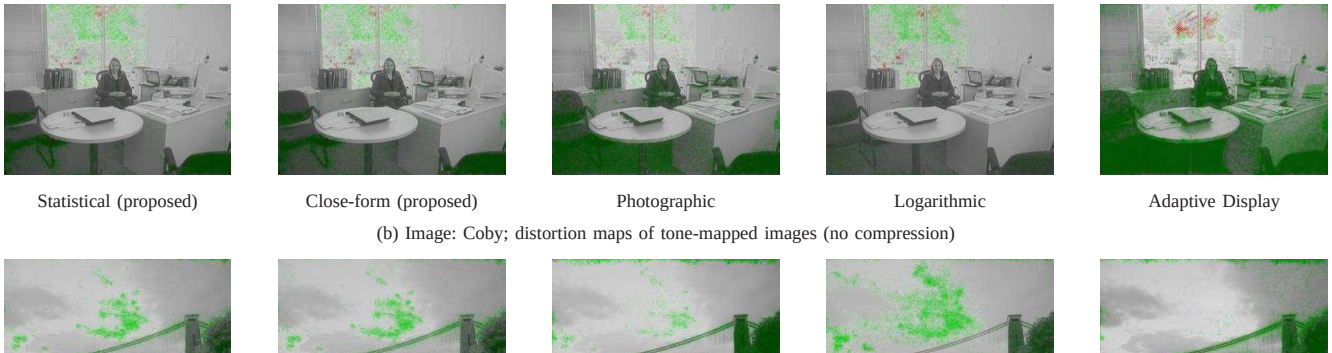

(b) Image: Coby; distortion maps of tone-mapped images (no compression)

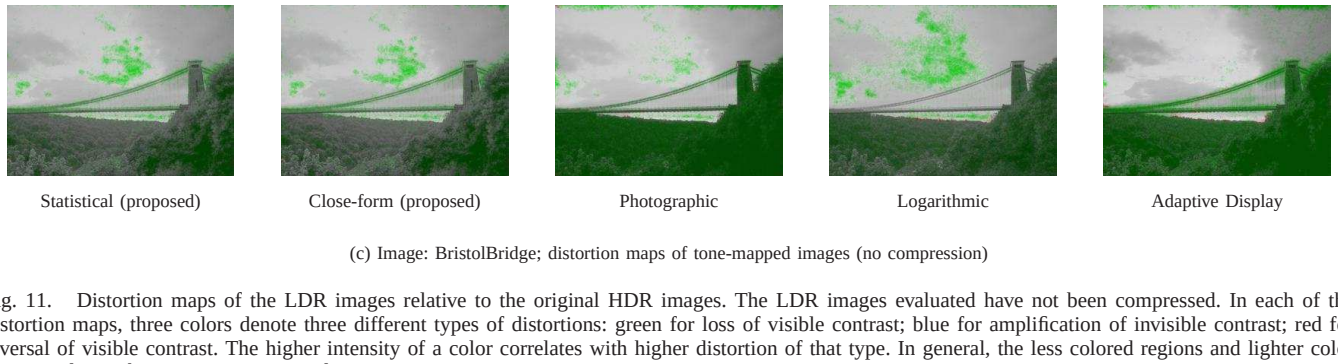

(c) Image: BristolBridge; distortion maps of tone-mapped images (no compression)

Fig. 11. Distortion maps of the LDR images relative to the original HDR images. The LDR images evaluated have not been compressed. In each of the distortion maps, three colors denote three different types of distortions: green for loss of visible contrast; blue for amplification of invisible contrast; red for reversal of visible contrast. The higher intensity of a color correlates with higher distortion of that type. In general, the less colored regions and lighter color intensity denote better LDR image quality.

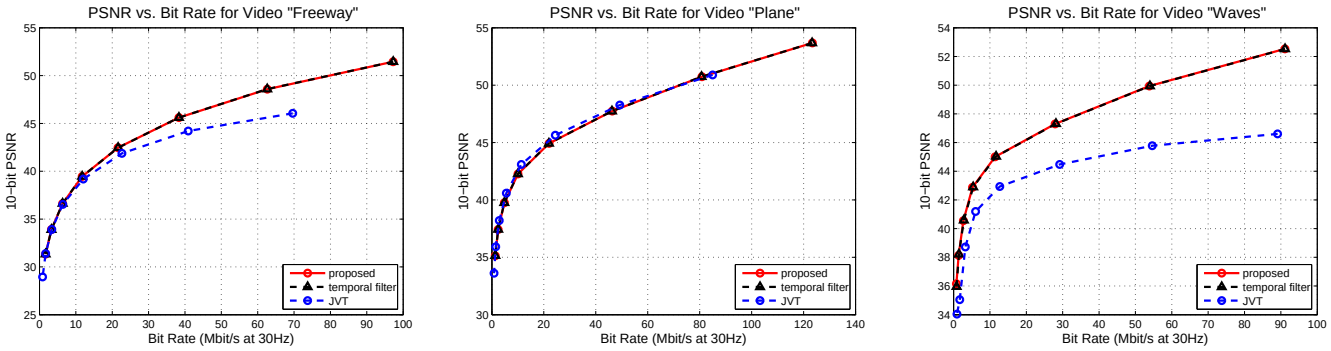

Fig. 12. Compression performance for the proposed method used with the the JVT bit-depth scalable encoding. The comparison is made between the closed-form solution of our method (red curve), the same method but with the temporal filter that prevents flickering (black), and the generic tone-mapping used in the JVT test sequences (dashed-blue). The x-axis denotes the bit rate in Mbit/s at 30Hz of frame rate, and the y-axis is the PSNR between the original 10-bit video and its inversely-tone-mapped version. 100 segments are used in the histogram for each of three sequences.

V. Discussion

Our proposed tone-mapping methods can directly improve the compression efficiency of bit-depth scalable video coding. The proposed methods are designed to produce a better reconstructed HDR representation. A direct consequence of this design objective is a reduction in the size of the higher bit-depth enhancement layer which contains the difference between the reconstructed HDR image and the original one. Thus, our method can lower the total bit rate for the bit-depth scalable coding.

Although the primary goal of tone-mapping algorithms is to optimize the visual quality of the displayed LDR image, we instead designed a tone-mapping operator that optimizes the compression performance in the backward-compatible encoding scheme. Numerous algorithms in the literature have considered the primary objective that explicitly focuses on

producing good quality tone-mapped images. This group of algorithms include the photographic [31], logarithmic [32] and display adaptive [25] TMOs considered in this paper. The tone-curve that meets both objectives can be approximated by the linear combination of the tone-curves produced by our method and these algorithms. Moreover, our study shows that the overall quality of images tone-mapped with our method is comparable to other tone-mapping algorithms and none of the tone-mapped images we generated was considered as unacceptable. This means that for the applications that do not require a finely adjusted backward-compatible layer, our method can be used directly.

Our considerations are limited to global tone-curves used for the entire frame while many modern tone-mapping methods use local (spatially varying) processing to retain more details and produce better looking images. However, the study in [30] showed that local tone-mapping operators result in worse compression performance than global operators. This suggests that there is a trade-off between choosing a tone-mapping operator that is optimal for preserving HDR information in compressed images and an operator that produces the best looking images. The study of such a trade-off and the application of local tone mapping for compression is an interesting topic for future work.

Finally, tone-mapping each frame of a video sequence independently can produce flickering since the tone-curve can change rapidly from frame to frame. This, however, can be avoided when a low-pass filter is applied on the sequence of computed tone-curves, as done in [25].

VI. Conclusion

In this paper, we showed that the appropriate choice of a tone-mapping operator (TMO) can significantly improve the reconstructed HDR quality. We developed a statistical model that approximates the distortion resulting from the combined processes of tone-mapping and compression. Using this model, we formulated a constrained optimization problem that finds the tone-curve which minimizes the expected HDR MSE. The resulting optimization problem, however, suffers from high computational complexity. Therefore, we presented a few simplifying assumptions that allowed us to reduce the optimization problem to an analytically tractable form with a closed-form solution. The closed-form solution is computationally efficient and has a performance compatible to our developed statistical model. Moreover, the closed-form solution does not require the knowledge of QP, which makes it suitable for cases where the compression strength is unknown. Although our models are designed to minimize HDR MSE, the extensive performance evaluations show that the proposed methods provide excellent performance in terms of SSIM and the LDR image quality, in addition to an outstanding performance in MSE.

Appendix A
H.264/AVC intra coding error model (p_C)

Let $v, \tilde{v} \in \{0, 1, \cdots, v_{max}\}$, be the original and decoded values of a LDR pixel, respectively. We denote by $p_C(v - \tilde{v})$ the probability that decoded pixel luma level has shifted by a factor $v - \tilde{v}$ from its original luma level.

For a specific image, $p(v - \tilde{v})$ can be estimated by subtracting each pixel value of the de-compressed image from that of the original image and then fitting a distribution for these differences by sampling such distribution on a large set of compression-distorted images. Let w be equal to $v - \tilde{v}$. We found that the probability can be well approximated with the General Gaussian Distribution (GGD):

$$p_C(w : \mu, \sigma, \alpha) = \frac{\lambda(\alpha, \sigma) \cdot \alpha}{2 \cdot \Gamma(1/\alpha)} \cdot e^{-[\lambda(\alpha,\sigma)\cdot|w-\mu|]^{\alpha}}, \quad (17)$$

where μ denotes the mean, σ is the standard deviation and α denotes the shape parameter. The functions λ and Γ are expressed as follows:

$$\lambda = \frac{1}{\sigma} \cdot \left[\frac{\Gamma(3/\alpha)}{\Gamma(1/\alpha)}\right]^{1/2}, \quad (18)$$

$$\Gamma(m) = \int_0^{\infty} t^{m-1} \cdot e^{-t} dt, z > 0, \quad (19)$$

and (19) is called the gamma function. In order to find a GGD fitting, the values of μ, σ and α need to be assigned. The distribution mean is set equal to 0 since all compression schemes make every effort to keep decoded pixel values unchanged. To find the standard deviation σ and the shape parameter α that best fit the image histograms, we use the least square regression. Note that σ and α vary for different images and different values of quantization parameters (QPs). Fig. 13 shows example error distributions and the resulting fitting curves.

We collected the estimated α and σ for a large number of images and for different QPs using H.264/AVC intra-frame mode. Fig. 14 demonstrates the results of α and σ vs. the value of QP for different images. We found that σ and α can be well described by the functions of QP:

$$\sigma = a \cdot \mathrm{QP}^2 - b \cdot \mathrm{QP} + c, \quad (20)$$

$$\alpha = 1 + e^{(d \cdot \mathrm{QP} + g)}, \quad (21)$$

where a, b, c, d and g are constants equal to 0.00625, 0.12457, 1.2859, -0.1 and 1.32, respectively.

We also modelled the error distributions for H.264/AVC predicted frames (P frames), bi-directional predicted frames (B frames) and JPEG compression. We found that their compression errors can be well estimated by GGD too. Please refer to the supplementary documents for details [40].

References

[1] H. Seetzen, W. Heidrich, W. Stuerzlinger, G. Ward, L. Whitehead, M. Trentacoste, A. Ghosh, and A. Vorozcovs, "High dynamic range display systems," *ACM Transactions on Graphics (Proc. SIGGRAPH)*, vol. 23, no. 3, pp. 757–765, 2004.

[2] IEC61966-2-4, "Colour measurement and management - part 2-4: Colour management - extended-gamut YCC colour space for video applications - xvYCC," 2006.

[3] Rafał Mantiuk, Grzegorz Krawczyk, Karol Myszkowski, and Hans-Peter Seidel, "Perception-motivated high dynamic range video encoding," *ACM Transactions on Graphics (Proc. SIGGRAPH)*, vol. 23, no. 3, pp. 730–738, 2004.

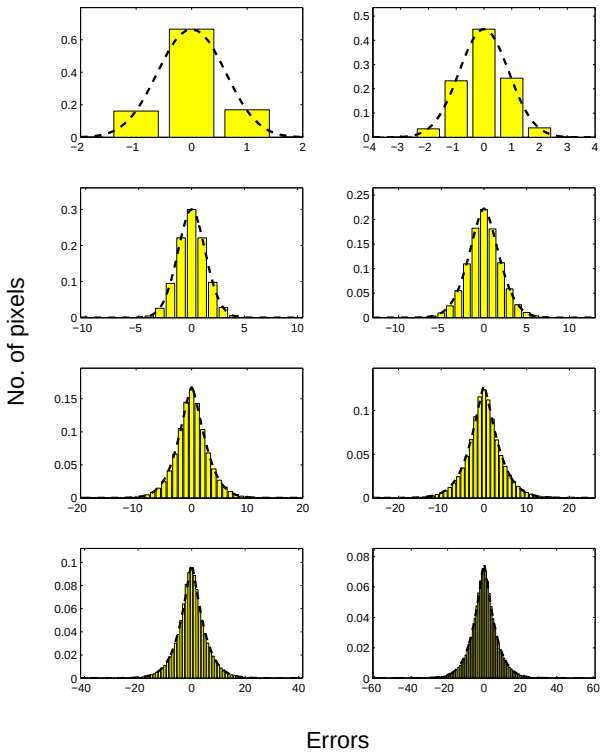

Fig. 13. Error distributions fitted with GGD distribution curves. Yellow bars denote the histogram of the compression error and dash black line is the fitted GGD curve.

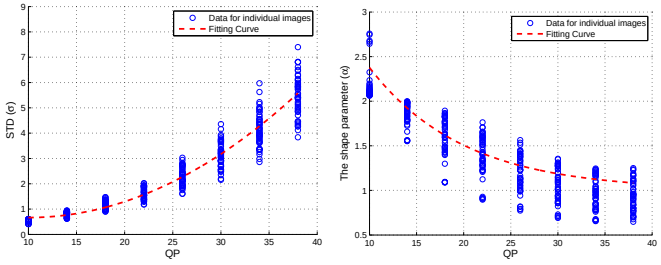

(a) (b)

Fig. 14. The standard deviation and the shape parameter vs. the value of QP for H.264/AVC encoding over a large pool of images. The blue circle corresponds to a single image at a certain QP. The red line is the averaging curve used to remove the image dependency.

[4] G. Ward, "Real pixels," *Graphics Gems II*, pp. 80–83, 1991.

[5] R. Bogart, F. Kainz, and D. Hess, "Openexr image file format," *ACM SIGGRAPH 2003,Sketches and Applications*, 2003.

[6] G. Ward Larson, "Logluv encoding for full-gamut, high-dynamic range images," *Journal of Graphics Tools*, vol. 3, no. 1, pp. 15–31, 1998.

[7] Y. Gao and Y. Wu, "Applications and requirement for color bit depth scalability," Tech. Rep., ISO/IEC JTC1/SC29/WG11 and ITU-T SG16 Q.6, JVT-U049, Oct. 2006.

[8] G.J. Sullivan, Haoping Yu, S.-i. Sekiguchi, Huifang Sun, T. Wedi, S. Wittmann, Yung-Lyul Lee, A. Segall, and T. Suzuki, "New standardized extensions of mpeg4-avc/h.264 for professional-quality video applications," *Image Processing, 2007. ICIP 2007. IEEE International Conference on*, vol. 1, pp. I–13–16, 2007.

[9] M. Winken, D. Marpe, H. Schwarz, and T. Wiegand, "Bit-depth scalable video coding," *Image Processing, 2007. ICIP 2007. IEEE International Conference on*, vol. 1, pp. I–5–8, 2007.

[10] A. Segall and J. Zhao, "Bit-stream rewriting for SVC-to-AVC conversion," in *Image Processing, 2008. ICIP 2008. 15th IEEE International Conference on*, 2008, pp. 2776–2779.

[11] Y. Gao and Y. Wu, "Bit depth scalability," Tech. Rep., ISO/IEC JTC1/SC29/WG11 and ITU-T SG16 Q.6, JVT-V061, Jan. 2007.

[12] Y. Ye, H. Chung, M. Karczewicz, and I. S. Chong, "Improvements to Bit Depth Scalability Coding," Tech. Rep., ISO/IEC JTC1/SC29/WG11 and ITU-T SG16 Q.6, JVT-Y048, Oct. 2007.

[13] M. Winken, H. Schwarz, D. Marpe, and T. Wiegand, "SVC bit depth scalability," Tech. Rep., ISO/IEC JTC1/SC29/WG11 and ITU-T SG16 Q.6, JVT-V078, Jan. 2007.

[14] Yong Yu, Steve Gordon, and Michael Yang, "Improving Compression Preformance in Bit Depth SVC with a Prediction Filter," Tech. Rep., ISO/IEC JTC1/SC29/WG11 and ITU-T SG16 Q.6, JVT-Z045, Jan. 2008.

[15] A. Segall and Y. Su, "System for bit-depth scalable coding," Tech. Rep., ISO/IEC JTC1/SC29/WG11 and ITU-T SG16 Q.6, JVT-W113, April 2007.

[16] Andrew Segall, Louis Kerofsky, and Shawmin Lei, "Tone Mapping SEI Message," Tech. Rep., ISO/IEC JTC1/SC29/WG11 and ITU-T SG16 Q.6, JVT-T060, July 2006.

[17] K. E. Spaulding, G. J. Woolfe, and R. L. Joshi, "Using a residual image to extend the color gamut and dynamic range of an sRGB image," in *Proc. of IS&T PICS Conference*, 2003, pp. 307–314.

[18] Greg Ward and Maryann Simmons, "JPEG-HDR: A backwards-compatible, high dynamic range extension to JPEG," in *Proceedings of the 13th Color Imaging Conference*, 2005, pp. 283–290.

[19] Rafał Mantiuk, Alexander Efremov, Karol Myszkowski, and Hans-Peter Seidel, "Backward compatible high dynamic range mpeg video compression," *ACM Transactions on Graphics (Proc. SIGGRAPH)*, vol. 25, no. 3, 2006.

[20] A. Segall, "Scalable coding of high dynamic range video," in *Image Processing, 2007. ICIP 2007. IEEE International Conference on*, 2007, vol. 1, pp. I –1–I –4.

[21] Y. Wu, Y. Gao, and Y. Chen, "Bit-depth scalability compatible to H. 264/AVC-scalable extension," *Journal of Visual Communication and Image Representation*, vol. 19, no. 6, pp. 372–381, 2008.

[22] Heiko Schwarz, Detlev Marpe, and Thomas Wiegand, "Overview of the Scalable Video Coding Extension of the H.264/AVC Standard," *Transactions on Circuits and Systems for Video Technology*, vol. 17, no. 9, pp. 1103–1120, Sepetember 2007.

[23] Shan Liu, Woo-Shik Kim, and Anthony Vetro, "Bit-depth scalable coding for high dynamic range video," in *Proc. of SPIE Visual Communications and Image Processing 2008*, William A. Pearlman, John W. Woods, and Ligang Lu, Eds., 2008, vol. 6822, p. 649212.

[24] Patrick Ledda, Alan Chalmers, Tom Troscianko, and Helge Seetzen, "Evaluation of tone mapping operators using a high dynamic range display," *ACM Transactions on Graphics (Proc. SIGGRAPH)*, vol. 24, no. 3, pp. 640–648, 2005.

[25] R. Mantiuk, S. Daly, and L. Kerofsky, "Display Adaptive Tone Mapping," *ACM Transactions on Graphics-TOG*, vol. 27, no. 3, pp. 68–68, 2008.

[26] Erik Reinhard, Greg Ward, Sumanta Pattanaik, and Paul Debevec, *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, Morgan Kaufmann, 2005.

[27] Yuanzhen Li, Lavanya Sharan, and Edward H. Adelson, "Compressing and companding high dynamic range images with subband architectures," *ACM Trans. Graph.*, vol. 24, no. 3, pp. 836–844, 2005.

[28] Chul Lee and Chang-Su Kim, "Gradient domain tone mapping of high dynamic range videos," in *Image Processing, 2007. ICIP 2007. IEEE International Conference on*, 2007, vol. 3, pp. III –461–III –464.

[29] R. Fattal, D. Lischinski, and M. Werman, "Gradient domain high dynamic range compression," *Computer Graphics Forum (Proc. of Eurographics)*, vol. 21, no. 3, pp. 249–256, 2002.

[30] R. Mantiuk and H.P. Seidel, "Modeling a Generic Tone-mapping Operator," *Computer Graphics Forum (Proc. of Eurographics'08)*, vol. 27, no. 2, pp. 699–708, 2008.

[31] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda, "Photographic tone reproduction for digital images," *ACM Transactions on Graphics (Proc. SIGGRAPH)*, vol. 21, no. 3, pp. 267–276, 2002.

[32] F. Drago, K. Myszkowski, T. Annen, and N. Chiba, "Adaptive logarithmic mapping for displaying high contrast scenes," *Computer Graphics Forum (Proc. of Eurographics)*, vol. 22, no. 3, pp. 419–426, 2003.

[33] Z. Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.

[34] R. Mantiuk, R. Mantiuk, A. Tomaszewska, and W. Heidrich, "Color correction for tone mapping," *Computer Graphics Forum (Proc. of EUROGRAPHICS)*, vol. 28, no. 2, pp. 193–202, 2009.

[35] Selig Hecht, "The visual discrimination of intensity and the weber-fechner law," *The Journal of General Physiology*, vol. 7, no. 2, pp. 235C267, November 1924.

[36] "H.264/AVC JM14.2 Reference Software," Available: http://iphome.hhi.de/suehring/tml/.
[37] T. O. Aydin, R. Mantiuk, K. Myszkowski, and H.-P. Seidel, "Dynamic range independent image quality assessment," *ACM Transactions on Graphics (Proc. SIGGRAPH)*, vol. 27, no. 3, 2008.
[38] R. Mantiuk, S. Daly, K. Myszkowski, , and H.-P. Seidel, "Predicting visible differences in high dynamic range images model and its calibration," in *Proc. of Human Vision and Electronic Imaging X*. 2005, vol. 5666 of *Proceedings of SPIE*, pp. 204–214, SPIE.
[39] P. Topiwala and H. Yu, "New Test Sequences in the VIPER 10-bit HD Data," Tech. Rep., ISO/IEC JTC1/SC29/WG11 and ITU-T SG16 Q.6, JVT-Q090, Oct. 2005.
[40] "Supplimentary documents and results," Available: http://www.ece.ubc.ca/∼zicongm/tmo_hdr_enc/HTMLReport.html.

Zicong Mai is currently a Ph.D. candidate with the Electrical and Computer Engineering Department at the University of British Columbia. His recent research focuses on developing processing and compression algorithms for new-generation visual signals, such as high dynamic range (HDR) images/video and stereoscopic video.

He received his Master's degree also from the Electrical and Computer Engineering Department at the University of British Columbia, and was the award holder of University Graduate Fellowship. His research involved realizing compatibility between interactive multimedia systems, including Blu-ray and DVB-based iTV.

Hassan Mansour (S'99 - M'09) received the B.E. degree in Computer and Communications Engineering from the American University of Beirut, Lebanon, in 2003. He received the M.A.Sc. and PhD degrees in Electrical and Computer Engineering from the University of British Columbia, Vancouver, BC, Canada, in 2005 and 2009, respectively, where his research focused on multimedia communication and video coding. He is currently a postdoctoral research fellow with the departments of Computer Science and Mathematics at the University of British Columbia where he is conducting research in compressed sensing and sparse signal reconstruction algorithms for distorted image and video signals.

Rafal Mantiuk received PhD from the Max-Planck-Institute for Computer Science (Germany) in 2006. He is a member of the Wales Research Institute of Visual Computing and Lecturer (Assistant Professor) in the School of Computer Science at Bangor University, North Wales, UK. His recent interests focus on designing imaging algorithms that adapt to human visual performance and viewing conditions in order to deliver the best images given limited resources, such as computation time, bandwidth or dynamic range. He investigates how the knowledge of the human visual system and perception can be incorporated within computer graphics and imaging algorithms. More on his research can be found at: http://goo.gl/UkCrt.

Panos Nasiopoulos received a Bachelor degree in Physics from the Aristotle University of Thessaloniki (1980), Greece, and a Bachelor (1985), Master (1988) and Ph.D. (1994) in Electrical and Computer Engineering from the University of British Columbia, Canada. He is presently the Director of the Institute for Computing, Information and Cognitive Systems (160 faculty members and 800 graduate students) at the University of British Columbia (UBC). He is also a Professor with the UBC department of Electrical and Computer Engineering, the Inaugural Holder of the Dolby Professorship in Digital Multimedia, and the current Director of the Master of Software Systems Program at UBC. Before joining UBC, he was the President of Daikin Comtec US (founder of DVD) and Executive Vice President of Sonic Solutions. Dr. Nasiopoulos is a registered Professional Engineer in British Columbia, the Associate Editor for IEEE Signal Processing Letters, and has been an active member of the Standards Council of Canada, the IEEE and the Association for Computing Machinery (ACM).

Rabab Kreidieh Ward is a Professor in the Electrical and Computer Engineering Department at the University of British Columbia, Vancouver, B.C., Canada. She is presently with the Office of the Vice President Research and International as the natural sciences and engineering research coordinator. Her research interests are in signal, image and video processing. She has made contributions in the areas of signal detection, image encoding, compression, recognition, restoration and enhancement, and their applications to infant cry signals, cable TV, HDTV, medical images, and astronomical images. She has published over 380 journal and conference papers and holds six patents related to cable television picture monitoring, measurement and noise reduction. Applications of her work have been transferred to U.S. and Canadian industries.

She was a Vice President of the IEEE Signal Processing Society, 2003-2005, the General Chair of the IEEE International Conference on Image Processing 2000, the IEEE Symposium on Signal Processing and Information Technology 2006, and the Vice Chair of the IEEE International Symposium on Circuits and Systems 2004. She is a Fellow of the Royal Society of Canada, the IEEE, Canadian Academy of Engineers and Engineering Institute of Canada, and a recipient of the UBC 1997 Killam Research Prize, YWCA Woman of Distinction Award (2008) , the R. A. McLachlan Memorial Award (the top award) of the Association of Professional Engineers and Geoscientists of British Columbia and the "Society Award "of the IEEE Signal Processing Society.

Wolfgang Heidrich is a Professor and the Dolby Research Chair in Computer Science at the University of British Columbia. He received a PhD in Computer Science from the University of Erlangen in 1999, and then worked as a Research Associate in the Computer Graphics Group of the Max-Planck-Institute for Computer Science in Saarbrucken, Germany, before joining UBC in 2000. Dr. Heidrich's research interests lie at the intersection of computer graphics, computer vision, imaging, and optics. In particular, he has worked on High Dynamic Range imaging and display, image-based modeling, measuring, and rendering, geometry acquisition, GPU-based rendering, and global illumination. Dr. Heidrich has written over 100 refereed publications on these subjects and has served on numerous program committees. He was the program co-chair for Graphics Hardware 2002, Graphics Interface 2004, and the Eurographics Symposium on Rendering, 2006.