

Temporal Coherency for Video Tone Mapping

Ronan Boitard^{a,b}, Kadi Bouatouch^b, Rémi Cozot^b, Dominique Thoreau^a and Adrien Gruson^b

^aTechnicolor R&D France, 1 av. Belle Fontaine, Rennes, France;

^bIRISA, 263 Avenue du Général Leclerc, Rennes, France

ABSTRACT

Tone Mapping Operators (TMOs) aim at converting real world high dynamic range (HDR) images captured with HDR cameras, into low dynamic range (LDR) images that can be displayed on LDR displays. Several TMOs have been proposed over the last decade, from the simple global mapping to the more complex one simulating the human vision system. While these solutions work generally well for still pictures, they are usually less efficient for video sequences as they are source of visual artifacts. Only few of them can be adapted to cope with a sequence of images. In this paper we present a major problem that a static TMO usually encounters while dealing with video sequences, namely the temporal coherency. Indeed, as each tone mapper deals with each frame separately, no temporal coherency is taken into account and hence the results can be quite disturbing for high varying dynamics in a video. We propose a temporal coherency algorithm that is designed to analyze a video as a whole, and from its characteristics adapts each tone mapped frame of a sequence in order to preserve the temporal coherency. This temporal coherency algorithm has been tested on a set of real as well as Computer Graphics Image (CGI) content and put in competition with several algorithms that are designed to be time-dependent. Results show that temporal coherency preserves the overall contrast in a sequence of images. Furthermore, this technique is applicable to any TMO as it is a post-processing that only depends on the used TMO.

Keywords: Video Tone Mapping, HDR Video, Temporal Coherency

1. INTRODUCTION

With the development of sensor technology, it is now possible to capture more information than can be displayed. While some emerging technologies are capable of displaying broader luminance ranges, they are still quite expensive and will not be available on the customer market soon. That is why some operations are still needed to convert High Dynamic Range (HDR) contents to Low Dynamic Range (LDR) ones. These operations are performed using Tone Mapping Operators (TMOs), also called Tone Reproducers. Technology regarding the tone mapping of still pictures has been thoroughly explored and several satisfying solutions¹ have been designed. Furthermore, these operators have various goals and their results can be quite different. Nevertheless, most of the existing operators are not designed to handle video sequences.

Adapting a tone mapping operator to a video sequence has been the focus of several papers.²⁻⁵ In these papers, temporal coherency has been addressed in different ways. The first one consists in preserving the overall tone mapped LDR values from frame to frame to avoid flickering. The second ensures the perception consistency of an object throughout the video, say similar levels of HDR value should be tone mapped to similar levels of LDR ones. The last way aims at preserving the temporal overall contrast consistency. Indeed, a difference in perceived HDR brightness levels should result in a similar difference of perception in LDR brightness levels. An illustration of these problems is depicted in Fig. 1. In this paper, we will deal with the last two ways of addressing temporal coherency as the first one has been widely coped with.²⁻⁵ Our solution consists of two steps. The first step applies a TMO to each image of the video separately. The second assigns a new LDR luminance range to each frame depending on some characteristics of the original HDR video.

The next section briefly presents related works in temporal adaptation and video tone mapping. Our solution is described in section 3 while section 4 presents some results. Finally, future work and remaining problems are addressed in the last section.

Further author information: (Send correspondence to first name last name)

a: E-mail: first.last@technicolor.com, Telephone: +33 (0)2 99 27 30 00

b: E-mail: first.last@irisa.fr, Telephone: +33 (0)2 99 84 71 00



Figure 1: Illustration of lack of temporal coherency in video tone mapping. On the false color representation (a), red represents the maximum HDR value of the video while blue its minimum. Both frames have the same HDR values between similar regions. However the scene (c) appears a lot brighter than the scene (b). Furthermore, the perception of the landscape is different between the two frames.

2. RELATED WORK IN VIDEO TONE MAPPING

Tone mapping aims at converting HDR luminance values (L_w) into LDR luma values (L_d). Luma is the weighted sum of gamma-compressed R'G'B' components of a color video. The 3 original HDR color channels RGB are scaled by L_d/L_w so as to preserve both the hue and the saturation in the tone mapped LDR image. The aims of TMOs can be quite different since they depend on the targeted applications. While our approach adapts easily to any TMO, in this work, similarly to other publications,^{3,4} we decided to build our solution upon Reinhard's TMO, called Photographic Tone Reproduction (PTR).⁶

In this section, the PTR algorithm is first introduced, before describing two methods that adapt this TMO to video tone mapping. Afterwards, two other TMOs designed for video tone mapping are presented. Finally, issues solved by these techniques are listed as well as the problems still unsolved.

2.1 Photographic Tone Reproduction

The PTR algorithm is based on photographic techniques and allows to choose the exposure of the tone mapped image. More precisely, this operator relies on the zonal system designed by Adams.⁷ This exposure is used to calibrate the input data (Equation (2)) depending on the key values of the image to be tone mapped. The key value of a scene, κ , is a subjective indication of its overall brightness. It is given by Equation (1):

$$\kappa = \exp\left(\frac{1}{N} \cdot \sum_{n=1}^N \log(\delta + L_w^n)\right), \quad (1)$$

$$L_s = \frac{a}{\kappa} \cdot L_w, \quad (2)$$

where a is the chosen exposure, δ is a small value to avoid singularity, L_w is the HDR luminance and L_s the scaled luminance. N is the number of pixels in the image. The tone mapping operation consists in using a tone map curve (sigmoid function) given by Equation (3):

$$L_d = \frac{L_s}{1 + L_s} \cdot \left(1 + \frac{L_s}{L_{white}^2}\right), \quad (3)$$

where L_{white} is used to burn out areas with a high luminance value and L_d is the LDR luma. Two parameters (a and L_{white}) are then necessary to compute the TMO results. For the sake of simplicity, these parameters are fixed for every sequence to $a = 18\%$ and $L_{white} = 2 \cdot L_{max}$.

Although this TMO has been well rated in several evaluations,⁸⁻¹⁰ it is not designed to work properly for video sequences. Indeed, if the key values change, the tone map curve also changes, resulting in different mappings for similar HDR luminance levels in successive frames, which is source of flickering. In the following section, two solutions are presented to avoid this flickering.




Figure 2: Evolution of the frame key value computed for every frame of two video sequences. On plot (a), an offset is added to avoid an overlap between the curves, the smoothing effect of both techniques are compared to Reinhard’s et al. technique. Kang’s et al. technique smoothes and propagates a change of value while Ramsey’s et al. technique only smoothes it. On plot (b), Ramsey’s et al. improvement reduces the temporal shifting compared to Kang’s et al. when sharp variations occurs.

2.2 Toward Video Tone Mapping

Kang et al.³ proposed a method based on Reinhard’s operator to reduce the flickering effect. Each image is tone mapped independently using a key value depending on a set of previous frames. More precisely, this key value is computed using Equation (1) with N the number of pixels of this set augmented with the one of the current image. As a consequence, this method smoothes abrupt variations of the frame key value throughout the video sequence. This technique is capable of reducing flickering for sequences with slow illumination variations. It is a first step toward temporal coherency as the processing of each frame depends on its previous ones.

However, for high variations, this technique fails because it considers a fixed number of previous frames. That is why, Ramsey et al.⁴ proposed a method that adapts dynamically this number. The adaptation process depends on the variation of the current frame key value and that of the previous frame. Moreover, the adaptation discards outliers using a min/max threshold. This solution performs better than that of Kang’s et al. for a wider range of video sequences. The computed key value for both of these techniques and the PTR algorithm is plotted in Fig. 2.

2.3 Display Adaptive Tone Mapping

Mantiuk et al.⁵ proposed a TMO that provides the least perceptually distorted LDR picture on a targeted display. Similarly to Tumblin and Rushmeier,¹¹ whose operator’s goal is to ensure that the scene and display brightnesses match, Mantiuk et al. compare the visual response of the Human Visual System (HVS) to the display-adapted LDR image with that of the original HDR image. The minimization of the residual error between these responses results in a piece-wise linear tone map curve. However, this minimization is computed on a per frame basis and consequently inherits all the problems related to video tone mapping.

To solve this issue, the authors propose a temporal adaptation using a low-pass filter applied to the nodes of the piece-wise tone curve. Similarly to the previous solutions,^{3,4} the tone map curve is smoothed. Nevertheless, this solution’s drawbacks are similar to those of Kang’s et al. technique. In addition, an optimization performed on a per frame basis is quite different from an optimization applied to the whole sequence. Consequently the minimization of the perceptual distortion is not preserved during the temporal adaptation.




Figure 3: Change of perception of an object when the TMO processes the frames independently. The constant HDR luminance of an object (green line) is tone mapped to high varying LDR luma values (green dashed line). The highly changing maximum HDR luminance value of every frame (orange line) is tone mapped to the same LDR luma level (orange dashed line).

2.4 Time Dependent Visual Adaptation Tone Mapping

Pattanaik et al.¹² designed a TMO called Time Dependent Visual Adaptation (TDVA). The aim of this operator is to simulate the adaptation of the HVS when a change of illumination occurs. Similarly to other HVS based operators,^{13,14} this TMO simulates both cone and rod photoreceptor responses using Hunt’s model:¹⁵

$$R_{rod} = B_{rod} \frac{L_{rod}^n}{L_{rod}^n + \sigma_{rod}^n}, \quad R_{cone} = B_{cone} \frac{L_{cone}^n}{L_{cone}^n + \sigma_{cone}^n}, \quad (4)$$

where R is the response to the HDR luminance L , n is a sensitivity constant and both B and σ are determined depending on the overall scene luminance. The combination of these responses is adapted to the targeted display range through an appearance/inverse appearance model. Finally, those responses are inverted to compute the tone mapped LDR luma.

While this TMO is called time-dependent, it was not designed for video sequences. It only simulates the time needed for the eye to adapt to a change of illumination condition, say when the photoreceptors are fully adapted. This adaptation time is modeled by exponential decay functions. This TMO computes each frame separately while taking into account the adaptation state of the previous frame.

2.5 Open Issues

As seen in the previous subsections, the first three proposed methods³⁻⁵ aim at reducing the flickering due to changes in the tone map curve throughout subsequent frames. Their solutions smooth the evolution of these curves either by adapting the key value to several frames^{3,4} or by filtering the tone map curve itself.⁵ On the other hand, the last method¹² simulates the behavior of the HVS during the transition between two fully adapted states. However, this method also reduces the flickering effect, as it prevents the variables, which control the tone map curve, from evolving too quickly.

While effectively reducing the flickering, these techniques do not solve all the problems inherent in video tone mapping such as the perception consistency of an object and the preservation of the overall temporal brightness (subjective value relative to luminance). The contribution of this paper is to address these two problems.

To ensure that the perception of an object is coherent throughout a video is our first aim. However, a TMO processes each frame separately and without information of the previous/next state of an object. That is why, an object with constant HDR luminance level throughout the whole video is tone mapped to several different LDR levels. This problem is illustrated in Fig. 3 where the green line represents an object with constant HDR luminance value over time whereas its associated tone mapped value varies.

Our second aim is to preserve the temporal brightness coherency. As a static TMO uses to the best the available display range, each frame is tone mapped to the best brightness level independently of the whole video sequence. Consequently, relative temporal HDR brightness information is not preserved during the course of the tone mapping process. In addition, the maximum HDR luminance value of each frame is tone mapped to the maximum LDR luma value, meaning that different HDR luminance values are tone mapped to the same LDR value. Fig. 3 illustrates this problem by the orange line which highly varies for the HDR sequence but is practically constant in the LDR sequence. Consequently, relative HDR brightness levels through subsequent frames are not preserved in the LDR frames. In the next section, we propose solutions to solve the two above mentioned problems.




Figure 4: Video HDR-LDR conversion workflow with temporal coherency post processing.

3. TEMPORAL COHERENCY SOLUTION

3.1 Overview

Contrary to the related work in video tone mapping, our approach preserves the relative luminance levels throughout the whole video. To this end, our method considers the whole video sequence rather than a subset of successive frames. As every TMO works for only certain applications, instead of creating a new video tone mapping operator for a limited set of applications, our goal is to design a general solution which adapts any TMO to video tone mapping.

A TMO processes every frame independently of the others. By adding a post processing operation at the output of the TMO, we can adapt it to video sequences. This operation needs information about the characteristics of the video. Consequently, a video analysis is needed.

Fig. 4 represents a usual workflow of a TMO, that we have extended to cope with video tone mapping. The HDR original luminance L_w is first computed from the input HDR RGB frames. Then the TMO computes the LDR luma L_d . The 3 RGB color channels are scaled by the ratio of the HDR luminance to the LDR luma, i.e. $\frac{L_d}{L_w}$. Finally, the resulting floating point values are converted to integer values within the range of the targeted display.

Our solution consists in adding two processing to the usual TMO workflow: a video analysis and a post processing operation. The video analysis process computes the characteristics of the video from the HDR luminance values. The post processing operation modifies each resulting LDR frame according to Equation (5). These two processing are described in the following subsections. We also propose a second post-processing technique.

3.2 Analysis of the video

The characteristic of the video used to express the relative brightness level in the HDR sequence is the key value κ_P of each frame. Let us recall that the key value of a scene is a subjective indication of its overall brightness. In addition, the key value κ_V associated with the whole video is also needed. We compute the key value κ_V using Equation (1), with N the total number of pixels of the whole video sequence.

3.3 Temporal Coherency Post Processing

Our post processing operation is designed to achieve two goals. First, an object with the same HDR luminance value throughout the HDR sequence should be tone mapped throughout the sequence to a similar LDR luma value. The second aim is to preserve the relative levels of brightness throughout the video. Fig. 5 illustrates the first aim which is the consistency of the perception of an object (green line). The preservation of relative brightness is represented by the new mapping of the maximum luma (orange line in Fig. 5).

Our proposed solution preserves the perception of an object throughout the video while ensuring that each frame is mapped with a brightness level that is relative to the whole video. Our technique reduces the range of a frame




Figure 5: Post processing. The LDR luma is the output of the TMO, see Fig. 3. The object luma (green line) variations have been reduced allowing a perception of the object through time stable. The maximum LDR luma (orange line) is reduce to a smaller range in order to preserve temporal brightness coherency.

to preserve the relative brightnesses. This reduction is performed using a scaling function (Equation (5)) whose shape is defined by the key video value κ_V and the current frame key value κ_F :

$$L'_d = \frac{\kappa_F}{\kappa_V + \kappa_F} \cdot L_d, \quad (5)$$

where L'_d is the post processed LDR luma value and L_d is the output of the TMO. However, in the case of low scaling value, the range is too small to represent a scene. To cope with this problem, we modify the slope of the scaling function and add an user defined offset $MinScale$ set by the user as shown in Equation (6):

$$L'_d = \left(MinScale + (1 - MinScale) \cdot \frac{\kappa_F}{\kappa_V + \kappa_F} \right) \cdot L_d. \quad (6)$$

This method performs well regarding the two aforementioned aims. However, the proposed scaling function works only for Reinhard's technique and its derivative. To adapt our solution to Kang's et al. or Ramsey's et al. techniques, we just need to compute new κ_F and κ_V values. However, adapting our method to another TMO is not straightforward but possible with some changes in the scaling function.

In case the preservation of the perception of an object is not considered, we propose in the next subsection, another method which works for every TMO.

3.4 Relative Brightness Preservation with any TMO

We want to devise a technique which adapts to any TMO. To this end, we propose a method which ignores the preservation of object perception but ensures the temporal brightness coherency between frames of the sequence. This solution works for any TMO.

Similarly to Tumblin and Rushmeier,¹¹ whose operator's goal is to ensure that the scene and display brightnesses match, in our method the relative difference of brightness between frames is preserved. To this end, the LDR luma L_d is modified to satisfy the following equation:

$$\frac{\kappa_F^{HDR}}{\kappa_{F,max}^{HDR}} = \frac{\kappa_F^{LDR}}{\kappa_{F,max}^{LDR}}, \quad (7)$$

where κ_F^{HDR} is the current HDR frame key value and $\kappa_{F,max}^{HDR}$ is the maximum frame key value of the video. The new scaling function is defined by:

$$L'_d = L_d \cdot \frac{\kappa_P^{HDR} \cdot \kappa_{P,max}^{LDR}}{\kappa_{P,max}^{HDR} \cdot \kappa_P^{LDR}} \quad (8)$$

where the κ_P are computed using Equation (1) and L'_d is the post-processed LDR luma. This solution preserves the relative HDR brightness levels in the LDR tone map results. Similarly to the previous section, where low scale ratio results in artifacts, we use the modifications added in Equation (6).




Figure 6: Let us consider one pixel. The plots represent the LDR value of this pixel throughout the video sequence and for different TMOs. Note the high variations of the plots, except the one corresponding to our technique.

In this section, two techniques were designed to deal with temporal coherency. The first technique uses Reinhard’s et al. operator and adjusts its output to preserve both the object’s perception consistency and the relative brightness level coherency. The second technique adapts to any TMO but only preserves the temporal brightness coherency.

4. RESULTS

4.1 Test Sequences

This section provides results of our methods for two video sequences. The first one, called UnderBridgeHigh, is an HDR panoramic picture¹⁶ of resolution 10656x5329. To convert it into a video sequence, a vertical traveling along the picture is performed. The frame’s resolution is 1920x1080 with a frame rate of 10 frames per second (fps). The traveling shift is of 40 pixels between 2 frames. This video sequence is used to test the object consistency issue. Indeed, pixels in successive HDR frames have exactly the same HDR values, so variations of LDR values can only be due to the TMO. The second video, a computer generated video called "ModernFlat", represents a furnished living room with three big windows in bright daylight. The camera records sequentially areas starting from the shadowy left side of the room to a direct view of the windows. This video has a resolution of 1280x720, is recorded at a frame rate of 10 fps. This second video sequence was designed to test the relative brightness coherency of tone mapped frames throughout the video.

4.2 Object Consistency Results

The consistency of the perception of an object throughout a video sequence is its ability to has low variations of LDR values when no change of illumination conditions occurs in the HDR video. Indeed, a small area of a scene is usually visible in many successive frames and have the same HDR luminance value. Consequently, a TMO which fully preserves the object consistency, tone maps every HDR luminance value of this area to the same LDR luma value.

To assess the performance of our solution, we used the first video sequence, namely UnderBridgeHigh. Let us consider a pixel with the same HDR luminance value along 28 frames. Fig. 6 shows the evolution of this pixel. For each already existing technique presented in section 2, the variations of the luma value of this pixel are high rather than being almost constant. Contrary to these solutions, our technique keeps this variation very low.

To show that our method performs well regarding the perception of an object throughout the video sequence, we took two HDR frames, number 25 and 35. We consider only a small area around the pixel studied in Fig. 6.




Figure 7: LDR images of the area surrounding the pixel studied in Fig. 6, resulting from different TMOs. Each column corresponds to one TMO, the upper image corresponding to frame 25 while the lower to frame 35. Our solution (e) is the only one that preserves the overall perception of this area.

In Fig. 7, we provide LDR images of this area resulting from some existing techniques. We see that our solution is the only one that preserves the overall perception of this area.

4.2.1 Temporal Brightness Coherency Results

Preserving the relative levels of brightness throughout a video is the aim of our two above mentioned solutions. We have tone mapped the video sequence ModernFlat with different TMOs as well as with our two solutions.

We have extracted three frames (number 5, 50, 70) from these results, as shown in Fig. 8. The HDR key values of these three frames are 0.005, 2.2236 and 0.0211 respectively. Consequently, from these values, the second frame should be brighter than the others. Contrary to what we would expect in the LDR results from the existing TMO solutions, the first and third frames are brighter than the second. Only Mantiuk’s et al. technique (Fig. 8(d)), preserves the relative brightness between those frames.

Our solutions relying on Reinhard’s et al. technique are presented in Fig. 8(f,g,h). Fig. 8(f) shows the results with our first solution (Equation (5)). Although the relative brightness is preserved, the first and third frames are very dim. Our enhancement (Equation (6)), provides better results as shown in Fig. 8(g). With our enhancement method, we trade a small loss of relative brightness for a better visual perception. Fig. 8(h) provides results of our second solution (Equation (8)) which are similar to those obtained with our first solution.

Finally, Fig. 8(i) is the ModernFlat sequence results for Tumblin & Rushmeier¹¹ operator which are compared to our second solution (Equation (8)) when adapted to this operator (Fig. 8(j)). The relative brightness is also preserved with our solution while Tumblin & Rushmeier technique fails to preserve it.

5. CONCLUSION

This paper introduced a workflow to preserve the temporal coherency of a tone mapped HDR video sequence. Our objective was to preserve the consistency of perception of an object as well as the relative temporal brightness. Although other techniques already dealt with temporal coherency,³⁻⁵ they only aim at solving the flickering issue. We proposed two different solutions for our workflow.

Our first contribution is built on Reinhard’s⁶ et al. technique. It efficiently preserves the overall perception of an object throughout a video as well as its relative brightness. The method finds a compromise between relative brightness coherency and loss of information in the scene thanks to a user defined parameter.




Figure 8: Images of 3 frames (number 5, 50, 70) from the ModernFlat video sequence provided by different TMOs. The upper image corresponds to frame 5, the middle to frame 50 and the last to frame 70. HDR images are given in false color: red for the maximum luminance of the HDR video and blue for its minimum. The model used in this sequence is courtesy of Jay.

Our second contribution adapts to any tone mapping operator. Only the relative brightness coherency is preserved in this solution but the same user-defined parameter (offset) is used to avoid loss of information. As it adapts to any TMO, our solution is capable of addressing any application.

Both solutions are based on the principle that video tone mapping should use the luminance range of the whole HDR video rather than those of the frames separately.

In future works, we would like to address the coherency of the perception of an object over time while considering its motion. Relative brightness coherency should also take into account change of illumination conditions as well as change of scene.

REFERENCES

- [1] Devlin, K., Chalmers, A., Wilkie, A., and Purgathofer, W., "Tone Reproduction and Physically Based Spectral Rendering," *Eurographics 2002* (2002).
- [2] Durand, F. and Dorsey, J., "Interactive tone mapping," *Eurographics Workshop on Rendering* (2000).
- [3] Kang, S. B., Uyttendaele, M., Winder, S., and Szeliski, R., "High dynamic range video," *ACM Trans. Graph.* **22**, 319–325 (July 2003).
- [4] Ramsey, S., III, J. J., and Hansen, C., "Adaptive temporal tone mapping," *Computer Graphics and Imaging - 2004* (3), 3–7 (2004).
- [5] Mantiuk, R., Daly, S., and Kerofsky, L., "Display adaptive tone mapping," *ACM Transactions on Graphics* **27**, 1 (Aug. 2008).
- [6] Reinhard, E., Stark, M., Shirley, P., and Ferwerda, J., "Photographic tone reproduction for digital images," *ACM Trans. Graph.* **21**, 267–276 (July 2002).
- [7] Adams, A., [*The Print: The Ansel Adams Photography Series 3*], Little, Brown and Compagny (1981).
- [8] Kuang, J., Yamaguchi, H., Liu, C., Johnson, G. M., and Fairchild, M. D., "Evaluating HDR rendering algorithms," *ACM Transactions on Applied Perception* **4**, 9–es (July 2007).
- [9] Ledda, P., Chalmers, A., Troscianko, T., and Seetzen, H., "Evaluation of tone mapping operators using a high dynamic range display," in [*ACM SIGGRAPH 2005 Papers*], *SIGGRAPH '05*, 640–648, ACM, New York, NY, USA (2005).
- [10] Yoshida, A., "Perceptual evaluation of tone mapping operators with real-world scenes," in [*Proceedings of SPIE*], **5666**, 192–203, SPIE (2005).
- [11] Tumblin, J. and Rushmeier, H., "Tone reproduction for realistic images," *Computer Graphics and Applications IEEE* (1993).
- [12] Pattanaik, S. N., Tumblin, J., Yee, H., and Greenberg, D. P., "Time-dependent visual adaptation for fast realistic image display," in [*Proceedings of the 27th annual conference on Computer graphics and interactive techniques*], *SIGGRAPH '00*, 47–54, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA (2000).
- [13] Ferwerda, J. A., Pattanaik, S. N., Shirley, P., and Greenberg, D. P., "A model of visual adaptation for realistic image synthesis," in [*Proceedings of the 23rd annual conference on Computer graphics and interactive techniques - SIGGRAPH '96*], 249–258, ACM Press, New York, New York, USA (1996).
- [14] Kuang, J., Johnson, G. M., and Fairchild, M. D., "iCAM06: A refined image appearance model for HDR image rendering," *Journal of Visual Communication and Image Representation* **18**, 406–414 (Oct. 2007).
- [15] Hunt, R., [*The Reproduction of Colour*], The Wiley-IS&T Series in Imaging Science and Technology, John Wiley & Sons (2005).
- [16] Panorama available at <http://www.openfootage.net/?p=144>.