Preserving color fidelity in real-time color image compression using a ranking naturalness criterion

Marina Nicolas, Fritz Lebowsky ; STMicroelectronics; Grenoble, France

Abstract

*High-end PC monitors and TVs continue to increase their native display resolution to 4k by 2k and beyond. Subsequently, uncompressed pixel amplitude processing becomes costly not only when transmitting over cable or wireless communication channels, but also when processing with array processor architectures. This paper follows a series of papers we presented earlier on a 4*4 block-based memory compression architecture for text, graphics, and video using a multi-dimensional vector representation with context sensitive control of visually noticeable artifacts. A key feature in the system is the sorting by magnitude of pixel amplitudes. To increase the compression ratio and simultaneously alleviate the limitation on block size, we analyze to which extent the sorting orders can be predicted and we consequently propose new schemes to transmit them efficiently. Depending on the compression ratio, the new cost function defined can be considered as a no-reference or reduced-reference ranking naturalness criterion. We show how pertinent our approach is to additionally correct specific visually noticeable compression artefacts thanks to its adaptive pixel positioning mechanism. Finally, we also provide hints on how to extend this new philosophy to support the optimization of future scalable architectures for transcoding or rendering on high quality displays.*

Introduction

From a simple point of view, color fidelity can be considered proportional to the number of bits per color component. However, as display size increases towards covering the entire visual field or even beyond, current color amplitudes with limited 8 to 10 bit depths can no longer ensure a smooth reproduction of subtle gradients. Consequently, higher resolution comes with the increase of the number of bits per color component, towards high dynamic range (HDR) formats. All these contributions to favor enhanced home theater experiences however aggravate overall cost of image processing devices, due to the increased image processing bandwidth in video applications.

We are thus interested in developing image compression architectures to cope with memory bandwidth requirements and alleviate the size of on-chip local image data storage. Compared to well-known image compression methods such as JPEG or MPEG standards, which can achieve high compression factors, we aim at low compression factors in the range between 1.5 and 4. However, the underlying architectural challenge arises from an implementation at a fraction of the cost of well-known compression methods. Furthermore, the targeted visual performance is also more challenging as this new compression should be visually lossless, whereas it has to operate on images that are often decompressed from MPEG streams, without interacting with potential MPEG residual artefacts.

A first proof of concept and successful implementation, the Parametric Functional Compression (PFC), has been described in [1]. The architectural simplicity was based on parametric exploration of a nonlinear domain using sorting of pixel amplitude values within a block of 4x4 pixels. The nonlinear system behavior enables randomized spreading of residual error amplitudes. Such spatially and temporally randomized error amplitudes are less visually noticeable as long as medium and high error amplitudes remain sparse. Without significant increase of algorithmic complexity, a new error minimization strategy improved PSNR rating up to 12 dB [2]. In [3], we introduced a complementary element based on vectorized linear interpolation (VLI), opening the door to new orientation-guided compression schemes.

The analysis results on the probability of orientation and curvature (POC) pairs associated to the pixels and representing the local structural content in [3] incline us to especially focus on amplitude correlations between the individual pixel and its direct neighbors. This brings the scope of the current paper on the ranking part of the PFC. From a principle point of view, ranking is the sustaining stone of our system so that lossless ranking transmission was first considered as mandatory. As a matter of fact, preservation of pixel positions in any video processing algorithm seems to be an evident constraint for an appropriate fidelity to the original image. Still, corroborating with on our past observations in [4], we find it worth to investigate how far pixels can be shuffled without hindering the understanding of local structural content.

Simultaneously, we are interested in finding ways to further increase the compression ratio without compromising on the image quality. Most PFC key elements have been studied to this purpose in our previous works. Ranking is a significant contributor to the bit budget and therefore deserves our full attention.

In the first part of the paper, we recall the fundamentals of the PFC, zooming in the ranking part. Then, we consolidate our past observations to assert a sufficient predictability of ranking orders that enables us to propose, in the following section, a partial ranking transmission scheme and associated ranking reconstruction. Next we give objective measurements validating the visual performance of the new method. One of the observed unforeseen benefit allows us to derive a similar strategy on the encoder side. We end up with some conclusions on the criteria used for optimization, namely a reduced-reference ranking naturalness criterion on decompression side or equivalently a reference ranking fidelity criterion on the encoder side.

Fundamentals of parametric functional compression (PFC)

The ever increase in resolution of high-end PC monitors and TV displays, together with the deployment of High Dynamic Range makes the processing of video streams very challenging in terms of memory bandwidth requirements, size of on-chip local image data storage, power consumption, and cost. Parametric functional compression (PFC) aims at efficiently leveraging these on-chip VLSI design constraints while maintaining visibility of artifacts below visual threshold. From a system architecture point of view, the ultimate goal of PFC is accomplishing a flexible image data compression method operating in time domain that considers perceptual models enabling enhanced visual quality control of the encoding scheme and facilitating a highly simplified decoding scheme. The current system, based on the method presented in [1], complemented with an efficient inter-component scheme, offers compression factors from 2 to 2.5 depending on the storage format, without compromising on image quality. Motivated by these results, we keep on further exploring information correlations and predictability to increase the compression ratios. To do this efficiently, we analyze the contribution of the different PFC key elements and their interactions to focus on the limiting elements. After optimizations on the function approximation, the quantization and error minimization strategy, ranking now draws our attention, as a next stage in a system where lossless ranking transmission was up to now considered as the sustaining stone.

Overview of key components involved in parametric functional compression

Figure1 illustrates the key elements involved in the parametric functional compression introduced in [1]. In the sorting/ranking component, pixel amplitudes are distributed in 2 sets of values (the “clusters”) and ordered in ascending order in each set. The ordering enables operating in a new vector space composed of monotonic functions. The functional approximation block estimates the optimal representation of the ordered values in a parametric function through error minimization [2]. Using successive linear segments for the parametric function turns out to be very efficient both from a budget point of view and from a visual performance point of view. Indeed, without any extra quantization, the segment lengths, the block extreme amplitudes and the segment slopes describing the functions will never require more than 92 bits to approximate a component block of 16 10-bit amplitude values. Furthermore, these parameters can undergo adaptive inter-cluster encoding in the quantization component with hardly any visible impact. This ensures a representation in less than 80 bits for the same 16 10-bit component block.

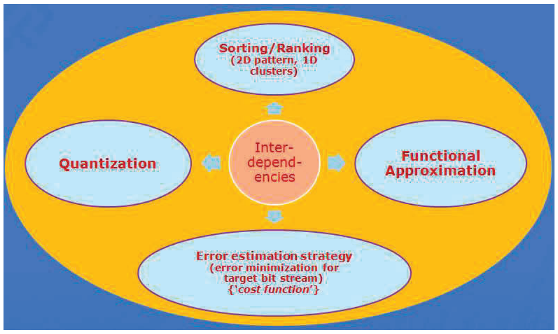

Figure 1. Overview of PFC components.

The obvious challenge in the PFC is the optimization of the component interdependencies. The choice of the pattern for the clustering is highly dependent on the error minimization strategy. Functional approximation and quantization components highly interact to ensure that the transmitted parameters will provide the best reconstruction of original values at the decompression side. The error estimation strategy component controls the functional approximation and the quantization, but also the distribution of the bit budget among the different parameters to be transmitted.

The work presented in this paper focused on the sorting/ranking component and the error estimation strategy component.

From the decompression diagram shown on Figure 2, we can list the content of the compressed block containing the parameters for the reconstruction:

- Some configuration parameters
- Inter-cluster encoded parameters describing the monotonic approximation
- Original pixel positions corresponding to the ordered pixel amplitudes
- Pattern type used for the clustering

Parameters related to configuration and pattern type fill a minor portion of the overall bit budget. Parameters for the monotonic approximation are already quantized and encoded very efficiently. So far, original pixel positions are transmitted in a lossless way, which makes the ranking an important contributor to the bit budget. This is one major reason why ranking draws our attention in this paper.

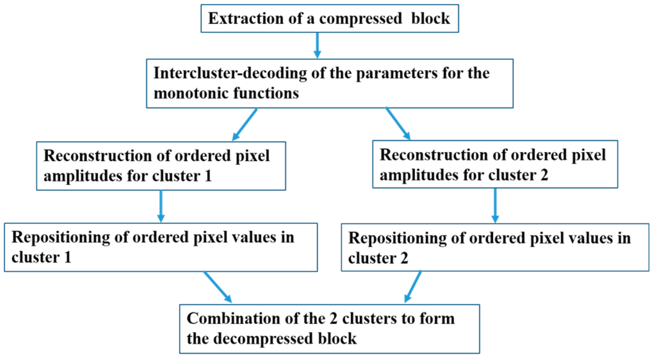

Figure 2. Reconstruction of approximated block at the decompression side.

A key feature in parametric functional compression: sorting by magnitude of image amplitude

In our previous work [2], we demonstrated the capability of the ranking of amplitude values to lead to an efficient parametric representation, but also to automatically enable randomized spreading of coding errors, which most often lowers visibility of residual errors.

Now, let us consider the contribution of the ranking in the overall bit budget. A lossless transmission of M pixel positions/ranking orders requires ceiling(log2(factorial(M))) bits. Consequently, thanks to the additional exploitation of inter-component correlations (outside the scope of this paper) the PFC finally achieves a compression factor of 2 on YcbCr 420 10 bit format. Increasing the block size, especially in a 4K context where the average granularity of information gets coarser, would help getting more efficient monotonic functions. Unfortunately, as shown on Figure 3, the increase of the relative contribution of ranking positions in the overall budget cancels the decrease of the contribution from the monotonic functions, meaning that the achievable compression ratio hardly increases with the block size. At the end, the achievable compression ratio would hardly exceed 2.5 whatever the block size chosen.

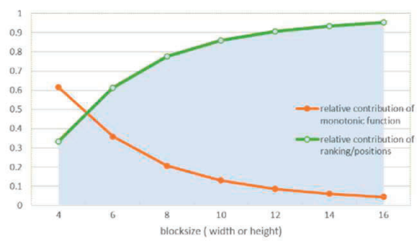

Figure 3. Evolution of the contribution of ranking when the block size increases.

Visually lossless reconstruction of ranking orders

To reach higher compression ratios, it seems then mandatory to alleviate the limitations on block size linked to the ranking budget. Therefore, we are interested in replacing the lossless transmission of ranking orders by a visually lossless reconstruction scheme, that would allow for a partial transmission of the costly ranking/positions information.

This problem can be expressed as: knowing the decompressed ordered pixel amplitudes for both clusters and the position associated to some of these values in the original block, can we compute the positions of the remaining values in a way to minimize the visual impact of wrong positioning? We choose to transmit the ranking of every second pixel following a quincunx pattern, whatever the original distribution in 2 clusters performed by the PFC. In the 4*4 block illustrated in Figure 4, we would know the values of the red/blue pixels and we then need to associate one blue/red position to each of the eight remaining candidate values.

1	1	**2**	2
3	3	**4**	4
5	5	**6**	6
7	7	8	8

*Figure 4. Indexes of pixel positions in PFC 4*4 block.*

The ranking orders transmitted are still computed in the original cluster distribution. Configuration data along with the partial ranking allows a correct positioning of M decompressed pixels in the block, where 2*M is the total number of pixels in the block. In the most favorable case where the pixels have been distributed following a quincunx pattern, this partial transmission requires ceiling(log2(factorial(M))) bits saving half of the ranking budget cost. In the least favorable case, the budget required is ceiling(2*log2(factorial(M) / factorial(M/2))). As an example, for blocks of 4*4 pixels, the lossless transmission of ranking orders would cost 32 bits, the partial transmission would cost from 16 to 22 bits only. Such a partial transmission scheme would significantly increase the compression factors, leading to factors in the range of 3.1 to 3.6 for most of the block sizes considered in Figure 3.

Predictability of ranking orders

In the prospect to reduce the bit budget required for the transmission of ranking orders, we are interested to see how far ranking orders could be only partially transmitted and reconstructed at the decompression side, and thus how far they can be predicted. Ranking correlation has been used by Harwood et al in [5] to characterize textures. Directional RANK-strength statistics have also been proposed by Patel et al in [6] to classify and segment textures from Brodatz's database. We wonder how far these statistical results can be adapted to any video content. Following our partial transmission scheme above, we aim to investigate whether the choice for one of the other candidate values at one pixel position to be filled can be given from its neighboring known pixel amplitude values.

Predictability of ranking orders in natural images

The vectorized linear interpolation presented in [3] took advantage of the fact that a significant portion of pixels in a natural image can be approximated without any visual impact by the averaging of two neighboring pixels. As a recall example, Figure 5 shows one color component of a test image and the associated error maps for error thresholds from 1 lsb up to 4 lsbs. Yellow positions correspond to pixels that can be averaged by a pair of neighboring pixels with an error up to the error threshold (VLI pixels) and pixels are marked as red when the averaging error exceeds the allowed threshold.

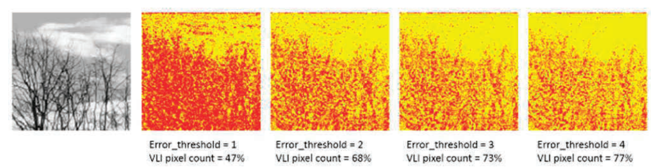

Figure 5. Test image and associated error maps showing pixels that can be approximated by the average of two neighboring pixels.

From this result, it seems meaningful to assume that the probability for a given pixel amplitude to be in the range defined by the lowest amplitude and the highest amplitude among its eight direct neighbors will be rather high on natural images. Following the quincunx pattern, we are interested in checking whether this assumption remains valid if we restrict the set of neighbors to the top, bottom, left and right pixels. Taking for example the pixel value at position 4 in the red cluster shown on Figure 5, we would like to know whether there is a high probability for this value to lie between the minimum and maximum value among values at positions 2, 3, 4 and 6 from the blue cluster.

We compute the probability for a pixel to be in the range defined by its top, bottom, left and right pixels for a variety of 4K content as shown on Figure 6.

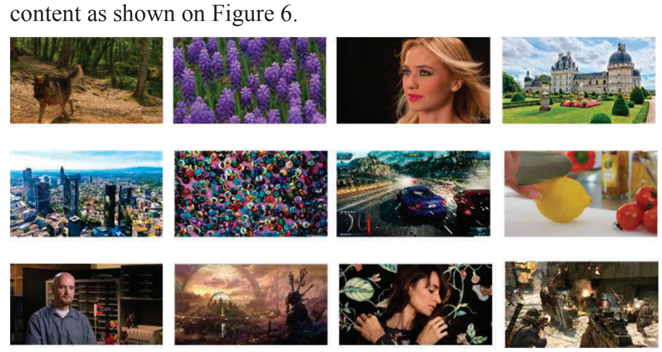

Figure 6. Set of 4K images used for our statistical analysis.

Figure 7 shows that a very large majority of pixels obey to this constraint, even in highly textured content. As a reference, a random distribution of the ranking orders would only lead to 60% of the pixels obeying this constraint.

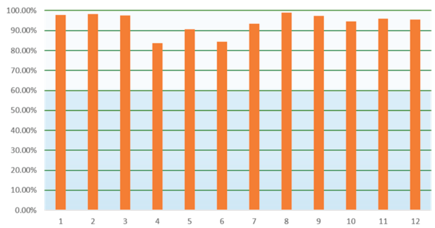

Figure 7. Percentage of pixels obeying the range constraint.

Predictability of ranking orders considering the block granularity of the processing

In our search towards an optimized solution, we also need to consider the block granularity of the processing as follows:

(1) The range taken into account at a given position has to be restricted to the range defined by the available neighbors in the block. For example, we will consider the probability of the amplitude at red position 1 to lie between the minimum and maximum value among values at blue positions 1, 2, 3. Making projections from the results in Figure 7, we can approximate the impact of the block borders. For the least favorable content, there is still a 73 % chance for a pixel to in the range associated to its available neighbors in a 4*4 block. Over the images tested, 83% of the pixels obey to this new constraint. It seems then meaningful to exploit this predictability to give preferences or penalties to one of the other candidate values when associating a position to each candidate value.

(2) Minimum and maximum (extreme) values in the block obviously will not obey the constraint, except for homogeneous blocks. On the other hand, assuming that most pixels are neither minimal nor maximal local values as suggested by the results in Figure 7, we can guess that minimal and maximal values are likely to be pushed towards the block borders. Therefore, we can probably still use separate predictability of positions for minimal and maximal amplitude values. Experiments on the set of images on Figure 6 shows that statistics for minimal values and maximal are very close. So we will refer indifferently to "extreme amplitude" positions. Results also do not evolve significantly from one image to the other. Figure 8 shows the probability of the extreme amplitudes to be at a given position in a 4*4 block. More than 40% of minimal and maximal values are located on one of the 4 corners and less than 10% on one of the 4 center pixels. This pattern compensates for a somewhat lower amount of information when block borders are considered, due to the fact that extreme values are by definition outside the range of neighboring values and allows us to give preferences or penalties to given (positions, extreme values) association pairs.

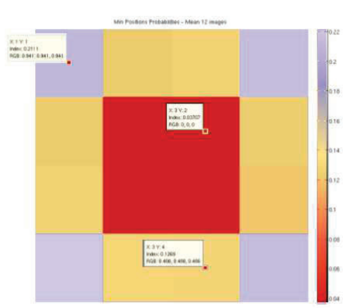

Figure 8. Probability for an extreme amplitude to be at a given position.

In conclusion, based on our statistics on (position, amplitude) associations in natural images, we propose to partly (transmit position, amplitude) associations and reconstruct the other ones at the decompression side.

Reconstruction of missing (position,amplitudes) associations

In our prospect to best distribute positions among a set of candidate amplitudes, we propose a probability model based on the predictability results above. We namely compute two matrices of size N*N where N is the number of candidates/positions:

(1) A first matrix *Ppos* where *Ppos(k,l)* gives the estimated probability that the position k is filled with candidate value Cl, Cl being the value at rank l in the ordered set of candidates.

(2) A second matrix *Prank* where *Prank(k,l)* gives the estimated probability that the candidate value Ck is at the position l, Cl being the value at rank l in the ordered set of candidates.

The reconstruction of the missing (position,amplitudes) associations consists of the following stages :

(1) At the encoder side, a decision is made to decide which set of positions/ranks to transmit (red or blue pixels in Figure 4) and which set will need to be reconstructed (blue or red pixels in Figure 4).

(2) *Ppos* and *Prank* matrices are computed separately, but following very similar rules which enable mutualizing most of the computational cost.

(3) Both matrices are combined to form a new N*N matrix *Pred* that will guide our strategy for the final distribution of positions to the candidates.

(4) Final distribution is performed following different strategies. Examples of strategies used are to (1) distribute first positions/ranks for which there is a very high adequacy prediction, or to (2) make decisions on unwanted positions/ranks first to reach best compromises and avoid risky associations.

Choice of the set to be transmitted/reconstructed

A key element for the computation of *Ppos* and *Prank* being the consideration of local ranges defined by the neighboring pixels, it is relevant to avoid the reconstruction of positions associated to the extreme values of the block. Therefore, when both the minimum and maximum values are located on the same quincunx grid, this grid is chosen as the set of positions transmitted. When minimum and maximum values lie on different grids, a distance estimation from one candidate set corresponding to one grid to the total range associated to the other one guides the choice for the set to be transmitted, namely the most risky one.

Computation of the Pos/Prank matrices

For each Cl candidate value, its probability to be at a given position k is estimated by the function profile shown on Figure 9:

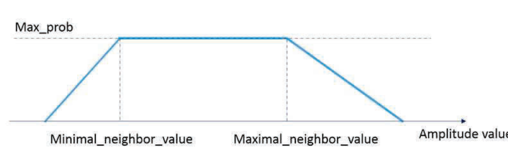

Figure 9. Ppos/Prank profile computation.

When the candidate value Cl falls in the range defined by its neighbors at a given position k [*Minimal_neighbor_value Maximal_neighbor_value*], we consider that Cl is among the most appropriate candidates at this position, thus setting *Ppos(k,l)* to its maximum value *Max_prob*. Obviously *Max_prob* depends on the number of candidates falling in the range, as these best-fitting candidates have to share the probability to be at the position k, and as the probability that the real amplitude at position k is outside the range gets lower when the number of candidates already present in the range is higher. Once the probability for a given position to be filled with a candidate in its local range has been computed, the probability for the position to be filled with a candidate outside the range is distributed among these remaining candidates. The decay rate will depend on the number of remaining candidates.

Prank matrix is defined following similar assumptions. Positions for which a candidate value Cl falls in the range of neighbors are considered as most appropriate positions for Cl and the normalization takes into account the number of best-fitting positions for the candidate. The probability for a given candidate to be outside this set of preferred positions is distributed among the remaining positions, taking into account a distance criterion to the corresponding local neighboring ranges.

Combination of the Pos and Prank matrices

At this stage, what we seek for is a synthesis of the prediction we can make on (positions,candidates) associations based on local information, taking into account the one-to-one nature of the problem. The multiplication of both matrices leading to a new matrix *Pred* gives an adequate representation where easy to fill positions, easy to position amplitudes, critical amplitudes or critical positions stand out. At this level, the *Pred* matrix is slightly modulated to take into account the preferred positions of extreme candidate values and to favor the maximization of the quality metric defined in the next paragraph.

Final distribution of (position,candidate) associations.

The final distribution of positions to each candidate can be performed following different strategies based on the quality of available local information gathered in the *Pred* matrix. For each possible distribution where a position k is filled with candidate Cl, we can define *select(k)*=l. A quality metric for this distribution can be computed as the cumulative sum of *Pred(k, select(k))* over all positions. A reference optimal distribution would then seek for a maximization of this quality metric.

To avoid this expensive maximization stage, we define two recursive distribution schemes as follows:

(1) The first strategy consists in associating first (position k,candidate l) pairs for which *Pred(k,l)* is very high whereas $Pred(k, i \neq l)$ and $Pred(j \neq k, l)$ are rather low. Basically, we handle easy cases first where there is little doubt to secure some key pixels positions. Each time a new association is set, the corresponding probability in the *Pred* matrix is distributed over the remaining elements, leading to an updated *Pred* matrix with one line and one column less. The process keeps on handling the easiest position/ candidate among the remaining ones until the dimension of the matrix becomes smaller than a number of predefined elements. Maximization of the quality metric on the reduced *Pred* matrix is performed for these last elements, avoiding to have loser positions or loser candidates. As a matter of fact, the last candidates are likely to end up at positions they do not prefer, but we take care that they do not end up at their worst position.

(2) The second strategy consists in handling difficult cases first, so that an unwanted position gets filled with its preferred candidate or so that an unwanted candidate settles at the least annoying place. To achieve this, we select the line or column in the *Pred* matrix with the lowest sum over the line or column. The computations involved are very similar to the ones used in the first strategy where maximum operators would be replaced by minimum ones.

Results of the reconstruction

When we estimate the visual quality of the final decompressed images with this partial ranking transmission/ reconstruction, both strategies give very similar results from a PSNR or SSIM [7] point of view. The increased performance provided by the reference optimal distribution maximizing *Pred* is not worth the extra complexity. Consequently, we do not report performances of the different strategy separately and reported numbers refer to the first strategy.

The SSIM metric is used as an indicator of the global quality to validate the fact that increasing compression factors following our approach leads to graceful degradation. Figure 10 shows the impact of partial ranking transmission/reconstruction on the set of images from Figure 6. Even on the most critical content, the degradation is hardly (visually) noticeable.

Figure 10. Impact of partial ranking transmission/reconstruction.

Besides the global benefit of the method, we would like to check the annoyingness of potential artefacts on isolated pixels. The quality of the reconstruction is highly linked to the amount of relevant information in the *Pred* matrix. Fundamentally, we cannot rule out the occurrence of isolated blocks where the real ranking is not natural and cannot be predicted. The higher risk is then to end up with a pair of inverted pixels in the final image. As long as the inversion is coherent from one component to the other and does not lead to isolated pixels with wrong colors (visible hue changes), this inversion has hardly any visual impact on a 4K content. Furthermore, the PFC contains a mechanism to exploit inter-component correlations, preventing most occurrences of potential wrong colors that could be caused by a partial ranking transmission.

We could imagine a ranking fidelity metric quantifying the ranking error done by the reconstruction, comparing the local rank in the original image to the one in the decompressed image. On the castle image, about 95 % of the pixels settle at the correct place. However, most position inversions remain fully acceptable. What visually matters is the rank of the individual pixel among its direct neighborhood. In this sense, the reconstruction proposed basically relies on a non-reference ranking metric to ensure coherency between the two grids, regardless of the pattern used to distribute the pixels for the computation of the functional approximation.

At a few places as shown on Figure 11, it turns out that our method leads to more pleasing results than the original method with a lossless transmission of ranking orders. The limitation on the middle image has already been reported in our previous paper [2] and is due to the fact that equal values within a block happen to be represented by different values in the decompressed image. In this case, the M different values are associated to M different positions arbitrarily. Our partial reconstruction scheme on the right image restores the most natural distribution of positions leading to a lower visibility of the artefact.

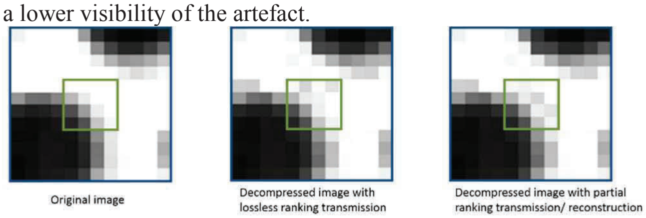

Figure 11. Visual impact of ranking reconstruction following ranking naturalness criterion.

Using the ranking naturalness criterion on the encoder side.

In the previous section, we showed the benefit of following a ranking naturalness criterion to cope with ambiguous ranking cases. The correction shown on Figure 11 was performed on the decompression side, on the grid for which ranking orders were not transmitted. In practice, the limitation described can be found as well on the grid for which the ranking orders are transmitted. Therefore, it seems meaningful to control the ranking orders transmitted with the same kind of ranking criteria. The *Ppos*, *Prank* and *Pred* matrices are then limited to the number of ambiguous elements.

The major difference is that the non-reference natural profile defined on Figure 9 is replaced by a full-reference local ranking profile depending on the real local context, as shown on Figure 12. At most positions, the correction will be guided by the middle profile. The two additional profiles will control the positioning of minimum and maximum values. Our non-reference naturalness ranking criterion now becomes a full-reference fidelity ranking criterion.

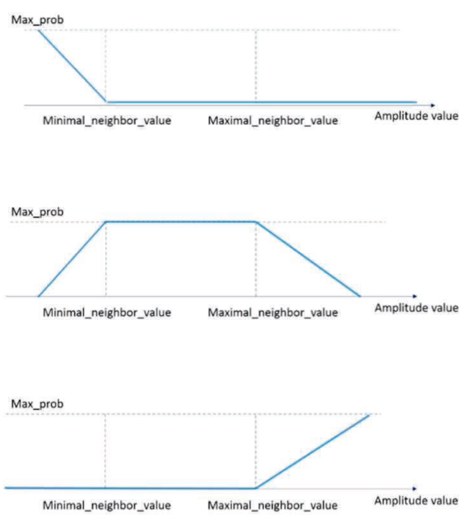

Figure 12. Ppos profile on the encoder side.

Taking into account the different configuration possibilities of the functional approximation, ambiguous ranking at the encoder side can only happen in one of the following situations:

(1) 2 identical values being compressed/decompressed to two different ones
(2) 3 identical values being compressed/decompressed to three different ones
(3) 2 couples of identical values being compressed/ decompressed to 2 couples of different ones

In all these cases, the number of possible shuffling is so limited that we opt for a direct maximization of the *Pred* matrix.

Conclusion

Parametric functional compression has proven to be a very promising approach to leverage chip design constraints, dealing with ever increasing bandwidth requirements and the size of local image data storage. In this paper, we propose a new partial ranking transmission scheme and a ranking reconstruction method that allows us to successfully alleviate the limitations on block size and consequently give room for a significant increase in the compression ratio, with hardly any noticeable impact for the human visual system. Our method relies on the high predictability of local ranking distribution and our reconstruction is therefore guided by a ranking naturalness criterion. An equivalent ranking fidelity criterion is used on the encoder side to improve the visual quality. Due to its simplicity, the ranking criterion allows a real time, perceptually oriented, control of locally acceptable errors in system-on-chip solutions embedding compression units. In the near future, we will further exploit the promising results presented to derive a flexible compression system targeting compression factors largely above 3. We also further plan to explore alternative uses of this new perceptual criterion for the optimization of future scalable architectures.

Acknowledgment

We would like to thank Kevin Picot, a former trainee at ST Microelectronics for carrying out deep statistical analysis on the positions of local minimum and maximum amplitudes, with a very refreshing state of mind.

This project was funded by NANO2017 framework.

References

[1] F. Lebowsky, "Optimizing color fidelity for display devices using contour phase predictive coding for text, graphics, and video content," in Proc. SPIE 8652, Color Imaging XVIII: Displaying, Processing, Hardcopy, and Applications, 86520X,2013

[2] F. Lebowsky and M. Nicolas, "Preserving color fidelity for display devices using scalable memory compression architecture for text, graphics, and video," in Proc. SPIE 9394, Human Vision and Electronic Imaging XX, pp. 939405, 2015.

[3] M. Nicolas and F. Lebowsky, "Optimizing color fidelity for display devices using vectorized interpolation steered locally by perceptual error quantities," in Proc. SPIE 9395, Color Imaging XX: Displaying, Processing, Hardcopy, and Applications, pp. 939502, 2015.

[4] F. Lebowsky and M. Nicolas, "Color discrimination problems in digital TV systems,"in Proc. SPIE 5667, Color Imaging X: Processing, Hardcopy, and Applications,2005.

[5] D. Harwood, M. Subbarao and LS Davis,"Texture classification by local rank correlation", in Computer Vision, Graphics and Image Processing, vol32,pp.404-411, 1985.

[6] D. Patel and T.J Stoneham,"Texture image classification and segmentation using rank-order clustering", in proc. 11th TCPR,Den Haag,1992.

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

Author Biography

Marina Nicolas received her MS in telecommunication system engineering from Grenoble National Polytechnical Institute, France in 1991. She entered the world of image processing working on LCD information panels at Stanley Electric Company in Yokohama, Japan. From 1992 to 2000, she worked as a research and development engineer developing multimedia processing chains and algorithms for video analysis and video enhancement at Philips LCD lab, Philips Advanced System lab in Eindhoven, the Netherlands and Philips TV unit in Singapore. She joined ST Microelectronics in Grenoble, France in 2001, where she is focusing on image quality improvement, compression algorithms and video quality metrics.

Fritz Lebowsky received his MS (1985) and PhD (1993) in electrical engineering from the Technical University of Braunschweig, Germany. He began his professional career as a research and teaching assistant at the Institute of Telecommunications of the Technical University of Braunschweig in 1985. From 1991 he worked as a research and development engineer in the field of digital video processing at Micronas in Freiburg, Germany. In 1995 he joined Thomson Consumer Electronics Components in Meylan, France, as a development engineer modeling video processor networks as well as digital acquisition sub-systems for DVD ROM drives. In 2000 he joined the Imaging and Display Division of STMicroelectronics Inc. in San Jose, CA, developing advanced display engines for the PC flat panel monitor market. Since 2004 he is with STMicroelectronics in Grenoble, France, working on image quality improvement for consumer TV products.