

Fast and Lightweight Rate Control for Onboard Predictive Coding of Hyperspectral Images

Original

Fast and Lightweight Rate Control for Onboard Predictive Coding of Hyperspectral Images / Valsesia, Diego; Magli, Enrico. - In: IEEE GEOSCIENCE AND REMOTE SENSING LETTERS. - ISSN 1545-598X. - 14:3(2017), pp. 394-398. [10.1109/LGRS.2016.2644726]

Availability:

This version is available at: 11583/2671723 since: 2017-05-23T11:18:38Z

Publisher:

Institute of Electrical and Electronics Engineers Inc.

Published

DOI:10.1109/LGRS.2016.2644726

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Fast and Lightweight Rate Control for Onboard Predictive Coding of Hyperspectral Images

Diego Valsesia, and Enrico Magli

Abstract—Predictive coding is attractive for compression of hyperspectral images onboard of spacecrafts in light of the excellent rate-distortion performance and low complexity of recent schemes. In this letter we propose a rate control algorithm and integrate it in a lossy extension to the CCSDS-123 lossless compression recommendation. The proposed rate algorithm overhauls our previous scheme by being orders of magnitude faster and simpler to implement, while still providing the same accuracy in terms of output rate and comparable or better image quality.

Keywords—Rate control, predictive coding, hyperspectral images

I. INTRODUCTION

Hyperspectral imaging from space-borne spectrometers has a multitude of applications, such as terrain analysis, material identification, military surveillance, etc. The ever growing spectral and spatial resolution of such instruments poses challenges in handling such wealth of information. In particular, onboard compression is of paramount importance to overcome the limited downlink bandwidth. This is a very active area of research as it poses peculiar challenges, not encountered elsewhere. In fact, onboard compression algorithms face strict complexity limitations due to constraints on the payload hardware. Several solutions based on different techniques have been proposed, such as low-complexity spatial [1] and spectral transforms [2], distributed source coding [3], compressed sensing [4], [5], and predictive coding [6]–[8]. Predictive coding is one of the most popular solutions, as it enables low-complexity, high-throughput solutions, and excellent rate-distortion performance. One of the most recent recommendations by the Consultative Committee for Space Data Systems (CCSDS) is a lossless compression algorithm based on predictive coding [9] and an extension to lossy compression is ongoing. Lossy predictive coding is typically operated in the so-called near-lossless mode, where the maximum absolute error on the reconstructed pixels is bounded by a constant. However, near-lossless compression has some drawbacks, the most notable being a variable output rate dependent on the image content. In our previous works [10], [11] we showed that it is possible to perform simultaneous rate and quality control with a complexity compatible with resources available onboard of spacecrafts. In this letter, we propose a significantly improved rate control algorithm with respect to [10]. The main contribution is that the new algorithm is extremely fast and lightweight with respect to [10] and this allows mainly to *i*) save hardware resources that could be dedicated to other purposes (or not used thus saving power), compared e.g. with the solution in [11]

which requires a dedicated board, running in parallel to the compressor; *ii*) enable very high throughput implementations, while the original algorithm may have caused bottlenecks in such scenarios because of some complex rate-distortion optimization operations. We test the proposed algorithm by extending the CCSDS-123 recommendation, but the method is general and can be applied to any predictive encoder.

II. RELATED WORK

This section reviews the rate control algorithm originally proposed in [10] and further developed in [11]. The goal of the algorithm is to control the output rate of a predictive encoder of hyperspectral and multispectral images, under low complexity and memory constraints. To be best of our knowledge this is the only work on rate control for onboard predictive coding. However rate control algorithms are often used either onground [12] or onboard [13] in transform coding schemes. The algorithm in [10] selects quantizers operating on the prediction errors in predefined spatial and spectral regions. The rate control algorithm works on a slice-by-slice basis, where we call “slice” a predefined number of lines with all their spectral channels. Each band of each slice is divided into nonoverlapping 16×16 blocks. The rate control algorithm assigns a quantization step size for each of the blocks in each spectral band in order to meet the target rate with the lowest distortion. This is done via a two-stage process consisting of:

- Training stage: unquantized prediction residuals in each block are modeled as realizations of independent and identically distributed Laplacian random variables. The sample variance is estimated in order to predict the rate as function of the quantization step size. The rate function is stored in a lookup table (LUT) as it is too complex to implement directly.
- Optimization stage: a greedy optimization algorithm is employed to select the quantization step sizes that allow achieving the target rate with the lowest distortion.

The training stage requires to run the predictor twice for each slice, once to estimate the variance of the unquantized prediction residuals, then one more time in order to perform the actual coding which also quantizes the residuals with the step sizes computed by the rate controller. This dual prediction stage can be performed in series to the coding process, as in the original algorithm, or in parallel, as in [11], by having an independent module that works ahead of the coding process.

Furthermore, the algorithm measures the actual rate produced by encoding the slice with the computed quantization step sizes and uses this information to update the target rate for the next slices. This mode of operation has been shown to

This work was supported by the European Space Agency (ESA-ESTEC) under grant 107104. The authors are with Politecnico di Torino, Torino, Italy.




Fig. 1. High-level diagram of proposed method. A different number of operations are performed depending on the current pixel position. Pixel in position D (last pixel of the current line in the last band) performs all the operations of pixels A, B or C and runs the rate control procedure to determine the quantization step size for the next line.

effectively correct inaccuracies in the model without reducing the rate-distortion performance.

In [14] the authors simplify the rate control algorithm by eliminating the rate-distortion optimization phase and working on line-by-line basis instead of blocks. This idea is also used in this work, which however solves major drawbacks of the method proposed in [14]. The method proposed in this letter improves over [10], [11], [14] by introducing the following novel points:

- choosing just one quantization step size per spectral line, which avoids complex rate-distortion optimization algorithms (also present in [14]);
- on-the-fly estimation of the residual statistics, which avoids running the predictor twice, with significant gains in terms of speed;
- simpler arithmetic that does not involve squaring operations to compute the statistical parameters;
- simpler arithmetic speeding up access to the rate LUT;
- reduced number of LUT lookups.

III. PROPOSED METHOD

In this section we explain the main novel contributions of this work. The goal is to design a lightweight rate control algorithm that is capable of providing accurate control of the output rate of a predictive encoder while being, at the same time, simple to implement in hardware, and fast in order to enable high-throughput implementations. Fig. 1 shows a high-level diagram of the proposed method and its interaction with a predictive encoder. The main operations are described in the following subsections. In the following the term *spectral line* refers to a row of pixels with all its spectral channels. It can be noticed that the algorithm proceeds one spectral line at a time and that different operations are performed depending on the spatial location of the pixel under coding. For the sake of simplicity, the following description assumes that the pixels are read and coded in band-interleaved-by-line (BIL) order. The objective of the algorithm is to encode the current spectral line, thus producing entropy-coded versions of quantized prediction

residuals and to decide the quantization step size to be used to encode next spectral line. A single quantization step size Q is used for each spectral line.

A. Prediction and parameter estimation from residuals

The use of a single quantization step size Q of a uniform scalar quantizer for each spectral line is at the core of the algorithm. It is necessary to have a statistical characterization of the prediction residuals in each band in order to choose the value of Q yielding the desired target rate. The statistics of prediction residuals of adjacent spectral lines are highly correlated and therefore it is possible to use the statistics of the current spectral line as a prediction for the statistics of next line. We assume that the unquantized prediction residuals are realizations of i.i.d. Laplacian random variables. The Laplacian distribution ($f_R(x) = \frac{\Lambda}{2} e^{-\Lambda|x|}$) is characterized by a single parameter Λ . Notice that the Λ parameter is also related to the variance σ^2 through $\Lambda = \sqrt{\frac{2}{\sigma^2}}$. The magnitude of a Laplace random variable follows an exponential distribution of parameter Λ , and the sample mean is the unbiased maximum likelihood estimator of Λ^{-1} . However, this is not a robust estimator [15]. Replacing it with the sample median has several advantages, including robustness to outliers (the residuals are far from perfectly i.i.d. Laplacian), and simpler arithmetic (avoids division by the length N_{cols} of the line). Computing the exact median of the prediction residuals of an entire row would require a sorting operation of complexity $\mathcal{O}(N_{\text{cols}} \log N_{\text{cols}})$. A good approximation can be obtained with the median of medians algorithm [16] which breaks down a row into non-overlapping subsets of L contiguous residuals. The median of each subset is computed and then the median of medians provides an approximation of the true median of the row. This approach reduces the complexity to $\mathcal{O}(N_{\text{cols}} \log L + \frac{N_{\text{cols}}}{L} \log \frac{N_{\text{cols}}}{L})$. The parameter L is set as a compromise of several objectives. In fact, L should be large enough that the approximation to the median by means of the median of median is satisfactory. However, a large L increases the complexity of the calculation. A good design should approximately balance the complexity of the local median operations (each costing $\mathcal{O}(L \log L)$) and the complexity of the final median performed at the end of the row (costing $\mathcal{O}(\frac{N_{\text{cols}}}{L} \log \frac{N_{\text{cols}}}{L})$) for the typical width of an image, so that throughput is maximized by the absence of bottlenecks.

The coding and estimation jobs proceed together. Following the notation of Fig. 1, pixels of type A entropy-code the current quantized prediction residual of the current pixel but also store in a temporary buffer the unquantized residual. The encoder reaches a type-B pixel after $L - 1$ type-A pixels and the temporary buffer holds the unquantized prediction residuals for that subset of type-A pixels. After the computation of the unquantized residual for the type-B pixel, the median of the buffer is computed and stored in the buffer of medians. This sequence of $L - 1$ pixels of type A terminated by a type B pixel is repeated until the end of the row. The row is terminated by a pixel of type C, which computes both the median for the local subset of pixels and the median of medians. Each band z stores its own median of medians m_z . The last pixel in the

last band is of type D and it uses all the medians of medians to run the procedure to determine the quantization step size to be used for the next spectral line.

There are several important differences with the algorithms in [10], [14]. In particular, a major shortcoming of those algorithms is the need to run the predictor twice, once for estimation of the statistics and once for the actual coding. Moreover, they adopt a more computationally expensive approach to parameter estimation by computing the sample variance σ^2 of the unquantized prediction residuals, and then using it to compute Λ . This process requires some computational effort when performed with integer arithmetic and is subject to numerical approximations.

B. Assignment of quantization step size

The procedure that assigns the quantization step sizes is extremely lightweight in order to provide the fastest implementation possible. It takes as input the set $\{m_z\}_{z=1}^{N_{\text{bands}}}$ of the medians of medians for all the bands. The procedure uses a rate LUT to find the odd integer-valued quantization step size $1 \leq Q \leq Q^{\max}$ such that:

$$Q = \arg \min_{q \in \{1, 3, 5, \dots, Q^{\max}\}} \left| \sum_{z=1}^{N_{\text{bands}}} R(m_z, q) - R_{\text{target}} \right|,$$

i.e., the one providing a total rate for the line and all its spectral channels closest to the target rate R_{target} . The rate LUT implements the following formula (from [10]), which is the entropy of an i.i.d. Laplacian source quantized with a uniform scalar quantizer of step size Q and parameter $\Lambda = \frac{1}{m}$.

$$R(m, Q) = - \left(1 - e^{-\frac{Q}{2m}} \right) \log_2 \left(1 - e^{-\frac{Q}{2m}} \right) + \frac{e^{-\frac{Q}{2m}}}{\log(2)} \left[\log \left(\frac{1 - e^{-\frac{Q}{m}}}{2} \right) + \frac{Q}{2m} - \frac{Q}{m \left(1 - e^{-\frac{Q}{m}} \right)} \right] \quad (1)$$

The proposed method uses two variables, namely m and $\delta = (Q - 1)/2$, to index the rate LUT that is treated as a two-dimensional matrix storing the values of $R(m, Q)$ precomputed using Eq.(1). Notice that using δ instead of Q as index for the LUT is slightly more efficient (one fewer bit of dynamic range). This technique is faster than the one used in [10] which uses a single-variable indexing. Indeed the rate depends only on the product ΛQ and the parameter estimation algorithm of [10] directly returns Λ . However, computing the correct position in the table in such manner is inefficient because it requires to perform a multiplication every time Q or Λ are changed. Moreover, the rate function was sampled at predefined values of ΛQ and a suitable mapping of the product ΛQ provided at runtime with the closest in the LUT had to be computed. The proposed method overcomes such limitations by exploiting the fact that m is discrete-valued and with limited dynamic range as well as not requiring any index calculation.

A further optimization is included to minimize the number of LUT lookups. The procedure that determines the quantization step size for the following spectral line uses the last chosen quantization step size as a prediction for the next one. Then

Algorithm 1 Coding and Rate Control

Inputs:
 $R_{\text{target}}, L, Q^{\max}$

Initialize:
 $Q_1 = 1$

for $y = 1, \dots, N_{\text{rows}}$ **do**
 for $z = 1, \dots, N_{\text{bands}}$ **do**
 $i \leftarrow 0$
 for $x = 1, \dots, N_{\text{cols}}$ **do**
 $\hat{s}_{x,y,z} \leftarrow$ prediction of pixel $s_{x,y,z}$
 Residual $r_{x,y,z} \leftarrow \hat{s}_{x,y,z} - s_{x,y,z}$
 if $x \bmod L = L - 1$ **then**
 $\tilde{m}_i \leftarrow$ median($r_{x,y,z}, \dots, r_{x-L+1,y,z}$)
 $i \leftarrow i + 1$
 end if
 Quantized $\hat{r}_{x,y,z} \leftarrow \text{sgn}(r_{x,y,z}) \cdot \left\lfloor \frac{|r_{x,y,z}| + \frac{Q-1}{2}}{Q_y} \right\rfloor$
 end for
 $m_z \leftarrow$ median($\tilde{m}_0, \dots, \tilde{m}_{(N_{\text{rows}}/L)}$)
 end for
 Measure actual output rate
 Update R_{target} using Eq. (14) from [10]
 $Q_{y+1} \leftarrow Q_y$
 $R \leftarrow \sum_{z=1}^{N_{\text{bands}}} R(m_z, Q_{y+1})$
 if $R \geq R_{\text{target}}$ **then**
 while $R \geq R_{\text{target}}$ AND $Q_{y+1} < Q^{\max}$ **do**
 $R_{\text{old}} \leftarrow R$
 $Q_{y+1} \leftarrow Q_{y+1} + 2$
 $R \leftarrow \sum_{z=1}^{N_{\text{bands}}} R(m_z, Q_{y+1})$
 end while
 if $|R - R_{\text{target}}| > |R_{\text{old}} - R_{\text{target}}|$ **then**
 $Q_{y+1} \leftarrow Q_{y+1} - 2$
 end if
 else
 while $R \leq R_{\text{target}}$ AND $Q_{y+1} > 1$ **do**
 $R_{\text{old}} \leftarrow R$
 $Q_{y+1} \leftarrow Q_{y+1} - 2$
 $R \leftarrow \sum_{z=1}^{N_{\text{bands}}} R(m_z, Q_{y+1})$
 end while
 if $|R - R_{\text{target}}| > |R_{\text{old}} - R_{\text{target}}|$ **then**
 $Q_{y+1} \leftarrow Q_{y+1} + 2$
 end if
 end if
end for

it checks if such quantization step size yields a rate above or below the target, consequently deciding whether to increase it or decrease it. Since the quantization step sizes used in successive spectral lines are expected to be close, this strategy minimizes the expected number of times the memory storing the LUT is queried to compute $R(m, Q)$, thus reducing latency.

As a final remark, in a tradeoff between efficiency and available features, some characteristics of the algorithm in [10] have been traded for increased speed. In particular, the original algorithm performed simultaneous control of rate and quality on image blocks, thus enabling the definition of spatial as well as spectral regions of interest with different rate-distortion

TABLE I. TEST IMAGES

Image	Rows	Columns	Bands	P	Bit depth	Entropy (bpp)
AVIRIS SC0_RAW	512	680	224	15	16	12.62
AIRS GRAN9	135	90	1501	10	14	11.16
CASI-T0477F06-NUC	1225	406	72	2	16	10.65
CRISM-SC167-NUC	510	640	545	3	16	10.54
LANDSAT MOUNTAIN	1024	1024	6	5	8	6.33
MODIS-MOD01DAY	2030	1354	14	2	12	7.99

TABLE II. ACCURACY OF RATE CONTROL

Image	Method	0.5 bpp	1 bpp	2 bpp	3 bpp	4 bpp
AVIRIS	Proposed	0.501	1.002	2.002	3.001	4.001
SC0_RAW	[10]	0.510	1.005	2.000	2.990	3.998
AIRS	Proposed	0.502	1.006	2.001	2.993	4.021
GRAN9	[10]	0.487	0.987	1.967	2.942	3.930
CASI	Proposed	0.501	1.001	2.001	3.001	4.000
T0477F06-NUC	[10]	0.402	1.142	2.011	3.007	4.007
CRISM	Proposed	0.501	1.003	2.004	3.001	4.002
SC167-NUC	[10]	0.508	1.003	1.993	2.986	3.992
LANDSAT	Proposed	0.502	1.001	2.001	3.002	3.736 ^(*)
MOUNTAIN	[10]	0.435	0.898	2.008	2.838	3.736 ^(*)
MODIS	Proposed	0.508	1.002	2.005	3.006	4.005
MOD01DAY	[10]	0.478	1.018	2.016	3.008	4.006

(*) : lossless

policies. This feature is partially lost in the faster version as a single quantization step size is used for all the bands of a line.

The full algorithm, including both estimation of the statistics of residuals and computation of the quantization step sizes, is reported in Algorithm 1. Notice that the initialization of Q_1 may be different from $Q_1 = 1$ and, depending on the particular target rate, a higher Q might be appropriate.

IV. EXPERIMENTAL RESULTS

In this section we study the performance of the proposed rate control method. The algorithm is integrated into a lossy extension of CCSDS-123 where the prediction residuals are quantized with a uniform scalar quantizer and entropy coded with a range encoder [17] (refer to [10] for more details). Since the focus of this paper is on the rate control algorithm we omit comparisons with other compression algorithms (e.g. M-CALIC [18]) which do not provide rate control on a predictive encoder. We refer the reader to [19] for comparisons between M-CALIC and CCSDS-123 in near-lossless and rate-controlled (using [10]) mode, where it is shown that the latter typically outperforms the former. All tests are conducted with the feedback version of the rate control algorithm, with the rate adjusted using Eq. (14) from [10] (parameters $\tau = 5$, $|Z| = 1$). The rate LUT stores the value of $R(m, Q)$ multiplied by 1000 as an integer for the following domains of the input variables: $m \in [0, 1023]$, $Q = \{1, 3, 5, \dots, 511\}$. The choice of $Q^{max} = 511$ allows the algorithm to reach very low rates (if needed according to the statistics of the residuals and the target rate). Choosing a lower value of Q^{max} is possible but should not conflict with the target rate otherwise the rate controller will be unable to reach the desired rate. A lower value that is compatible with the target rate has the effect of a constraint on the maximum error and can often improve quality, as shown in [10] and in Table V, by forbidding excessive quantization in presence of inaccurate statistics of the residuals. The LUT uses 16-bit integers to represent the rate values, thus requiring 512 kB of storage in a software implementation. Notice that this

TABLE III. IMAGE QUALITY

Image	Rate	Proposed				[10]	
		SNR (dB)	MAD	SNR (dB)	MAD		
AVIRIS SC0_RAW	0.5 bpp	39.09	255	39.99	255		
	1 bpp	44.60	255	47.50	228		
	2 bpp	57.79	25	57.90	25		
	3 bpp	64.10	25	63.27	25		
	4 bpp	71.13	3	71.07	3		
AIRS GRAN9	0.5 bpp	38.12	255	46.41	53		
	1 bpp	53.86	18	54.93	10		
	2 bpp	63.04	4	63.31	4		
	3 bpp	67.28	3	69.64	1		
	4 bpp	79.25	1	76.78	1		
CASI-T0477F06-NUC	0.5 bpp	32.37	255	27.32	255		
	1 bpp	41.69	49	38.12	77		
	2 bpp	50.92	7	49.45	13		
	3 bpp	57.64	4	56.51	8		
	4 bpp	62.03	3	62.89	3		
CRISM-SC167-NUC	0.5 bpp	37.53	144	36.95	84		
	1 bpp	44.14	48	44.17	27		
	2 bpp	52.59	7	52.67	9		
	3 bpp	59.37	3	58.43	4		
	4 bpp	64.25	2	64.67	4		
LANDSAT MOUNTAIN	0.5 bpp	21.06	41	19.90	35		
	1 bpp	27.13	10	22.88	58		
	2 bpp	34.17	3	32.94	5		
	3 bpp	39.37	3	38.49	4		
	4 bpp	∞	0	∞	0		
MODIS-MOD01DAY	0.5 bpp	30.72	255	29.51	255		
	1 bpp	38.70	230	37.38	214		
	2 bpp	49.59	134	49.12	68		
	3 bpp	58.66	17	58.22	28		
	4 bpp	63.90	70	66.09	10		

requirement could be reduced by subsampling the available quantization step sizes (trading some accuracy in the rate) or choosing a Q^{max} lower than 511 (if the expected target rate is high). The median estimation uses subsets of length $L = 17$ pixels. This value was chosen to approximately equalize the complexity of the median computed at type-B pixels and the median computed at type-C pixels as explained in Sec. III. However, results in terms of image quality and rate are not very sensitive to the choice of this parameter.

All the following tests were performed on images extracted from the corpus¹ defined by the Multispectral and Hyperspectral Data Compression (MHDC) Working Group of the CCSDS for performance evaluation and testing of compression algorithms. We will use an ultraspectral image from the AIRS sensor, hyperspectral images from the AVIRIS, CRISM and CASI sensors and multispectral images from Landsat and MODIS sensors. Table I reports the test images and their dimensions, bit-depth and entropy, as well as the number of bands P used by the predictor.

The first test is concerned with the accuracy of the rate control algorithm in terms of output rate. Table II reports the output rate for various target rates. As can be seen, the proposed algorithm achieves remarkable accuracy for all target rates and for both the hyperspectral and multispectral cases.

Table III shows a comparison with the algorithm in [10] in terms of SNR defined as

$$\text{SNR} = 10 \log_{10} \frac{\sum_{i=1}^{N_{\text{pixel}}} s_i^2}{\sum_{i=1}^{N_{\text{pixel}}} (s_i - \tilde{s}_i)^2}$$

being s_i and \tilde{s}_i the i -th pixel in the original image and in the decoded image, respectively, and in terms of Maximum

¹ Available at <http://cwe.ccsds.org/sls/docs/sls-dc/123.0-B-Info/TestData>

TABLE IV. AVERAGE LUT LOOKUPS PER MEGAPIXEL AND RATE CONTROL RUNTIME (SEC.)

Image	Algorithm	0.5 bpp		1 bpp		2 bpp		3 bpp		4 bpp	
		Lookups	Time								
AVIRIS SC0_RAW	Proposed	13,485	0.0031	6,250	0.0026	3,412	0.0016	3,349	0.0019	3,553	0.0017
	[10]	615,567	11.8774	371,675	7.4340	257,333	5.4265	235,065	5.0735	240,722	5.9426
AIRS GRAN9	Proposed	79,506	0.0024	24,198	0.0008	23,704	0.0008	24,115	0.0008	22,058	0.0008
	[10]	369,916	1.5309	288,757	1.3689	273,748	1.5250	276,991	1.6256	242,905	1.0447
CASI-T0477F06-NUC	Proposed	8,149	0.0015	5,503	0.0015	5,125	0.0015	5,626	0.0014	5,001	0.0014
	[10]	767,673	6.7921	738,115	4.6862	576,056	3.7174	386,564	2.5567	362,650	2.4206
CRISM-SC167-NUC	Proposed	5,863	0.0022	3,713	0.0017	3,321	0.0015	3,398	0.0013	3,165	0.0015
	[10]	389,630	33.4665	282,061	20.1007	237,631	16.1764	234,546	16.8475	240,384	17.0099
LANDSAT MOUNTAIN	Proposed	2,418	0.0007	2,175	0.0009	2,244	0.0009	1,957	0.0009	1,951	0.0007
	[10]	382,011	0.4723	274,881	0.3482	246,634	0.3190	234,428	0.2991	209,264	0.2672
MODIS-MOD01DAY	Proposed	5,559	0.0024	5,155	0.0023	3,445	0.0020	1,729	0.0019	1,893	0.0019
	[10]	1,461,049	11.8816	1,179,269	8.2141	990,538	6.8846	1,337,818	9.3214	606,623	3.9860

TABLE V. IMPACT OF MAXIMUM ERROR CONSTRAINT

Image	Q^{max}	Rate (bpp)	SNR (dB)	MAD
AVIRIS SC0_RAW	511	1.002	44.60	255
	127	1.002	49.26	63
1 BPP TARGET	63	1.001	50.36	31
	11	2.656	63.25	5

Absolute Distortion (MAD), the largest deviation of a reconstructed pixel from the original one. It can be noticed that the image quality is comparable with the one produced by the original algorithm, which provided a significant improvement over some low-complexity transform coding approaches [2]. Remarkably, the proposed algorithm sometimes even outperforms the reference despite its lower complexity.

In order to test the increased efficiency of the algorithm we measured the total number of lookups to the rate LUT needed by the rate control algorithm and compared it with the algorithm in [10]. Table IV shows the number of lookups normalized by the image size in millions of pixels. Notice that the proposed algorithm improves over the old one by reducing the number of lookups by roughly two orders of magnitude. Notice that the number of lookups is directly related to the total number of operations performed by the rate controller, implying that a speedup of at least two orders of magnitude is expected. In order to confirm this we also measured the total execution time for the portion of code responsible for rate control, also reported in Table IV, under the Time column. The results confirm that the proposed algorithm is several orders of magnitude faster, needing milliseconds when the original required seconds to perform its operations.

V. CONCLUSION

We proposed a novel rate control algorithm for predictive coding of hyperspectral images onboard of spacecrafts and tested it by integrating it into a lossy extension of the CCSDS-123 recommendation. The proposed algorithm loses some flexibility with respect to [10] in terms of spatial modulation of quantization step sizes, but it is orders of magnitude faster enabling simpler and higher throughput implementations, while providing comparable or better image quality.

REFERENCES

- [1] Consultative Committee for Space Data Systems (CCSDS), "Image Data Compression," *Blue Book*, November 2005. [Online]. Available: <https://public.ccsds.org/Pubs/122x0b1c3.pdf>
- [2] I. Blanes and J. Serra-Sagrìstà, "Pairwise orthogonal transform for spectral image coding," *IEEE Trans. Geosci. Remote Sens.*, vol. 49, no. 3, pp. 961–972, 2011.
- [3] A. Abrardo, M. Barni, E. Magli, and F. Nencini, "Error-resilient and low-complexity onboard lossless compression of hyperspectral images by means of distributed source coding," *IEEE Trans. Geosci. Remote Sens.*, vol. 48, no. 4, pp. 1892–1904, 2010.
- [4] D. Valsesia and P. T. Boufounos, "Universal encoding of multispectral images," in *2016 IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, March 2016, pp. 4453–4457.
- [5] A. Barducci, D. Guzzi, C. Latri, V. Nardino, I. Pippi, and V. Raimondi, "Compressive sensing for hyperspectral Earth observation from space," *International Conference on Space Optics*, vol. 7, p. 10, 2014.
- [6] F. Rizzo, B. Carpentieri, G. Motta, and J. A. Storer, "Low-complexity lossless compression of hyperspectral imagery via linear prediction," *IEEE Signal Process. Lett.*, vol. 12, no. 2, pp. 138–141, 2005.
- [7] B. Aiazzi, P. Alba, L. Alparone, and S. Baronti, "Lossless compression of multi/hyper-spectral imagery based on a 3-d fuzzy prediction," *IEEE Trans. Geosci. Remote Sens.*, vol. 37, no. 5, pp. 2287–2294, 1999.
- [8] A. B. Kiely and M. A. Klimesh, "Exploiting calibration-induced artifacts in lossless compression of hyperspectral imagery," *IEEE Trans. Geosci. Remote Sens.*, vol. 47, no. 8, pp. 2672–2678, 2009.
- [9] Consultative Committee for Space Data Systems (CCSDS), "Lossless Multispectral and Hyperspectral Image Compression," *Blue Book*, no. 1, May 2012. [Online]. Available: <https://public.ccsds.org/Pubs/123x0b1ec1.pdf>
- [10] D. Valsesia and E. Magli, "A novel rate control algorithm for onboard predictive coding of multispectral and hyperspectral images," *IEEE Trans. Geosci. Remote Sens.*, vol. 52, no. 10, pp. 6341–6355, Oct 2014.
- [11] —, "A hardware-friendly architecture for onboard rate-controlled predictive coding of hyperspectral and multispectral images," in *2014 IEEE Int. Conf. on Image Processing*, Oct 2014, pp. 5142–5146.
- [12] A. Said and W. A. Pearlman, "A new, fast, and efficient image codec based on set partitioning in hierarchical trees," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 6, no. 3, pp. 243–250, Jun 1996.
- [13] B. Huang, *Satellite data compression*. Springer Science+ Business Media, 2011.
- [14] M. Conoscenti, R. Coppola, and E. Magli, "Constant-SNR, rate control and entropy coding for predictive lossy hyperspectral image compression," *IEEE Trans. Geosci. Remote Sens.*, (accepted for publication).
- [15] P. J. Huber, *Robust statistics*. Springer, 2011.
- [16] M. Blum, R. W. Floyd, V. Pratt, R. L. Rivest, and R. E. Tarjan, "Time bounds for selection," *Journal of Computer and System Sciences*, vol. 7, no. 4, pp. 448 – 461, 1973.
- [17] G. N. N. Martin, "Range encoding: an algorithm for removing redundancy from a digitised message," in *Video and Data Recording Conference*, 1979, pp. 24–27.
- [18] E. Magli, G. Olmo, and E. Quacchio, "Optimized onboard lossless and near-lossless compression of hyperspectral data using CALIC," *IEEE Geosci. Remote Sens. Lett.*, vol. 1, no. 1, pp. 21–25, Jan 2004.
- [19] I. Blanes, E. Magli, and J. Serra-Sagrìstà, "A Tutorial on Image Compression for Optical Space Imaging Systems," *IEEE Geosci. Remote Sens. Mag.*, vol. 2, no. 3, pp. 8–26, Sept 2014.