Isorange Pairwise Orthogonal Transform

Ian Blanes, *Member, IEEE,* Miguel Hernández-Cabronero, Francesc Aulí-Llinàs, *Senior Member, IEEE,*
Joan Serra-Sagristà, *Senior Member, IEEE,* and Michael W. Marcellin, *Fellow, IEEE*

Abstract—**Spectral transforms are tools commonly employed in multi- and hyperspectral data compression to decorrelate images in the spectral domain. The Pairwise Orthogonal Transform (POT) is one such transform that has been specifically devised for resource-constrained contexts like those found on board satellites or airborne sensors. Combining the POT with a 2D coder yields an efficient compressor for multi- and hyperspectral data. However, a drawback of the original POT is that its dynamic range expansion —i.e., the increase in bit depth of transformed images— is not constant, which may cause problems with hardware implementations. Additionally, the dynamic range expansion is often too large to be compatible with the current 2D standard CCSDS 122.0-B-1. This paper introduces the Isorange Pairwise Orthogonal Transform, a derived transform that has a small and limited dynamic range expansion, compatible with CCSDS 122.0-B-1 in almost all scenarios. Experimental results suggest that the proposed transform achieves lossy coding performance close to that of the original transform. For lossless coding, the original POT and the proposed isorange POT achieve virtually the same performance.**

Index Terms—**On-board multi- and hyperspectral image coding, progressive lossy-to-lossless, limited dynamic range expansion, Pairwise Orthogonal Transform (POT).**

I. INTRODUCTION

IMAGERS are remote-sensing devices that often produce abundant amounts of data, and more so with the constant advent of new generations of sensors, each with a higher resolution than the previous one. For imagers carried on satellites (i.e., space-borne), the acquisition of large amounts of data is often juxtaposed with the necessity to transmit these data to ground stations. Transmission is made difficult by the bandwidth constraints of the downlink channel, and this makes on-board data compression a fundamental part of satellite signal processing hardware [1].

Image compression is particularly important for sensors that acquire multi- and hyperspectral images, i.e., those that capture light intensities at multiple wavelength ranges for a given spatial location. The memory and bandwidth resources required to store and transmit such acquired data can be significantly reduced by means of compression techniques [2]. Depending on the needs of a given mission, the acquired data can be either compressed without any fidelity loss, or

by trading fidelity losses for smaller file sizes. The former and latter modes are respectively called lossless and lossy coding.

On-board lossless compression of multi- and hyperspectral images has been thoroughly studied in the literature. For instance, [3] and [4] extend the Context-based Adaptive Lossless Image Coding (CALIC) scheme [5] to 3D-CALIC and M-CALIC, respectively. A predictive coder based on look-up tables is introduced in [6], and variations of this technique are presented in [7] and [8]. Another relevant method is the crisp and fuzzy adaptive predictor called S-FMP [9]. In 2012, the Consultative Committee for Space Data Systems (CCSDS) [10] —a consortium of major space agencies— published the recommended standard CCSDS 123.0-B-1 [11] for lossless multi- and hyperspectral image compression, which is based on the Fast Lossless method introduced in [12], and further investigated in [13].

CCSDS 123.0-B-1 provides excellent lossless coding performance, but does not support lossy coding. Transform coding is a frequently used approach to provide for both lossy and lossless coding within the same coding algorithm. Such algorithms typically suffer a small performance penalty in lossless coding performance compared to algorithms that focus exclusively on lossless coding. A common approach that achieves good performance in lossy-to-lossless coding is to utilize a one-dimensional transform to exploit the correlation between spectral bands, followed by encoding each resulting band with a lossy-to-lossless 2D image coder [2], [14], [15], [16], [17], [18].

The one-dimensional transforms that typically achieve the best performance in such a coding scheme for multi- and hyperspectral imagery are the Karhunen Loeve Transform (KLT) and other similarly data-dependent transforms [19]. As originally formulated, KLT-based transforms are usually not suitable for on-board devices due to their high computational complexity, both because they involve a high number of arithmetic operations, and because they are data dependent, which in turn, necessitates large memory buffers for training.

The reduction of computational complexity for the KLT has been addressed in the literature in different forms [20]. The Pairwise Orthogonal Transform (POT) [21] is a spectral transform devised for on-board image compression whose main insight is a divide-and-conquer strategy that approximates the KLT, while requiring only a fraction of the

computational resources. Other divide-and-conquer strategies, which do not address on-board image coding, are the recursive approach [22], the two level approaches [23], and the multi-level structures [24].

While the CCSDS has yet to create a standard for lossy multi- and hyperspectral image compression, the natural approach to an on-board lossy-to-lossless coder would use a one-dimensional transform, such as the POT, together with CCSDS 122.0-B-1 [25]. CCSDS 122.0-B-1 is a lossy-to-lossless 2D image coder (not to be confused with the similarly named lossless coder CCSDS 123.0-B-1, discussed previously).

The use of CCSDS 122.0-B-1 for the 2D image coder is both natural and convenient as it is a proven space-qualified technology and several hardware implementations exist [26], [27], [28]. The POT, on the other hand, is suitable for on-board devices due to its reduced computational complexity. The major issue that may arise in such an approach is related to the bit depth expansion that occurs naturally in the POT (and most others transforms). Let us explain further. The CCSDS 122.0-B-1 standard supports an image bit depth of up to 16 bits. This maximum was chosen since, in general, imaging sensors do not produce data with bit depths in excess of this value. However, when a 1D spectral transform is employed as a preprocessing stage to high bit-depth multi- or hyper-spectral images, the resulting data may exceed this maximum. Thus, even though bit depth expansion may not be problematic in most scenarios, in the context of leveraging existing investments in CCSDS 122.0-B-1 hardware implementations, bit depth expansion becomes a critical issue.

When applied to multi- or hyper-spectral imagery, the 1D KLT (and transforms derived from the 1D KLT) tend to accumulate most of the image energy in just a few transformed bands that then require a large dynamic range. If not addressed, this may restrain the use of KLT-based transforms in combination with CCSDS 122.0-B-1.

This paper introduces a version of the POT that reduces and strictly bounds dynamic range expansion while preserving coding performance for both lossy and lossless regimes. Since our goal is to preserve dynamic range as much as possible, we refer to such transform as the *isorange* POT, using a composition of the prefix *iso-* and the word *range*. The proposed isorange POT maintains the advanced features of the original POT such as its ability to operate in progressive lossy-to-lossless mode, i.e., to smoothly progress from lossy to lossless bitstreams as the bit budget for the compressed image increases. To the best of our knowledge, this is the first KLT-based isorange transform.

The paper is organized as follows. Section II overviews the POT and introduces necessary notation. Section III presents the proposed isorange transform based on the POT. Section IV provides a mathematical derivation that bounds its dynamic range expansion. Section V provides experimental results, and the last section draws some conclusions.

II. THE ORIGINAL TRANSFORM

The POT is a divide-and-conquer approximation to the KLT in which the spectral bands (or components) of an image are transformed using multiple low-cost *pairwise operations* instead of a single and computationally-expensive transform.

Each pairwise operation applied in the POT is a KLT applied on only two components of the image, and these operations are organized in a multi-level structure. In each level, components are decorrelated in pairs and only the resulting component (of the pair) with the most energy is further decorrelated in a next level. This procedure is repeated until only one pairwise operation is applied on a final level. The multi-level decomposition of the POT is illustrated in Fig. 1a.

When the number of components to decorrelate in a level is odd, there is an *unpaired component*. That component is selected to be the last or the first in successive levels, so that any unpaired component traverses at most one level without being decorrelated (see Fig. 1b).
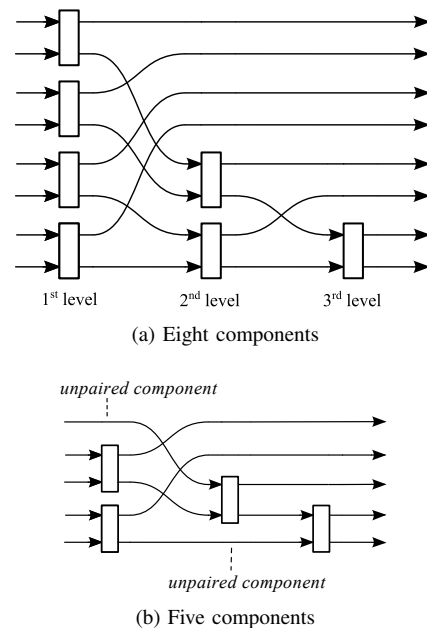

1st level   2nd level   3rd level

(a) Eight components

unpaired component

unpaired component

(b) Five components

Figure 1. Multi-level structure of the POT.

The computational cost of the KLT is in practice $O(n^2)$, with n being the number of components. The complexity of the KLT is dominated by an –usually– $O(n^3)$ eigendecomposition [29]; however, in practice many more than n vectors are transformed for each eigendecomposition. Thus,

the n^3 contribution of each eigendecomposition to the total cost of the transform is negligible (see [16], [24]). On the other hand, by working on only two components at once for a limited number of times, the POT results in a complexity of only $O(n)$, as each pairwise operation has a constant cost and a POT has $n-1$ pairwise operations.[1]

Other advantages of the POT are that it improves component scalability (yielding better error resiliency), that it has reduced memory requirements, and that it is easier to implement than the KLT. A single POT applied over the spatial extent of the entire image is described in the following discussion, although applying different POTs to different regions of a segmented image is straightforward.

Like the KLT, the POT assumes that the intensity distribution of each component is centered around zero. Since the data of a component do not commonly exhibit this distribution, input components should be *mean-corrected* by subtracting from their values an integer approximation of their original mean. This operation increases the bit depth of the input image by one bit at most, always producing a signed image with approximately zero arithmetic mean.
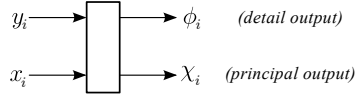

Figure 2.   A single pairwise operation.

Within each level of the structure, each pairwise operation takes two input components, denoted by x_i and y_i, and produces two output components: the principal output component χ_i —containing most of the input energy— and the detail output component ϕ_i —containing the remaining energy— (see Fig. 2). Ideally, a pairwise operation would use a KLT matrix to exactly produce χ_i^+ and ϕ_i^+ from x_i and y_i by

$$\begin{pmatrix} \chi_i^+ \\ \phi_i^+ \end{pmatrix} = Q \begin{pmatrix} x_i \\ y_i \end{pmatrix}, \qquad (1)$$

where Q is the transposed KLT matrix, which can be represented as a generic rotation matrix

$$Q = \begin{pmatrix} \sqrt{1-t^2} & t \\ -t & \sqrt{1-t^2} \end{pmatrix}, \qquad (2)$$

where $t \in [-1, 1]$ is a scalar parameter derived from the covariance matrix of the inputs x_i and y_i. That is,

$$t = f\left(\sigma_{x_i}^2, \sigma_{y_i}^2, \sigma_{x_i, y_i}\right). \qquad (3)$$

The definition of the function f used to compute t is given in [21]. The value of t needs to be stored as side information

---

[1]The latter can be proven by induction over the total number of pairwise operations, say $g(n)$, as a combination of the number of pairwise operations in the first level and the number of pairwise operations in the remaining ones, i.e., $g(n) = \lfloor n/2 \rfloor + g(\lceil n/2 \rceil) = n - 1$, with $g(1) = 0$.
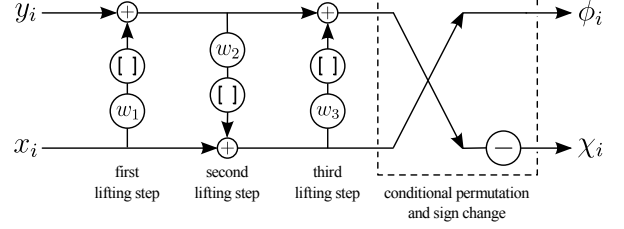

Figure 3.   Lifting structure for a pairwise operation.

to allow the inverse transform procedure to be performed in the decoder.

As described in [30], Q can be decomposed into three lifting steps (see [31], [32]) with an optional permutation and sign change according to

$$Q = \begin{pmatrix} 1 & 0 \\ 0 & s \end{pmatrix} P \begin{pmatrix} 1 & 0 \\ w_3 & 1 \end{pmatrix} \begin{pmatrix} 1 & w_2 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ w_1 & 1 \end{pmatrix}, \qquad (4)$$

which admits two solutions:

$$\begin{cases} P = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \ w_1 = w_3 = \frac{p-1}{t}, \ w_2 = t, \ s = 1 \\ P = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \ w_1 = w_3 = \frac{1-t}{p}, \ w_2 = -p, \ s = -1 \end{cases} \qquad (5)$$

where $p = \sqrt{1 - t^2}$.

The lifting network associated with these solutions is depicted in Fig. 3, including the permutation P and sign change s of the second solution. In this work, finite precision arithmetic and rounding are used within the lifting network to obtain the integer outputs χ_i and ϕ_i which approximate χ_i^+ and ϕ_i^+. Rounding is indicated by $[\ ]$ in the figure. It is worth noting that if real number arithmetic and no rounding were used, the outputs of the lifting network would be exactly χ_i^+ and ϕ_i^+. In subsequent discussions, when we refer loosely to "errors introduced by the lifting network," we mean "errors introduced by the finite precision arithmetic and rounding used in the lifting network."

The lifting implementation of the transform provides several advantages over a straightforward matrix multiply implementation. One such advantage is a reduced implementation complexity. Specifically, the lifting implementation of the POT uses one fewer multiplication per sample. A more significant advantage is that the finite precision arithmetic and rounding described above do not destroy the invertibility of the transform. Invertible transforms that map integers to integers are often called "reversible transforms" [33]. Such transforms are useful for lossless (as well as lossy) coding.

Ignoring errors caused by the lifting network (finite precision arithmetic and rounding), the dynamic range expansion in each output of each pairwise operation is 0.5 bits, since the largest output value can be as large as

$$\max_{-1 \le t \le 1} \|Q\|_\infty = \sqrt{2} \qquad (6)$$

times larger than the largest input value [21]. Thus, the dynamic range expansion for the entire POT is 0.5 bits times the number of levels, plus one extra bit from the mean-correction stage. This results in an expansion of $0.5 \cdot \lceil \log_2(n) \rceil + 1$ bits, where n is the number of input components. For example, the dynamic range expansion is 5 bits for an image acquired with the EO-1 Hyperion sensor ($n = 242$), and 8 bits for data acquired with one of the MetOp IASI sounders ($n = 8461$). As seen below, this expansion is further increased to accommodate errors introduced by the lifting network.

III. THE ISORANGE POT

In the following, some modifications to the POT formulation are described. The aim of these modifications is to bound the dynamic range expansion of the transform to two bits (three bits if mean correction is required).

The idea behind the isorange POT stems from the fact that the dynamic range expansion of the original POT is mostly caused by the accumulation of the expansion of the principal output (χ_i) over multiple levels. If such accumulation can be redressed, the expansion of the entire transform becomes just the expansion of ϕ_i over a single pairwise operation.

A straightforward solution might appear to be achievable by dividing χ_i by $\sqrt{2}$ after each pairwise operation. Unfortunately, this would preclude an integer transform. If the division were applied directly to χ_i at the output of the lifting network, followed by rounding, it would not be possible to recover χ_i from the rounded value $\left\lceil \chi_i / \sqrt{2} \right\rceil$ for use at the input of the inverse transform. Attempting to incorporate the division into the original transform as

$$Q' = \begin{pmatrix} 1/\sqrt{2} & 0 \\ 0 & 1 \end{pmatrix} Q \qquad (7)$$

is also fruitless, due to the fact that the determinant of Q' is $1/\sqrt{2}$, while lifting decompositions exist only for transforms having determinant of ± 1 [30, p. 2317].

This problem can be ameliorated by employing a transform of the form

$$Q'' = S_b Q, \qquad (8)$$

where S_b is a *shift matrix* of the form

$$S_b = \begin{pmatrix} 2^{-b} & 0 \\ 0 & 2^{+b} \end{pmatrix}. \qquad (9)$$

The shift matrix results in the multiplication of χ_i^+ by 2^{-b} and ϕ_i^+ by 2^b, which decreases the dynamic range of χ_i^+ by b bits, but increases that of ϕ_i^+ by b bits. Since $\text{Det}(S_b) = 1$, a lifting decomposition can still be obtained for Q''.

Employing the shift matrix with $b = 0.5$ exactly cancels the 0.5 bit expansion incurred by χ_i^+ of the original POT,

entirely eliminating the accumulated expansion associated with multiple levels. On the other hand, the expansion incurred by ϕ_i^+ is increased to a total of 1.0 bits. The matrix of the resulting pairwise operation is defined as

$$Q_R = S_{0.5} \cdot Q, \qquad (10)$$

which has the following two solutions for the lifting network:

$$
\begin{cases}
P = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \; w_1 = \frac{p - \sqrt{2}}{t}, \; w_2 = \frac{t}{\sqrt{2}}, \\
\qquad\qquad w_3 = \frac{2p - \sqrt{2}}{t}, \; s = 1 \\
P = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \; w_1 = \frac{\sqrt{2} - 2t}{2p}, \; w_2 = -\sqrt{2}\,p, \\
\qquad\qquad w_3 = \frac{\sqrt{2} - t}{2p}, \; s = -1
\end{cases} \qquad (11)
$$

where, as before, $p = \sqrt{1 - t^2}$.

Employing the lifting steps of (11) results in a multilevel transform with dynamic range expansion determined only by the expansion of the detail output of each pairwise operation, which is 1 bit. As discussed subsequently in section IV, additional expansion is caused by the errors introduced in the lifting networks.

The transformed image components resulting from the application of this modified POT are quasi-linearly scaled versions of those from the original POT. The scaling factor, or *gain*, introduced by the shift matrices must be compensated in any distortion estimates that might be employed in rate allocation for lossy coding. Small non-linearities are introduced due to rounding errors in the lifting steps, but otherwise the scaling is strictly linear.

The gain of a given transformed component can be obtained by multiplying the individual gains incurred in each pairwise operation involved to produce it. See Fig. 4 for an example for the case of 8 input components.
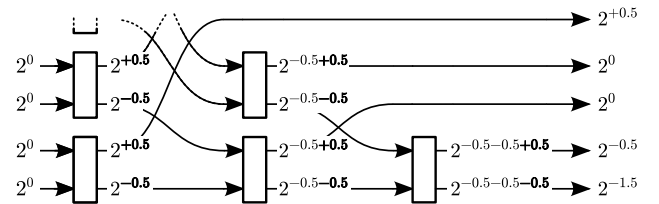

Figure 4. Accumulated transform gain for a three level transform of 8 input components. For the sake of space, two pairwise operations are not shown at the first level. The two omitted pairwise operations are identical to the two shown. Similarly, three of the detail outputs from the first level are not shown. Each of these has a gain of $2^{+0.5}$, identical to the one shown.

The unbalanced case

When the number of input components is not a power of 2, a pairwise operation may involve an unpaired component (see

Fig. 1b). This, in turn, may result in two inputs of different energy gains: a regular component and the unpaired one, which has not been processed in the immediate previous level. These components have a gain difference of $2^{0.5}$ that, if not handled properly, may significantly bias any distortion estimates used for rate allocation.

An *unbalanced* pairwise operation is now introduced to address this issue. Assume that y_i is the component not processed in the previous level. Otherwise, x_i and y_i can be swapped. The transform matrix for the unbalanced operation is then

$$Q_{\mathrm{U}} = S_{0.25} \cdot S_{0.5} \cdot Q \cdot S_{-0.25} = S_{0.25} \cdot Q_{\mathrm{R}} \cdot S_{-0.25}. \quad (12)$$

The $S_{-0.25}$ shift matrix compensates the $2^{0.5}$ gain difference between the inputs by increasing the gain of x_i by $2^{0.25}$ and decreasing the gain of y_i by $2^{0.25}$ (see Fig. 5). This operation equalizes the gain of both inputs at a gain $2^{0.25}$ times higher than that of a paired input, which if left uncorrected, would be transfered to the principal output and cause similar unbalancing problems in subsequent levels. This aspect is addressed by means of an additional shift matrix that displaces a gain of $2^{0.25}$ from χ_i to ϕ_i, so that the gain of χ_i is the same as that of a normal pairwise operation. The gain for ϕ_i is then 2^1, which results from the $2^{0.25}$ gain increase used to equalize inputs, multiplied by the original gain of $2^{0.5}$ from Q_{R}, multiplied by the $2^{0.25}$ gain to decrease the gain of χ_i.
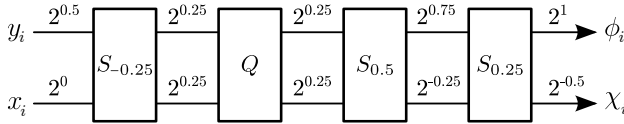

Figure 5. Accumulated transform gain for Q_{U}, broken down by shift matrix. The initial gain of y_i is $2^{0.5}$ because it is not processed in the previous level. Notice that $S_{0.25}$ corrects the gain of χ_i from $2^{-0.25}$ to $2^{-0.5}$.

The transform matrix of (12) can be decomposed into a lifting network with the following two solutions:

$$\begin{cases} P = \left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix} \right), \ w_1 = \frac{\sqrt{2}\,p-2}{t}, \ w_2 = \frac{t}{2}, \\ \qquad w_3 = \frac{2\sqrt{2}\,p-2}{t}, \ s = 1 \\ \\ P = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix} \right), \ w_1 = \frac{\sqrt{2}(1-2t)}{2p}, \ w_2 = -\sqrt{2}\,p, \\ \qquad w_3 = \frac{\sqrt{2}(2-t)}{4p}, \ s = -1 \end{cases} \quad (13)$$

In this case, the parameter t is computed as

$$t = f\left(\sigma_{x_i}^2, \left(2^{-0.5} \right)^2 \cdot \sigma_{y_i}^2, 2^{-0.5} \cdot \sigma_{x_i, y_i} \right). \quad (14)$$

which takes into account the fact that the gain of the original y_i is greater than that of x_i by a factor of $2^{0.5}$ (cf. Eq. 3).

Without taking into account the expansion caused by the errors introduced in the lifting networks, the aforementioned modifications to the POT result in a total dynamic expansion of the transform of 1.5 bits. This expansion is determined by the detail output of the unbalanced operations, in which 0.5 bits correspond to the original expansion and 1 bit to gain modifications. Errors introduced in the lifting steps may result in additional dynamic range expansion. In the next section, it is demonstrated that the total expansion of a multilevel POT (including errors from the lifting steps) can be limited to no more than two bits.

IV. ANALYSIS OF ROUNDING ERRORS

As mentioned earlier, the finite arithmetic and rounding steps introduced in the lifting network yield an integer-valued transform that approximates the real-valued version of the POT. The associated errors alter the dynamic range of the transform in a non-trivial manner. In this section, we prove the existence of a 2-bit bound on the total dynamic range expansion and provide necessary conditions to guarantee this bound.

The analysis first finds a bound for errors introduced in the lifting network associated with each pairwise operation. Using this bound, an overall 2-bit bound for the whole transform is then found.

There are two factors that introduce errors in the result of a pairwise operation: rounding and finite precision arithmetic.

Assuming rounding to the nearest integer, such operations introduce an absolute error of at most 0.5. This error is further modulated by subsequent lifting steps within the same pairwise operation.

Regarding finite precision arithmetic, in practice it is achieved via the use of finite precision weights within the lifting network. Suppose that a weight $\tilde{w}_j$ has Θ bits of fractional precision. Again, assuming nearest-integer rounding, the error in the weight $\tilde{w}_j$ is $|\tilde{w}_j - w_j| \leq 2^{-\Theta-1}$. Suppose further that the inputs are signed integers with bit depth B. Then, $|x_i|$ and $|y_i|$ are bounded by 2^{B-1}.

Formally, the expression within the norm of (15) represents the error in each of the two outputs χ_i and ϕ_i for a given pair of inputs x_i and y_i. The infinity-norm selects the largest of these two errors, which represent the difference between the computed outputs —with finite precision arithmetic and rounding— as compared to the ideal outputs —without finite precision arithmetic or rounding. The permutation and sign change in (11) are not considered here because they do not affect the dynamic range expansion. From this expression, a bound $\mathbb{E}(\cdot)$ on the absolute error introduced by the lifting network can be derived as given in (16). See the Appendix for a derivation of this bound.

$$\left\| \begin{pmatrix} 1 & 0 \\ w_3 & 1 \end{pmatrix} \begin{pmatrix} 1 & w_2 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ w_1 & 1 \end{pmatrix} \begin{pmatrix} x_i \\ y_i \end{pmatrix} - \begin{pmatrix} x_i + [\tilde{w}_2 (y_i + [\tilde{w}_1 x_i])] \\ y_i + [\tilde{w}_1 x_i] + [\tilde{w}_3 (x_i + [\tilde{w}_2 (y_i + [\tilde{w}_1 x_i])])] \end{pmatrix} \right\|_\infty \tag{15}$$

$$\mathbb{E}(B, \Theta, w_1, w_2, w_3) =$$
$$= \max \Big\{ 2^{B-\Theta-2} (|w_1| + |w_2| + 1) + 2^{-1}|w_2| + 2^{B-2\Theta-3} + 2^{-\Theta-2} + 2^{-1},$$
$$\left(2^{-3\Theta-2}\right)\left(2^{B+\Theta}(|w_1| + |w_2| + 1) + 2^{2\Theta+1}(|w_2|+1) + 2^\Theta + 2^{B-1}\right)\left(2^\Theta|w_3| + 2^{-1}\right)$$
$$+ 2^{B-\Theta-2}(|w_2 w_1| + |w_2|) + 2^{B-\Theta-1} + 1 \Big\} \tag{16}$$

While this bound is applicable to any choices of the weights w_1, w_2, and w_3, the weights corresponding to the balanced and unbalanced pairwise operations presented above are of particular interest. Recall that for a particular value of t, all pairwise operations presented above have two solutions to the decomposition of their transform matrices into a lifting network. Let solution $\boldsymbol{\alpha}$ be the solution in which $P = \left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right)$ in (11), and let solution $\boldsymbol{\beta}$ be the solution in which $P = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$.

Figures 6a and 6b show the weights for the balanced pairwise operation Q_R, as a function of t, for solutions $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$, respectively. As seen there, solution $\boldsymbol{\alpha}$ presents non-removable discontinuities in w_1 and w_3 at $t = 0$. On the other hand, solution $\boldsymbol{\beta}$ presents singularities in w_1 and w_3 at $t = \pm 1$. Figures corresponding to the unbalanced pairwise operation Q_U are similar and are not presented here.
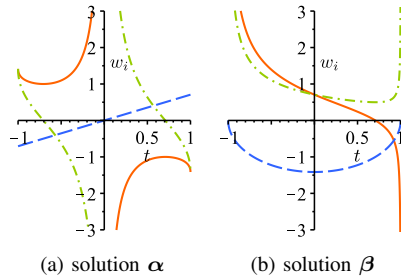

(a) solution $\boldsymbol{\alpha}$ (b) solution $\boldsymbol{\beta}$

Figure 6. Lifting solutions for Q_R. Values of —— w_1, -- w_2, and -·- w_3 are plotted as functions of t.
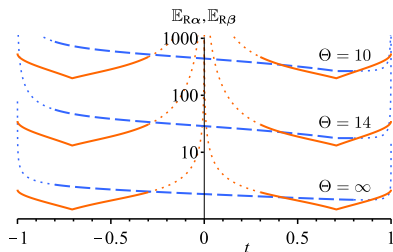
Figure 7. Bounds $\mathbb{E}_{R\boldsymbol{\alpha}}$ (——) and $\mathbb{E}_{R\boldsymbol{\beta}}$ (-- -) with $B = 17$ and $\Theta = \{10, 14, \infty\}$.

For a given value of t, the two resulting sets of weights for the balanced case can be substituted into the bound of (16), to obtain two bounds, each a function of B, Θ, and t, say $\mathbb{E}_{R\boldsymbol{\alpha}}(B, \Theta, t)$ and $\mathbb{E}_{R\boldsymbol{\beta}}(B, \Theta, t)$ (see Fig. 7).

Given the discontinuities discussed with respect to Fig. 7, it seems reasonable to consider solution $\boldsymbol{\alpha}$ only when $|t| \geq 0.3$ and solution $\boldsymbol{\beta}$ only when $|t| \leq 0.8$.[2]

With this restriction, a bound over all t for the balanced case can be found as

$$\mathbb{E}_R(B, \Theta) = \max \left\{ \max_{|t| \geq 0.3} \mathbb{E}_{R\boldsymbol{\alpha}}(B, \Theta, t), \right.$$
$$\left. \max_{|t| \leq 0.8} \mathbb{E}_{R\boldsymbol{\beta}}(B, \Theta, t) \right\}. \tag{17}$$

The bound for the unbalanced case $\mathbb{E}_U(B, \Theta)$ is found in the same way.

For any given Θ and B, exact values for $\mathbb{E}_R(B, \Theta)$ and $\mathbb{E}_U(B, \Theta)$ can be calculated using a computer algebra system. These bounds are dominated by a factor of $2^{-(\Theta-B)}$. Indeed, Fig. 8 shows $\mathbb{E}_R(B, \Theta)$ as a function of $\Theta - B$, for many values of Θ and B. As can be seen from this figure, different values of the bound are very similar for equal differences $\Theta - B$, indicating that the bound is strongly dependent on $\Theta - B$, and only loosely dependent on Θ and B, individually.

For a finite bit depth B, perfect representation of the weights (i.e., $\Theta = \infty$) yields the lowest value of the bounds and allows us to distinguish the error introduced by the finite weight representation from that introduced by the rounding operations within the lifting network. Thus, bounds on the error resulting from rounding operations in the balanced and unbalanced cases are $\mathbb{E}_R(B < \infty, \Theta = \infty) = (110 + 37\sqrt{2})/60 \simeq 2.71$ and $\mathbb{E}_U(B < \infty, \Theta = \infty) = 3.4$. It can also be observed, that for $\Theta - B \geq 4$, the error introduced by the use of finite precision weights is essentially negligible.

With the error for a single pairwise operation now bounded, the error for the whole transform can be addressed. This er-

---

[2]In practice, when data are being filtered, a single threshold T, $0.3 \leq T \leq 0.8$, may be used to select solution $\boldsymbol{\alpha}$ when $|t| > T$ or solution $\boldsymbol{\beta}$ when $|t| \leq T$.
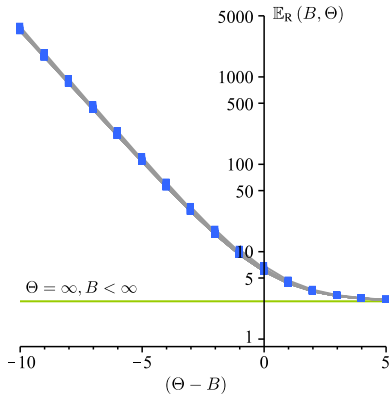
Figure 8. Bound $\mathbb{E}_R(B, \Theta)$ as a function of $\Theta - B$ for values of B from 1 to 18 and for values of Θ from 1 to 21. Similar results for the unbalanced case are omitted for brevity.

ror depends on the number of transform levels employed. In this work, 16 levels are considered sufficient, as this allows the POT to be performed on up to 2^{16} input components. For a 16-level structure, the worst case in dynamic range expansion happens in the detail output of an unbalanced operation at the 16^{th} level. For this case, the maximum absolute value of an output sample, say peak_{out}, is the maximum absolute input value, say peak_{in}, plus the lifting error accumulated at each of the previous 15 levels, increased by the implicit dynamic range expansion of the detail output of the last unbalanced pairwise operation and by its lifting error. Considering an equal maximum error E_p for each pairwise operation, this results in

$$\text{peak}_{\text{out}} \leq \left(\text{peak}_{\text{in}} + 15\, E_p\right) 2^{1.5} + E_p. \qquad (18)$$

Assuming two's complement signed inputs and outputs, $\text{peak}_{\text{in}} = 2^{B-1}$. Assuming further a maximum desired dynamic range expansion of 2 bits, we must then have $\text{peak}_{\text{out}} = 2^{B+1} - 1$. Note that we cannot allow $\text{peak}_{\text{out}} = 2^{B+1}$ as this would permit an output of $+2^{B+1}$ which would not fit in $B + 2$ bits. Thus, the assumption of two's complement representation is more restrictive than the assumption of sign magnitude representation. Substituting these values into (18) and solving for E_p yields

$$E_p \leq \frac{2^{B+1} - 2^{B+1/2} - 1}{30\sqrt{2} + 1}. \qquad (19)$$

The right side of this inequality, referred to hereafter as E, can be interpreted as the maximum allowable error (per level) that can be tolerated without exceeding a total of 2 bits of dynamic range expansion. For a given bit depth B, it then suffices to choose Θ so that

$$\mathbb{E}_R(B+1, \Theta) \leq E \text{ and } \mathbb{E}_U(B+1, \Theta) \leq E. \qquad (20)$$

Notice that the input bit depth in the pairwise operations at the second and higher levels is larger than B because of the

expansion in previous levels. Of the 2 bits of total expansion, 1.5 bits are intrinsic in the transform (in the unbalanced pairwise operation of the last level), and less than 0.5 bits of expansion can be produced before the inputs of the last level. Hence, it is sufficient to consider $B + 1$ instead of B in (20).

Using (19), (20) and the bounds provided by (17), the minimum values of Θ that allow an isorange transform with 2 bits dynamic range expansion are shown in Table I. It can be observed that for images with bit depths of at least 8 bits, it is always possible to find a value of Θ that satisfies (20). Bit depths below 8 are not of particular interest for this study. However, we discuss them here for the sake of completeness. For bit depths of 7 and below, the necessary value of E to contain the dynamic range expansion is below the threshold set by $\mathbb{E}_R(B < \infty, \Theta = \infty)$ and $\mathbb{E}_U(B < \infty, \Theta = \infty)$, and is thus unachievable. In practice, input data of low bit depth can be presented to the transform network using the least-significant part of an 8-bit word. Outputs of the transform can then be guaranteed to fit in no more than $8 + 2 = 10$ bits.

Table I
MAXIMUM ERROR PER LEVEL AND MINIMUM Θ REQUIRED AS A
FUNCTION OF B.

B	E	Min. Θ
7	1.70	-
8	3.43	12
9	6.88	10
10	13.78	10
11	27.60	10
12	55.22	10
13	110.48	10
14	220.98	9
15	441.99	9
16	884.00	9
17	1768.03	9

To recapitulate, the dynamic range expansion of the isorange POT has been proven to be 2 bits for all images of up to 2^{16} components, with bit depths of at least 8 bits, and where the intensity distributions of each component is centered around zero. Or in general, the dynamic range expansion has been proven to be 3 bits for all images of up to 2^{16} components, with bit depths of at least 8 bits, regardless of the intensity distribution of each component.

In case of being necessary, (18) and (19) can be trivially adapted for ultraspectral images of more than 2^{16} components, at the expense of larger values of minimum Θ. For a number of levels L,

$$E_p \leq \frac{2^{B+1} - 2^{B+1/2} - 1}{(L-1)2\sqrt{2} + 1}. \qquad (21)$$

Thus, for example, a 10-bit image of 2^{20} components requires $10.93 < \boldsymbol{E}$. Evaluating $\mathbb{E}_R$ and $\mathbb{E}_U$ for $\Theta = 10$ yields $\mathbb{E}_R(11, 10) < 9.23$ and $\mathbb{E}_U(11, 10) < 9.31$, which guarantees the bound for the given example.

Two practical considerations conclude this analysis. The first is that the bounds provided in this section are rather conservative. In Fig. 9, for a given value of t, a sample of 4096 input pairs, each a 17 bit signed integer, are processed by an isorange pairwise operation corresponding to that value of t. The weights employed in each pairwise operation have a fractional bit depth of 16. The resulting absolute peak error is then plotted as a function of t. For comparison, the corresponding bounds are also included in the figure. As can be seen in the figure, the actual peak error is typically significantly less than that implied by the bounds.
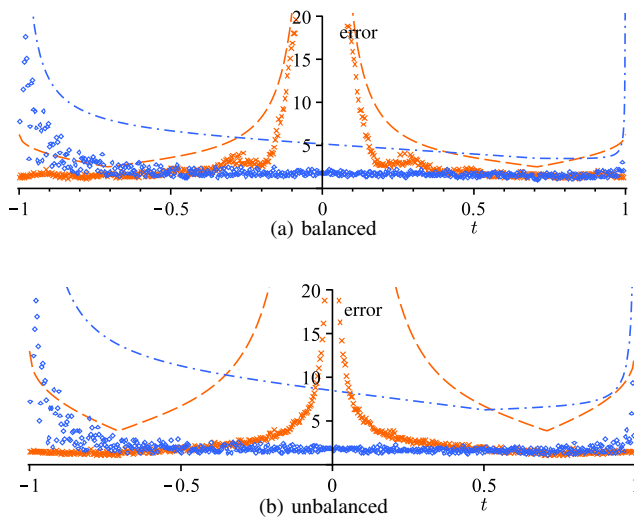

Figure 9. Peak absolute error of a sample of 4096 points with $\Theta = 16$ for solutions $\boldsymbol{\alpha}$ (✕) and $\boldsymbol{\beta}$ (◇), and their respective bounds (– – and – · –). Input pairs are uniformly distributed with $B = 17$, and t is sampled at steps of 2^{-8} (except for 0 and ± 1 that are not sampled).

The second practical consideration is in relation to the computational cost of the transform. In this regard, since the transform application still uses the same lifting network, its application costs remain the same as those of the original POT. However, the computational cost required to obtain the values of $\tilde{w}_j$ from t is changed. To assess the feasibility of this operation in an on-board environment, a simple look-up table using combinatorial logic was synthesized for a rad-hard FPGA from a leading provider of semiconductor solutions. For 1024 values of $t \in [-1, 1)$ spaced in steps of 2^{-9}, and for weights $\tilde{w}_j$ with 9 fractional bits ($\Theta = 9$), the look-up table required 10% of the available space for combinatorial logic and had a delay of 20 ns, which implies that the solution is feasible in most practical cases.

V. EXPERIMENTAL RESULTS

This section assesses the experimental performance of the proposed transform for multi- and hyperspectral on-board image coding. First, the test data set is described, then the experimental setup is detailed, and afterwards results are provided and analyzed.

Experiments have been performed on the comprehensive *Lossy Image Corpus*, as defined by the CCSDS *Space Link Services-Multispectral & Hyperspectral Data Compression* (SLS-MHDC) working group for the purposes of evaluation and testing of lossy coding algorithms. The corpus contains 47 images representative of 14 different sensors comprising multi-, hyper- and ultraspectral sensors. Due to page constraints, results are reported only for a representative selection of 24 images. Technical details of these images are provided in Table II.

The bit depth of images transformed with the original or the isorange POT varies from sensor to sensor, as the hardware capabilities of each sensor determine image bit depth and number of components. In addition, raw images (as acquired by the sensor) may be further processed with non-uniformity-correction (NUC) or calibration stages that may increase their bit depth. No distinction is made between signed and unsigned images, as all images of both types required mean-correction stages.

Table II includes the bit depth of the images when transformed with the original POT as well as with the proposed isorange POT. As expected, the bit depth increase of the isorange transform is limited to three bits (one from mean correction, and two from the lifting network). On the other hand, the original POT results in widely varying bit depth increases, ranging from 2 to 9. Additionally, images transformed with the original POT exceed the bit depth of 16 as supported by CCSDS 122.0-B-1 for most hyperspectral sensors and all the ultraspectral sensors. In comparison, all images having original bit depths no greater than 13 can be easily supported by the standard in concert with the proposed isorange POT.

There are a few images in the database with bit depths greater than 13. These include the noise-filtered radiance SFSI image and those produced by the non-spaceborne AVIRIS sensor. Such images are not seamlessly supported by the proposed scheme. However, several alternatives exist, such as dealing with some of the least-significant bit planes by storing them uncoded or by discarding them — both options would pose non-trivial side effects that would need to be considered in the rate allocation procedure.

Regarding the experimental setup, the POT and the isorange POT are implemented using fixed-point weights of 12 fractional bits, which due to inaccuracies in their computation from t, have a guaranteed accuracy of $\Theta = 9$ (i.e.,

$|\tilde{w}_j - w_j| \le 2^{-10}$). Side information containing the values of t (13-bit fixed-point values) and mean-correction factors is stored uncompressed. For comparison, experimental results are also provided for the reversible CDF 5/3 wavelet transform [34], implemented in a lifting scheme of similar limited dynamic range expansion [33, p. 435].

The proposed isorange POT has been coupled with a software CCSDS 122.0-B-1 2D encoder [25] configured to use the Integer 2D wavelet transform (with 3 levels). We stress again here that the CCSDS 122.0-B-1 2D encoder supports a maximum bit depth of 16, rendering the original POT essentially unusable in this scenario. For the purpose of comparison, we have modified the encoder to extend its bit depth support beyond 16 bits per sample so that experimental results can be presented for the original POT. However, these results would not be obtainable in practice.

The CCSDS 122.0-B-1 encoder is set to produce segments every 8 lines for each transformed component. Rate allocation is performed jointly across segments and bands with a post-compression rate-distortion optimizer [35] using truncation points at the end of each of the five stages of the CCSDS 122.0-B-1 bitplane encoding process [25, p. 4-4].

Experimental results are provided in Table III for the proposed isorange POT, the original POT and the CDF 5/3 wavelet, each applied to exploit inter-component correlation. Lossy results are reported in terms of Signal-to-Noise (SNR; in dB), and as the SNR difference (dB) with respect to the isorange POT. SNR is here computed based on the energy of the input signal, $\text{SNR}_{\text{Energy}} = 10 \cdot \log_{10}\left(\sum(I_{i,j,k})^2 / \sum(I_{i,j,k} - \widehat{I}_{i,j,k})^2\right)$, being I the original image and $\widehat{I}$ the recovered image.

The percentage of additional rate needed to match the SNR of the isorange POT is provided as well. For example, for the AVIRIS-Hawaii-raw image, the proposed isorange POT provides an SNR of 46.21 dB at a bit rate of 1.0 bits per pixel per band (bpppb). At the same rate, the original POT yields an SNR of $46.21 + 0.48 = 46.69$ dB. Alternatively, the original POT provides an SNR of 46.21 dB at a bit rate of $1.0 - 0.16 = 0.84$ bpppb.

The results suggest that the isorange POT suffers a small performance penalty in relation to the original POT. This penalty is on average 0.34 dB in SNR, or 6.5% in rate at 1.0 bpppb, and 0.20 dB in SNR, or 2.3% in rate at 2.0 bpppb. Nonetheless, the proposed transform achieves significantly higher coding performance than that of the CDF 5/3 wavelet transform, averaging 1.28 dB in SNR, or 21.5% in rate at 1.0 bpppb, and 1.17 dB in SNR, or 13.3% in rate at 2.0 bpppb.

It is worth noting that differences in performance between images can be non-negligible. This suggests that particular image or sensor characteristics influence the performance

results. For example, the performance improvements provided by the POTs (with respect to the wavelet transform) are much larger for the three AVIRIS images than for the M3 images. Such differences can be clearly seen in Fig. 10, which plots the rate distortion performance for selected images using each scheme discussed above.

As mentioned previously, CCSDS 122.0-B-1 is a lossy-to-lossless 2D codec. When used in concert with one of the three 1D transforms considered here, the resulting 3D codec is also lossy-to-lossless. Lossless coding results for this scheme are presented in the right portion of Table III. For the purpose of comparison, lossless results also included the lossless-*only* method CCSDS 123.0-B-1 configured as described in [36]. CCSDS-123 is selected as a reference of lossless-only methods, because it is the outcome of a careful algorithm selection among state-of-the-art methods, and because its multiple available implementations (such as [37]) make possible the production of experimental results for the corpus employed. Lossless results are reported in rate required for lossless coding (bpppb). When comparing results for lossless coding, both POTs provide extremely similar performance, often somewhat better than the 5/3 wavelet. As expected, the purely lossless method CCSDS-123 is notably better, at the expense of not providing the capability for lossy coding. Other pure lossless methods may yield similar qualitative results; see [13] for a comparison among pure lossless methods, and see [21] for a comparison between the original POT and 7 other approaches providing lossless compression (recall that the performance of isorange POT as compared to original POT for lossless compression is almost identical).

VI. CONCLUSIONS

This article examines the dynamic range expansion issue of a spectral transform for on-board image coding, and proposes a derived spectral transform —the isorange POT— that inherits all positive features of the original POT while having a limited dynamic range expansion.

The new transform is defined —including lifting network weights and the method to calculate transform gains— and the problem of unpaired components is addressed with a custom operation that takes into account the higher gain of these elements.

In addition, a comprehensive analysis is carried out to mathematically prove a worst-case bound of 3 bits on the dynamic range expansion for any reasonable input image, and the necessary weight precision to achieve said bound is obtained.

Finally, experimental results over a corpus of 47 images indicate that the isorange POT has a small performance penalization in relation to the original POT for lossy image

Table II

TECHNICAL DETAILS OF THE IMAGES EMPLOYED IN THE EXPERIMENTAL RESULTS. IMAGES ARE UNSIGNED UNLESS THEIR BIT DEPTH IS MARKED WITH A †. NUC STANDS FOR (SIMULATED ON-BOARD) NON-UNIFORMITY CORRECTION.

Name	Type	Bands	Height	Width	Bit Depth	Bit Depth After Transform	
						Proposal	Original POT
AIRS-Granule-9	raw	1501	135	90	13	16 (+3)	20 (+7)
AVIRIS-Hawaii-raw	raw	224	512	614	10	13 (+3)	16 (+6)
AVIRIS-Yellowstone-rad	radiance	224	512	677	14 †	17 (+3)	20 (+6)
AVIRIS-Yellowstone-raw	raw	224	512	680	15	18 (+3)	21 (+6)
CASI-t0477f06-nuc	nuc	72	1225	406	13	16 (+3)	18 (+5)
CASI-t0477f06-raw	raw	72	1225	406	12	15 (+3)	17 (+5)
CRISM-sc214-nuc	nuc	74	2700	64	11	14 (+3)	16 (+5)
CRISM-sc214-raw	raw	74	2700	64	11	14 (+3)	16 (+5)
Hyperion-GeoSample-ff	flatfielded raw	242	1024	256	13 †	16 (+3)	19 (+6)
Hyperion-GeoSample-nuc	nuc	242	1024	256	12	15 (+3)	18 (+6)
Hyperion-GeoSample-raw	raw	242	1024	256	12	15 (+3)	18 (+6)
IASI-Desert	calibrated	8461	60	66	12	15 (+3)	21 (+9)
Landsat-Coast	raw	6	1024	1024	8	11 (+3)	11 (+3)
Landsat-Mountain	raw	6	1024	1024	8	11 (+3)	11 (+3)
M3-globalA-nuc	nuc	86	512	320	11	14 (+3)	16 (+5)
M3-globalA-raw	raw	86	512	320	12	15 (+3)	17 (+5)
MODIS-250m-raw	raw	2	8120	5416	12	15 (+3)	14 (+2)
MODIS-500m-raw	raw	5	4060	2708	12	15 (+3)	15 (+3)
MSG-RC15	calibrated	11	3712	3712	10	13 (+3)	14 (+4)
Pleiades-Montpellier	HR, simulated	4	2456	224	12	15 (+3)	15 (+3)
SFSI-Mantar-rad-rmnoise	noise-filtered radiance	240	140	452	16 †	19 (+3)	22 (+6)
SFSI-Mantar-raw	raw	240	140	496	11	14 (+3)	17 (+6)
SPOT5-Toulouse-1	HRG, processed	3	1024	1024	8	11 (+3)	11 (+3)
VEGETATION-1-1c	raw	4	10080	1728	10	13 (+3)	13 (+3)

coding, whereas there is no significant performance difference for lossless compression.

APPENDIX
DERIVATION OF THE ERROR BOUND

The absolute error introduced in the lifting network is bounded by the infinity norm of a 2-dimensional vector $\left(\begin{smallmatrix} v_1 \\ v_2 \end{smallmatrix}\right)$, as given in (15). The bound is derived by bounding each vector element individually.

Throughout the derivation, the following properties of the absolute value operator are employed:

1) sub-additivity, i.e., $|a + b| \leq |a| + |b|$,   (22)
2) multiplicativeness, i.e., $|a \cdot b| = |a| \cdot |b|$,   (23)
3) triangle inequality, i.e., $|a - b| \leq |a - c| + |c - b|$, (24)

4) a bound on the distance between two elements when one is rounded, i.e.,

$$|a - [b]| \leq |a - b| + 2^{-1}, \tag{25}$$

5) and a direct result from (24) and (23),

$$|a \cdot b - c \cdot d| \leq |a \cdot b - c \cdot b| + |c \cdot b - c \cdot d| \\ = |b| \cdot |a - c| + |c| \cdot |b - d|. \tag{26}$$

The magnitude of the first vector element in (15) is

$$|v_1| = \big| (x_i + (w_2 (y_i + (w_1 x_i)))) \\ - (x_i + [\tilde{w}_2 (y_i + [\tilde{w}_1 x_i])]) \big|.$$

Canceling both x_i, employing (25), and distributing w_2 and $\tilde{w}_2$ in the outermost products yields

$$|v_1| \leq |w_2 y_i - \tilde{w}_2 y_i + w_2 (w_1 x_i) - \tilde{w}_2 [\tilde{w}_1 x_i]| + 2^{-1}.$$

By (22), taking common factor of y_i, and (23),

$$|v_1| \leq |w_2 - \tilde{w}_2| |y_i| + |w_2 (w_1 x_i) - \tilde{w}_2 [\tilde{w}_1 x_i]| + 2^{-1}.$$

Then, (26) can be applied on the second term of the summation above to yield

$$|v_1| \leq |w_2 - \tilde{w}_2| |y_i| + |w_1 x_i| |w_2 - \tilde{w}_2| \\ + |\tilde{w}_2| |w_1 x_i - [\tilde{w}_1 x_i]| + 2^{-1}.$$

By (25), taking common factor of x_i, and (23),

$$|v_1| \leq |w_2 - \tilde{w}_2| |y_i| + |w_1 x_i| |w_2 - \tilde{w}_2| \\ + |\tilde{w}_2| \left(|w_1 - \tilde{w}_1| |x_i| + 2^{-1} \right) + 2^{-1}.$$

Table III
PERFORMANCE COMPARISON OF THE ISORANGE POT (PROPOSED), THE ORIGINAL POT, AND THE WAVELET CDF 5/3. RESULTS FOR CCSDS 123.0-B-1 ARE INCLUDED AS REFERENCE FOR PURE LOSSLESS.

Image Name	1 bpppb			2 bpppb			Lossless			
	Prop. (dB)	Orig. POT (Δ dB, %)	CDF 5/3 (Δ dB, %)	Prop. (dB)	Orig. POT (Δ dB, %)	CDF 5/3 (Δ dB, %)	Prop. (bpppb)	Orig. POT (bpppb)	CDF 5/3 (bpppb)	123.0 (bpppb)
AIRS-Granule-9	57.56	0.43 −11%	−2.23 +56%	60.92	0.44 −8%	−1.58 +23%	4.64	4.62	5.01	4.30
AVIRIS-Hawaii-raw	46.21	0.48 −16%	−1.45 +48%	49.85	−0.60 +10%	−2.01 +24%	3.22	3.21	3.51	2.71
AVIRIS-Yellowstone-rad	46.78	0.56 −8%	−2.52 +39%	51.16	0.60 −8%	−1.45 +19%	4.58	4.58	4.86	3.96
AVIRIS-Yellowstone-raw	51.32	0.09 −1%	−4.77 +63%	55.87	0.25 −3%	−2.72 +30%	6.65	6.65	7.13	6.19
CASI-t0477f06-raw	42.76	0.27 −3%	−0.25 +3%	47.92	0.44 −5%	0.08 −1%	5.64	5.64	5.61	5.02
CASI-t0477f06-raw-nuc	43.34	0.34 −4%	−0.70 +9%	48.35	0.48 −6%	−0.21 +3%	5.56	5.56	5.59	4.93
CRISM-sc214-nuc	51.41	0.73 −12%	−1.91 +28%	56.72	−0.25 +3%	−1.67 +27%	3.73	3.74	3.99	2.75
CRISM-sc214-raw	42.51	0.20 −1%	−1.80 +16%	51.27	0.77 −6%	−1.20 +10%	4.69	4.69	4.84	2.74
Hyperion-GeoSample-ff	25.29	0.64 −14%	−1.30 +35%	28.52	0.52 −8%	−0.98 +15%	4.21	4.20	4.39	4.32
Hyperion-GeoSample-raw	44.27	0.30 −4%	−0.85 +15%	49.15	0.32 −4%	−0.99 +13%	4.81	4.81	4.94	4.32
Hyperion-GeoSample-nuc	46.73	0.48 −11%	−0.92 +24%	50.44	0.63 −10%	−0.87 +11%	4.44	4.43	4.60	4.32
IASI-Desert	47.34	0.20 −4%	−0.23 +4%	51.56	0.33 −4%	0.21 −3%	5.65	5.65	5.53	4.76
Landsat-Coast	36.50	0.40 −16%	−0.49 +10%	39.63	−0.05 +1%	−0.58 +6%	3.15	3.14	3.16	2.79
Landsat-Mountain	27.14	0.42 −7%	−1.33 +28%	31.36	0.21 −4%	−1.51 +24%	4.20	4.21	4.37	3.80
M3-M3globalA-nuc	54.59	1.11 −14%	−0.34 +6%	59.40	−0.13 +3%	−0.76 +7%	3.29	3.31	3.34	2.26
M3-M3globalA-raw	45.24	0.47 −4%	−0.76 +6%	54.50	0.33 −2%	−0.33 +2%	4.30	4.38	4.31	2.25
MODIS-250m-raw	36.61	−0.37 +3%	−3.21 +36%	44.26	0.33 −2%	−2.87 +25%	6.16	6.17	6.71	6.43
MODIS-500m-raw	29.81	−0.51 +3%	0.30 −3%	38.20	−0.26 +2%	−1.59 +13%	7.66	7.68	8.02	6.91
MSG-RC15	40.32	0.49 −6%	−1.18 +15%	46.15	0.03 −0%	−1.50 +15%	4.00	4.00	4.17	3.60
Pleiades-Montpellier	29.48	0.00 0%	0.31 −6%	36.65	−0.14 +1%	0.26 −2%	7.91	7.91	7.76	7.43
SFSI-Mantar-rad-rmnoise	50.72	0.93 −13%	−1.76 +23%	56.66	0.05 −0%	−1.56 +20%	3.69	3.64	3.92	3.71
SFSI-Mantar-raw	37.44	0.34 −11%	−0.43 +15%	41.06	0.11 −1%	−0.79 +11%	4.88	4.86	4.97	4.77
SPOT5-Toulouse-1	25.72	0.16 −3%	−2.34 +36%	31.11	0.31 −3%	−2.69 +23%	5.30	5.31	5.70	5.14
VEGETATION-1-1c	35.86	0.02 −0%	−0.67 +10%	42.01	0.14 −2%	−0.75 +5%	5.24	5.25	5.32	5.05

Using the substitutions

$$|\tilde{w}_i - w_i| \le K, |\tilde{w}_i| \le |w_i| + K, |x_i| \le M, |y_i| \le M, \tag{27}$$

where $K = 2^{-\Theta-1}$ and $M = 2^{B-1}$, yields

$$|v_1| \le KM + |w_1|MK + (|w_2| + K) \cdot (KM + 2^{-1}) + 2^{-1}. \tag{28}$$

The magnitude of the second vector element in (15), after canceling y_i, is

$$|v_2| = |(w_1 x_i) + (w_3 (x_i + (w_2 (y_i + (w_1 x_i)))))$$
$$- [\tilde{w}_1 x_i] - [\tilde{w}_3 (x_i + [\tilde{w}_2 (y_i + [\tilde{w}_1 x_i])])]|$$

Using (22) and applying (25) to both summands yields

$$|v_2| \le |(w_1 x_i) - \tilde{w}_1 x_i|$$
$$+ |(w_3 (x_i + (w_2 (y_i + (w_1 x_i)))))$$
$$- \tilde{w}_3 (x_i + [\tilde{w}_2 (y_i + [\tilde{w}_1 x_i])])| + 1.$$

Distributing w_3 and $\tilde{w}_3$ and applying (22) and (23) yields

$$|v_2| \le |w_1 - \tilde{w}_1||x_i| + |w_3 - \tilde{w}_3||x_i|$$
$$+ |w_3(w_2 (y_i + (w_1 x_i)))$$
$$- \tilde{w}_3 [\tilde{w}_2 (y_i + [\tilde{w}_1 x_i])]| + 1. \tag{29}$$

Focusing on the third term of (29)

$$T_1 = |w_3(w_2 (y_i + (w_1 x_i))) - \tilde{w}_3 [\tilde{w}_2 (y_i + [\tilde{w}_1 x_i])]|,$$

and applying (26) and (25) yields

$$T_1 \le |w_2(y_i + (w_1 x_i))| \cdot |w_3 - \tilde{w}_3|$$
$$+ |\tilde{w}_3| \cdot (|w_2 (y_i + (w_1 x_i))$$
$$- \tilde{w}_2 (y_i + [\tilde{w}_1 x_i])| + 2^{-1}). \tag{30}$$

Similarly, focusing on

$$T_2 = |w_2 (y_i + (w_1 x_i)) - \tilde{w}_2 (y_i + [\tilde{w}_1 x_i])|$$

from (30), distributing w_2 and $\tilde{w}_2$, and applying (22) and (23) yields

$$T_2 \le |w_2 - \tilde{w}_2| \cdot |y_i| + |w_2 (w_1 x_i) - \tilde{w}_2 [\tilde{w}_1 x_i]|,$$

on which applying (26) on the second summand and (25) afterwards yields

$$T_2 \le |w_2 - \tilde{w}_2| \cdot |y_i| + |w_1 x_i| \cdot |w_2 - \tilde{w}_2|$$
$$+ |\tilde{w}_2| \cdot (|w_1 - \tilde{w}_1||x_i| + 2^{-1}). \tag{31}$$

Substituting (30) and (31) back into (29) and using the substitutions in (27) gives

$$|v_2| \le KM + KM + (|w_2| + |w_2 w_1|)MK$$
$$+ (|w_3| + K) \cdot (KM + |w_1|MK$$
$$+ (|w_2| + K) \cdot (KM + 2^{-1})$$
$$+ 2^{-1}) + 1. \tag{32}$$

With $K = 2^{-\Theta-1}$ and $M = 2^{B-1}$, (28) and (32) can be rearranged to yield the two arguments of the max function in (16). ∎

(a) AVIRIS-Yellowstone-raw

(b) CASI-t0477f06-raw

(c) MODIS-500m

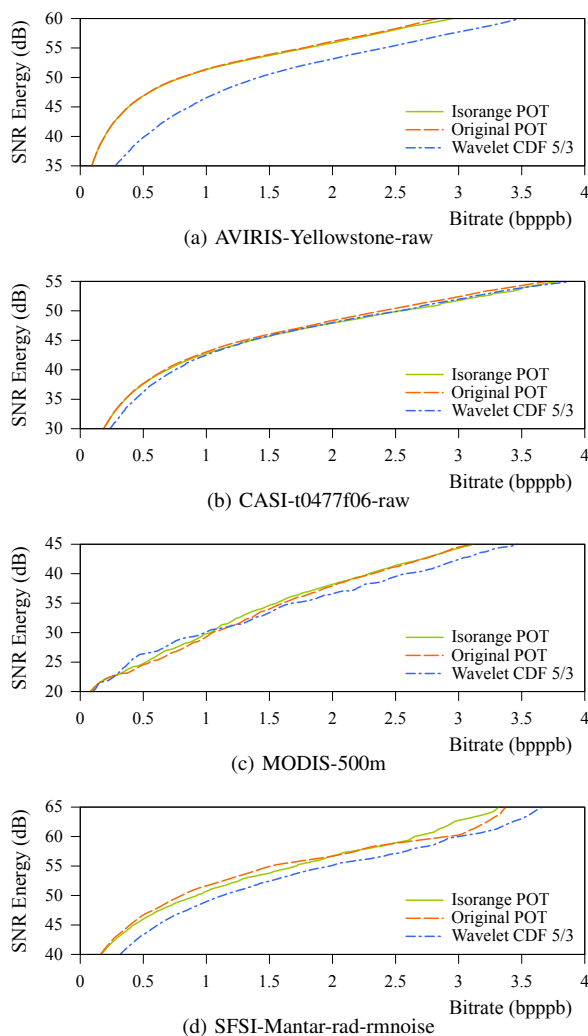(d) SFSI-Mantar-rad-rmnoise

Figure 10.   Rate-distortion performance for selected images. Rate is sampled every 0.01 bpppb. Plots have the same scale.

REFERENCES

[1] S. Qian, *Optical Satellite Signal Processing and Enhancement*. SPIE, Nov. 2013.

[2] J. Zhang, J. E. Fowler, and G. Liu, "Lossy-to-lossless compression of hyperspectral imagery using three-dimensional TCE and an integer KLT," *IEEE Geosci. Remote Sens. Lett.*, vol. 5, pp. 814–818, Oct. 2008.

[3] X. Wu and N. Memon, "Context-based lossless interband compression – extending CALIC," *IEEE Trans. Image Process.*, vol. 9, no. 6, pp. 994–1001, Jun 2000.

[4] 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.

[5] X. Wu and N. Memon, "Context-based, adaptive, lossless image coding," *IEEE Trans. Commun.*, vol. 45, no. 4, pp. 437–444, Apr 1997.

[6] J. Mielikainen, "Lossless compression of hyperspectral images using lookup tables," *IEEE Signal Process. Lett.*, vol. 13, no. 3, pp. 157–160, Mar. 2006.

[7] B. Huang and Y. Sriraja, "Lossless compression of hyperspectral imagery via lookup tables with predictor selection," *Proc. SPIE*, vol. 6365, pp. 63 650L–63 650L-8, Oct. 2006.

[8] J. Mielikainen and P. Toivanen, "Lossless compression of hyperspectral images using a quantized index to lookup tables," *IEEE Geosci. Remote Sens. Lett.*, vol. 5, no. 3, pp. 474–478, July 2008.

[9] B. Aiazzi, L. Alparone, S. Baronti, and C. Lastri, "Crisp and fuzzy adaptive spectral predictions for lossless and near-lossless compression of hyperspectral imagery," *Geoscience and Remote Sensing Letters, IEEE*, vol. 4, no. 4, pp. 532–536, Oct 2007.

[10] Consultative Committee for Space Data Systems (CCSDS). [Online]. Available: http://www.ccsds.org

[11] ——, *Lossless Multispectral & Hyperspectral Image Compression CCSDS 123.0-B-1*, ser. Blue Book. CCSDS, May 2012. [Online]. Available: http://public.ccsds.org/publications/archive/123x0b1ec1.pdf

[12] M. Klimesh, "Low-complexity adaptive lossless compression of hyperspectral imagery," *Proc. SPIE*, pp. 63 000N-1–63 000N-9, Sept. 2006.

[13] A. Kiely and M. Klimesh, "Exploiting calibration-induced artifacts in lossless compression of hyperspectral imagery," *IEEE Trans. Geosci. Remote Sens.*, vol. 47, no. 8, pp. 2672–2678, Aug. 2009.

[14] B. Penna, T. Tillo, E. Magli, and G. Olmo, "A new low complexity KLT for lossy hyperspectral data compression," *IGARSS 2006. IEEE International Conference on Geoscience and Remote Sensing Symposium, 2006.*, pp. 3525–3528, 2006.

[15] ——, "Progressive 3-D coding of hyperspectral images based on JPEG2000," *IEEE Geosci. Remote Sens. Lett.*, vol. 3, no. 1, pp. 125–129, Jan. 2006.

[16] ——, "Transform coding techniques for lossy hyperspectral data compression," *IEEE Trans. Geosci. Remote Sens.*, vol. 45, no. 5, pp. 1408–1421, May 2007.

[17] J. E. Fowler and J. T. Rucker, *Hyperspectral Data Exploitation: Theory and Applications*. Hoboken, NJ, USA: John Wiley & Sons Inc., 2007, ch. "3D wavelet-Based Compression of Hyperspectral Imager", pp. 379–407.

[18] I. P. Akam Bita, M. Barret, and D.-T. Pham, "On optimal orthogonal transforms at high bit-rates using only second order statistics in multicomponent image coding with JPEG2000," *Elsevier Signal Processing*, vol. 90, no. 3, pp. 753 – 758, 2010.

[19] ——, "On optimal transforms in lossy compression of multicomponent images with JPEG2000," *Elsevier Signal Processing*, vol. 90, no. 3, pp. 759–773, 2010.

[20] I. Blanes, J. Serra-Sagrista, M. Marcellin, and J. Bartrina-Rapesta, "Divide-and-conquer strategies for hyperspectral image processing: A review of their benefits and advantages," *IEEE Signal Process. Mag.*, vol. 29, no. 3, pp. 71–81, 2012.

[21] I. Blanes and J. Serra-Sagristà, "Pairwise orthogonal transform for spectral image coding," *IEEE Trans. Geosci. Remote Sens.*, vol. 49, no. 3, pp. 961 – 972, Mar. 2011.

[22] Y. Wongsawat, S. Oraintara, and K. R. Rao, "Integer sub-optimal Karhunen-Loève transform for multi-channel lossless EEG compression," *European Signal Processing Conference*, Sept. 2006.

[23] J. A. Saghri, S. Schroeder, and A. G. Tescher, "Adaptive two-stage Karhunen-Loeve-transform scheme for spectral decorrelation in hyperspectral bandwidth compression," *SPIE Optical Engineering*, vol. 49, p. 057001, May 2010.

[24] I. Blanes and J. Serra-Sagristà, "Cost and scalability improvements to the Karhunen-Loève transform for remote-sensing image coding," *IEEE Trans. Geosci. Remote Sens.*, vol. 48, no. 7, pp. 2854–2863, Jul. 2010.

[25] Consultative Committee for Space Data Systems (CCSDS), *Image Data Compression CCSDS 122.0-B-1*, ser. Blue Book. CCSDS, Nov. 2005. [Online]. Available: http://public.ccsds.org/publications/archive/122x0b1c3.pdf

[26] P. Winterrowd, C. Orbe, J. Venbrux, S. Whitaker, E. Cameron, R. Nelson, G. Maki, D. Fisher, and P.-S. Yeh, "A 320 Mbps flexible image data compressor for space applications," in *Aerospace Conference 2010*, March 2010, pp. 1–6.

[27] R. Vitulli and J.-L. Poupat, "CWICOM: the new CCSDS image compression ASIC," in *3rd Int. WS on On-Board Payload Data Compression (OBPDC)*, Barcelona, Spain, Oct. 2012.

[28] L. Li, G. Zhou, B. Fiethe, H. Michalik, and B. Osterloh, "Efficient implementation of the CCSDS 122.0-B-1 compression standard on a space-qualified field programmable gate array," *Journal of Applied Remote Sensing*, vol. 7, no. 1, pp. 074 595–074 595, 2013.

[29] G. Golub and C. van Loan, *Matrix Computations*. The Johns Hopkins University Press, Oct. 1996.

[30] P. W. Hao and Q. Y. Shi, "Matrix factorizations for reversible integer mapping," *IEEE Trans. Signal Process.*, vol. 49, no. 10, pp. 2314–2324, 2001.

[31] F. A. M. L. Bruekers and A. W. M. van den Enden, "New networks for perfect inversion and perfect reconstruction," *IEEE J. Sel. Areas Commun.*, vol. 10, no. 1, pp. 130–137, 1992.

[32] I. Daubechies and W. Sweldens, "Factoring wavelet transforms into lifting steps," *Journal of Fourier Analysis and Applications*, vol. 4, no. 3, pp. 247–269, 1998.

[33] D. Taubman and M. Marcellin, *JPEG2000: Image Compression Fundamentals, Standards, and Practice*. Kluwer International Series in Engineering and Computer Science, 2002, vol. 642.

[34] A. Cohen, I. Daubechies, and J. Feauveau, "Biorthogonal basis of compactly supported wavelets," *Comm. Pure Appl. Math.*, vol. 45, pp. 485–560, 1992.

[35] D. Taubman, "High performance scalable image compression with EBCOT," *IEEE Trans. Image Process.*, vol. 9, no. 7, pp. 1158–1170, 2000.

[36] E. Augé, J. E. Sánchez, A. Kiely, I. Blanes, and J. Serra-Sagristà, "Performance impact of parameter tuning on the CCSDS-123 lossless multi- and hyperspectral image compression standard," *Journal of Applied Remote Sensing*, vol. 7, no. 1, pp. 074 594 1–074 594 16, 2013.

[37] GICI, "Emporda," http://gici.uab.cat/GiciWebPage/downloads.php, 2012.

Miguel Hernández-Cabronero received two B.S. degrees in computer science and mathematics from the Universidad Autónoma de Madrid in 2010 and the M.S. in computer science from the Universitat Autònoma de Barcelona in 2011. Since 2010, he has been with the Group on Interactive Coding of Images of the Universitat Autònoma de Barcelona, where he is currently pursuing his Ph.D. degree. His main research interest is in data compression, with special focus on medical and biological image coding.

Francesc Aulí-Llinàs (S'06–M'08–SM'14) is a Ramón y Cajal Fellow with the Department of Information and Communications Engineering, Universitat Autònoma de Barcelona. He received the Ph.D. (cum laude) in computer science from Universitat Autònoma de Barcelona in 2006. Since 2002, he has been consecutively funded with doctoral and postdoctoral fellowships in competitive calls. He develops and maintains BOI, a free-source JPEG2000 implementation. In 2000 and 2002, he received two awards of Bachelor given to the first students of the promotion. In 2013, he was awarded with a distinguished R-Letter given by the IEEE Communications Society for a paper co-authored with M. Marcellin. He is reviewer for various magazines and symposiums and has authored numerous papers in the top journals and conferences of his field.

Joan Serra-Sagristà (S'97–M'05–SM'11) received the Ph.D. degree in computer science from Universitat Autònoma Barcelona (UAB), Spain, in 1999. He is currently an Associate Professor at Department of Information and Communications Engineering, UAB. From September 1997 to December 1998, he was at University of Bonn, Germany, funded by DAAD. His current research interests focus on data compression, with special attention to image coding for remote sensing and telemedicine applications. He serves as Associate Editor of IEEE Trans. on Image Processing. He has co-authored over one hundred publications. He was the recipient of the Spanish Intensification Young Investigator Award in 2006.

Ian Blanes (S'05–M'11) received the B.S., M.S., and Ph.D. in computer science from the Universitat Autònoma de Barcelona, in 2007, 2008, and 2010, respectively. Since 2003, he has been with the Group on Interactive Coding of Images of the Universitat Autònoma de Barcelona, where he currently holds a postdoctoral position. In 2010, he was a visiting postdoctoral researcher at the Centre National d'Etudes Spatiales, France. He participates in the CCSDS Multispectral Hyperspectral Data Compression Working Group since 2011. Dr. Blanes was the recipient of the 2007 awards by the Spanish Ministry of Education as second-best computer-science student of Spain.

Michael W. Marcellin (S'81–M'87–SM'93–F'02) received the B.S. in Electrical Engineering from San Diego State University in 1983, and the MS and PhD in Electrical Engineering from Texas A&M University in 1985 and 1987, respectively. Since 1988, Dr. Marcellin has been with the University of Arizona, where he currently holds the title of Regents' Professor, and is the International Foundation for Telemetering Chaired Professor. His research interests include digital communication and data storage systems, data compression, and signal processing. He has authored or coauthored more than two hundred publications in these areas. He has received numerous honors, including six teaching awards.