

ONBOARD/ON-GROUND IMAGE PROCESSING CHAIN FOR HIGH-RESOLUTION EARTH OBSERVATION SATELLITES

E. Cucchetti^a*, C. Latry^a, G. Blanchet^a, J.-M. Delvit^a, M. Bruno^a

^aCentre National d'Etudes Spatiales - CST, 18 Av. Edouard Belin, 31000 France

Commission III, WG III/Iva

KEY WORDS: Compression, Denoising, Image processing chain, Remote Sensing, Restoration

ABSTRACT:

Over the last decade, the French space agency (CNES) has designed and successfully operated high-resolution satellites such as *Pléiades*. High-resolution satellites typically acquire panchromatic images with fine spatial resolutions and multispectral images with coarser samplings for downlink constraints. The multispectral image is reconstructed on the ground, using pan-sharpening techniques. Onboard compression and ground processing affect however the quality of the final product. In this paper, we describe our next-generation onboard/on-ground image processing chain for high-resolution satellites. This paper focuses on onboard compression, compression artefacts correction, denoising, deconvolution and pan-sharpening. In the first part, we detail our fixed-quality compression approach, which limits compression effects to a fraction of the noise, thus preserving the useful information in an image. This approach optimises the bitrate at the cost of image size, which depends on the scene complexity. This technique requires however pre- and post-processing steps. The noisy HR images obtained after decompression are suited for non-local denoising algorithms. We show in the second part of this paper that non-local denoising outperforms previous techniques by 15% in terms of root mean-squared error when tested on simulated noiseless references. Deconvolution is also detailed. In the final part of this paper, we put forward an adaptation of this chain to low-cost CMOS Bayer colour matrices. We demonstrate that the concept of our image chain remains valid, provided slight modifications (in particular dedicated transformations of the colour planes and demosaicing). A similar chain is under investigation for future missions.

1. INTRODUCTION

The use of high resolution (HR) remote sensing images has significantly increased in the past years either for civilian (urbanism, disaster management, meteorology), scientific (hydrology, bathymetry, vegetation oversight) and military uses. Space agencies and private companies have therefore launched Earth observation satellites providing high-resolution images to cover this demand. The French space agency (CNES) has for instance operated HR satellites such as *Pléiades* (Gleyzes et al., 2012) or *CSO* over the last decade, and developed before that HR civilian (*SPOT*) and military programs (*HELIOS*).

Current-generation CNES Earth observation satellites are based on push-broom detectors. To obtain a sufficient signal to noise ratio, high resolution images are often acquired using large panchromatic bands (typically with wavelengths ranging between 400 and 900nm). To recover the colour information of the observed scenes, these finely sampled panchromatic images are coupled with coarser colour images (in general red, green and blue bands) acquired with lower resolution detectors. Under the assumption that most spatial high frequencies in the colour bands are correlated, pan-sharpening techniques recover a final colour image with the native panchromatic resolution (Latry et al., 2012; Vivone et al., 2014).

A drawback of HR digital images is the overall size of the information that needs to be downlinked to ground stations. Despite efforts to reduce the total size, image demand requires thousands of scenes to be acquired daily, closing the door to a lossless downlink of images. This feature drives the recent improvements in data compression algorithms (Zhou et al., 2015). It follows that the quality of the final image obtained on the

ground is highly dependent not only on the hardware of the satellite itself (in particular mirror diameter, satellite stability, detector efficiency), but also on the overall image processing chain applied to raw products on the ground, which provides compression, radiometric and geometric corrections. In this paper, we present the most recent in-flight and on-ground image processing chain developed by CNES for its future satellite applications. In particular, solutions for compression, denoising, image restoration and pan-sharpening are reviewed.

In Section 2, we present the onboard compression solution retained for future CNES Earth observation satellites. Unlike its predecessors, which operate on the image as a whole, this technique takes advantage of the intrinsic complexity of a given scene to provide a local compression rate driven by the image quality (Camarero et al., 2012). It therefore adapts compression thresholds locally to maintain the same level of information over the scenes, while obtaining an average compression rate similar to previous fixed bitrate techniques. However, this approach creates a variable size of compressed images and requires additional pre- and post-processing steps.

Section 3 presents the radiometric restoration chain, which needs to be adapted to the fixed-quality compression used on board. In particular, a noise restoration step is needed before denoising to recover the original noise distribution, followed by denoising, deconvolution and pan-sharpening if needed. Examples of this restoration chain are provided.

Finally, future space missions are not likely to use the same push-broom technology as the current satellites. The development of cheap, low-noise CMOS matrices opens new possibilities for HR imaging, as already demonstrated by constellations such as *Dove* (Wilson et al., 2017), *SuperDove* (Pritchett et al., 2020), or *Skysat* (Murthy et al., 2014) and even for astronomy (CNES nanosatellite *Eyesat*, Carret 2018). In Section 4, we

* Corresponding author: edoardo.cucchetti@cnes.fr

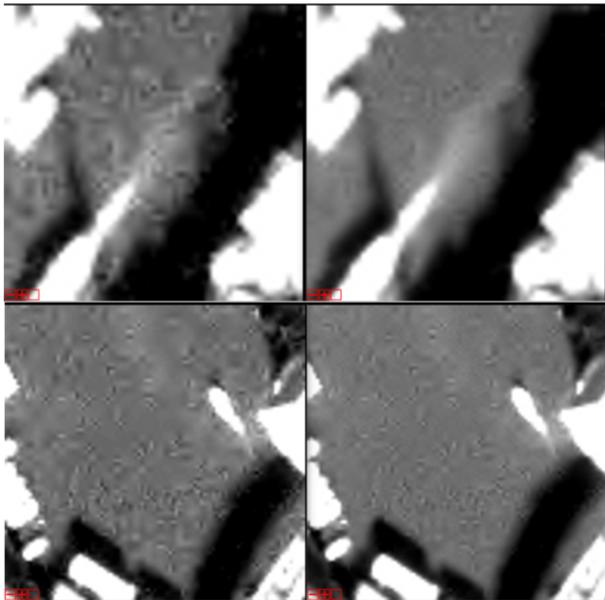


Figure 1. *Pléiades* images using the PAN detector (70cm) over Venice obtained without (Left) and with (right) noise restoration.

extend the chain presented in this paper to the case of matrix detectors, in particular colour filter arrays (CFAs). Examples for Bayer CFAs (Bayer, 1975) are provided.

2. FIXED-QUALITY COMPRESSION

2.1 Fixed-bitrate and fixed-quality compression

High-resolution images acquired by Earth observation satellites count millions of pixels with high bit encoding (typically 10 to 16 bits). Compression becomes therefore needed to ensure a proper downlink of all image products. Typical compression techniques use an orthonormal basis of functions of the 2D plane (Zhou et al. 2015, e.g. discrete cosine functions or discrete wavelet transforms, based on Cohen-Daubechies-Feauveau 9/7 or Le Gall-Tabatabai 5/3 wavelets, see Dua et al. 2020 for a review). The idea for a $N \times M$ image is to separate it into smaller samples and apply the selected transform to project the information onto a orthonormal basis of functions $\Psi_{i,j}$. This decorrelation step is reversible. For a subset of the image of size $K \times K$ the projected sub-image is

$$I_{K \times K}(x, y) = \sum_{i,j \leq K} a_{i,j} \Psi_{i,j}(x, y) \quad (1)$$

Once projected, $a_{i,j}$ coefficients can be downlinked. Each is either transmitted directly (the orthonormal projection reduces the image entropy, thus the total bitrate), or compared to a threshold T and set to zero if smaller (series of zeros are efficiently coded and reduce the size of the bitstream). This is the quantification step which makes the compression lossy and non-reversible. Even when all coefficients are downlinked, if floating point functions $\Psi_{i,j}$ (typically 9/7 wavelets) are used, the quantification causes a truncation error which introduces a loss. For multiple colour planes, spectral decorrelation transforms (Saghri et al., 2010) can be applied before compression to exploit inter-band correlations and reduce the total bitrate.

Fixed-bitrate compression uses a threshold on the $a_{i,j}$ values. For instance, all high-frequency wavelet coefficients of

the image are set to zero, and only low-frequency ones are encoded. A bitrate regulation can also be applied on the encoded $a_{i,j}$ values to enhance compression. Transmitted images then have exactly the same size, regardless of the observed scene, but without a certain level of high frequency information. This compression is extremely efficient and simplifies interfaces by providing images of exactly the same size once decoded.

A similar approach was used on *Pléiades*, whose compressor divided images into blocks, each compressed over the same fixed number of bits through bitrate regulation. Yet, this has a disadvantage of not accounting for the local complexity of the scene. Therefore, quasi-uniform landscapes will be encoded with no losses, while complex scenes will be strongly affected by compression (see artefacts on Figure 1). Instead, fixed-quality compression computes a threshold locally (e.g. a local average) and compares it to a fraction of instrumental noise. Then, only coefficients $a_{i,j}$ which are statistically meaningful (i.e. with high signal-to-noise ratio) are transmitted, while others are set to zero. This technique is extremely flexible, as it adapts compression locally and can be tuned over time. Compression thresholds can also be adapted per block. For space applications, encoding coefficients at the level of the noise is sufficient, but any other fraction can in principle be used.

2.2 Detector noise and variance stabilising transforms

Measurements taken using optical detectors are affected by multiple sources of noise. The total noise for a given signal S (in digital counts) entering the compressor can be modelled using a random Gaussian process of mean $\mu = 0$ and of standard deviation $\sigma = \sqrt{a^2 + bS}$. This model represents a combination of a Poissonian shot noise characterised by the constant b , and a white additive Gaussian noise, where a represents the standard deviation related to the dark current and the readout noise of the detectors. The latter is in fact dependent on integration time and temperature of the detector.

Before compression, each pixel has an intrinsic noise which depends on the signal S . This makes the comparison of coefficients $a_{i,j}$ to the noise level possible (e.g. by using the local average of $I_{K \times K}$) albeit imprecise. An elegant solution is to use variance-stabilising transforms (VSTs), which transform the digital counts of the image into a random variable with constant variance. Values can therefore always be compared to this constant. In our imaging chain, the Anscombe transform is chosen as it is well adapted for Poisson processes such as photon detections. For a given number of counts S in a pixel, the new count estimate S' is obtained by:

$$S' = AT(S) = 2\sqrt{\frac{a^2}{b^2} + \frac{S}{b} + \frac{3}{8}} \quad (2)$$

where a and b are the noise coefficients detailed previously. After this operation, the new variance of the noise is $\sigma = 1$ across the image and becomes independent of the signal in the pixel. This VST is performed on board before the compression and can be reversed on the ground using the inverse formula:

$$S = AT^{-1}(S') = b \left(\frac{S'}{2} \right)^2 - \frac{3}{8}b - \frac{a^2}{b} \quad (3)$$

2.3 Fixed-quality compression implementation

CNES has retained fixed-quality compression for its future missions (Camarero et al., 2012). Images are divided locally into

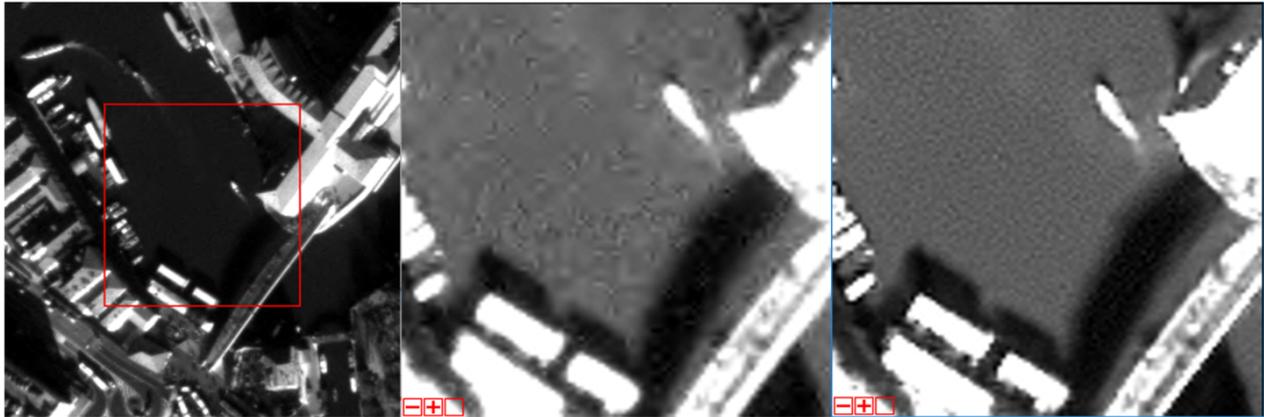


Figure 2. Denoising using NL-Bayes. (Left) Input image over Venice observed by *Pléiades* using the PAN detector (70cm resolution) with a zoom over the central parts of the image before (Middle) and after application of NL-Bayes (Right).

8×8 $9/7$ wavelet coefficients blocks, corresponding to a three-level wavelet decomposition (zones of 57×57 pixels). These include one DC coefficient (the average of the image subset) and 63 AC coefficients resulting from the $9/7$ 2D wavelet local description of the scene.

Instead aiming for a given number of bits, fixed-quality compression compares the basis coefficient to the noise variance locally, assuming a white Gaussian noise of standard deviation $\sigma = 1$ after VST. If $a_{i,j} < k$, coefficients are set to zero, otherwise they are transmitted. k is called the quality parameter and is often taken between 0.5 and 1, to ensure that the wavelet coefficients comparable to the noise are properly encoded ($k = 1$ represents a compression of all wavelet coefficients whose value is of the same level as the instrumental noise). The strength of fixed-quality compression is to perform this comparison locally, thus adapting the compression to the entropy of the image. VSTs are not mandatory when single planes are used, but become an interesting solutions when multiple correlated colour planes need to be downlinked, in particular if spectral decorrelation transforms are used (as discussed in Section 4)

Though providing images of variable sizes, this technique tunes the quality to lose the same level of information in every scene. In average, the obtained bitrate is similar to fixed-bitrate compression algorithms. However, complex scenes are now encoded with more bits (transcribing their higher complexity) while uniform landscapes can be strongly compressed thus gaining on the total bitrate. The final image quality is therefore enhanced using this type of compression.

3. RADIOMETRIC PROCESSING CHAIN

Once received on the ground, the compressed image can be decoded, decompressed and finally restored into its original digital counts through inverse VST, if needed. However, once decoded, the observed scene is affected by both noise (reducing the available information) and the modulation transfer function (or MTF), which attenuates high frequencies and/or introduces aliasing on the observed scenes if the MTF is significantly higher than 0 above the Nyquist frequency.

3.1 Noise restoration and denoising

The overall noise model of a given instrument is well-known and calibrated on the ground before flight using uniform target sources. Additional corrections can be performed in flight,

using for instance uniform scenes. For typical radiances, the quality of the image is related to its signal-to-noise ratio (SNR, i.e. the ratio S/σ). In dark areas, where this ratio is low, denoising becomes necessary to enhance the final image products. We assume that the noise can be approximated (per colour plane) by a white Gaussian noise. Structured noise, such as repeatable features or peaks visible in the spectrum of each image, is ignored and must in any case be considered separately depending on its value. Similarly, we also assume that the image is normalised ahead of our compression step to remove pixel-to-pixel response non-uniformity (harmful for compression).

Most denoising techniques assume a input image with a white noise, independent from the observed scene. This assumption is verified for the raw images, but is unfortunately no longer true once compression has been applied (Section 2). Raw decompressed images must therefore pass through two separate denoising steps:

- First, before radiometric processing, noise must be restored to obtain a noise distribution of the image compatible with the one observed before compression. In particular a Gaussian behaviour must be restored.
- Second, the image goes through the actual denoising process to reduce the noise whilst maintaining the high-frequency information of the scene.

Each of these two steps is detailed in the following subsections.

3.1.1 Noise restoration

Once received on the ground, the image has lost part of its information due to compression. Although the information removed is lower than a fraction of the instrumental noise in our fixed-quality approach, compression modifies the noise distribution which is no longer Gaussian. Yet, to ensure a proper application of denoising algorithms, a Gaussian distribution is required. A noise restoration step is thus introduced.

Knowing the compression threshold set in flight, it is possible to estimate the resulting noise distribution on the ground. The noise can then be restored by drawing random values of $a_{i,j}$ at the same level of the suppressed ones. This operation is performed in wavelet space using the normal distribution of the in-flight noise model (whose variance is 1 if a VST is used). After inverse discrete wavelet transform (DWT) and inverse VST the noise distribution is once again Gaussian, recovering the original behaviour of the scene (Delvit et al., 2018). Though this

step purposely introduces noise, it improves the subsequent denoising, based on Gaussian assumptions (Figure 1 and 2).

3.1.2 Denoising

Denoising aims at reducing the noise of an image using either thresholding, interpolation or filtering methods. Denoising techniques on previous CNES satellites included complex thresholding after deconvolution using wavelet packet transforms. Over the last decade, non-local (NL) methods started nonetheless to provide a more robust and accurate way for denoising (Menon and Calvagno, 2011).

The underlying assumption of these methods is to consider that scenes are self-similar, i.e. that any given image contains – at small scales – similar features. In this case, for any subset of pixels (e.g. a square of $k \times k$ pixels, called a “patch”) it is possible to find within the image similar patches (at the sense of the L_2 norm for instance) which can be considered drawn from the same random Gaussian process, whose average is the patches’ average and whose covariance matrix Σ can be computed analytically. Under these assumptions, for each family of similar patches one can compute and apply a good approximation of the optimal Wiener filter, thus denoising the images in a locally-optimal way in terms of likelihood. Several of these techniques (NL-means Buades et al. 2005, NL-Bayes Lebrun et al. 2013, BM3D Dabov et al. 2009) are currently implemented for space applications and constitute state-of-the-art denoising techniques. NL-Bayes is retained in our image processing scheme and was used in the *Pléiades* ground segment (see Figure 2 for an example).

3.2 Deconvolution

A known effect of telescopes is to behave as low-pass filters causing the attenuation – and ultimately the cancellation – of frequencies beyond the spatial cut-off frequency, smoothing the observed scenes and damaging the quality of the final image. This feature is linked to the instrument’s impulse response, the point spread function (PSF) in detector space, or equivalently the MTF in frequency space.

Provided a good knowledge of the MTF, telescope effects can be corrected by the image processing chain through a deconvolution step to recover sharpness. For remote sensing applications, the MTF is low at the Nyquist frequency to reduce aliasing on the final products and to maintain reasonable telescope diameters. After denoising, the image obtained by the detectors appears however blurred, and must be post-processed through deconvolution to restore high frequencies.

For noiseless images, deconvolution is a direct multiplication in the frequency space by the inverse of the system MTF, noted MTF^{-1} . However, a low MTF near the Nyquist frequency implies an oscillating deconvolution filter in detector space with infinite spatial extent. Further, for noised data, a direct multiplication by MTF^{-1} can result into extreme amplifications of the high spatial frequencies, whose signal-to-noise ratio is low in general. This can then increase the noise and reduce the final quality. Deconvolution must therefore be adapted to the properties of the optical system.

A compromise can be found by aiming for prolate or arctan functions, which enhance mid-range frequencies, while avoiding an excessive increase of high-frequency noise. Wiener-Tikhonov regularisation technique is used to obtain such filters \mathcal{D} (see Figure 3), defined in frequency space (f_x, f_y) by:

$$\mathcal{D}(f_x, f_y) = \frac{MTF^*(f_x, f_y)}{|MTF(f_x, f_y)|^2 + \frac{f_x^2 + f_y^2}{s}} \quad (4)$$

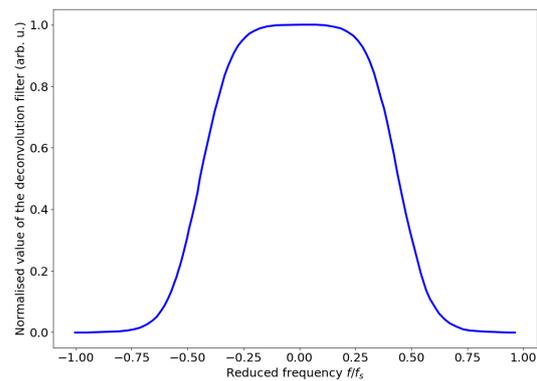


Figure 3. Target deconvolution filter \mathcal{D} (1D slice) obtained with the Wiener-Tikhonov method as a function of the reduced spatial frequency (frequency f over the Nyquist frequency f_s)

where MTF^* stands for the complex conjugate of MTF measured in-flight on ground targets or stars. The s parameter weights the regularisation term penalising the L_2 norm of the gradient. High values provide a sharp deconvolution; low values lead instead to smoother images. In our chain, $s = 6$ was found to provide the best visual results for human operators.

3.3 Pan-sharpening

Most HR satellites acquire both a panchromatic spectral band and red, green, blue, near infrared bands at a coarser resolution than the panchromatic band. This is justified by the constraints upon the onboard downlink rate, combined with a lower SNR of high-resolution multispectral bands due to narrower spectral bandwidths. However, multispectral HR images can be synthesised on ground through pan-sharpening techniques.

Problems may arise from aliasing, which is usually significant in colour bands due to the Pan/colour band resolution ratio (typically 4). CNES pan-sharpening technique minimises the impact of multispectral aliasing by computing, from the panchromatic band, low-resolution (LR) aliased panchromatic images $Pan_i^{LR, alias}$ with the same MTF and sampling grid as each multispectral band B_i . The idea is that the ratio $B_i/Pan_i^{LR, alias}$ may be resampled in the high-resolution Pan geometry to produce a high resolution B_i^{HR} band using (Latry et al., 2012):

$$B_i^{HR} = Pan^{HR} \times \frac{B_i}{PAN_i^{LR, alias}} \quad (5)$$

Applying a low-pass filtering stage to the aliased initial multispectral band B_i is advised to get rid of aliasing artefacts, but induces a slight blurring to the final product. Finally, applying previous noise restoration and NL-Bayes denoising stages not only on the panchromatic band as initially planned, but also on the multispectral aliased bands, proves to be very valuable for the final pan-sharpened image quality. An example of a pan-sharpened raw image, its post-processed version and the result after pan-sharpening is provided Figure 4. Benefits of noise restoration on a pan-sharpened image are illustrated Figure 5.

4. EXTENSION TO MATRIX-BASED DETECTORS

Over the last decade, the paradigm of Earth remote sensing is shifting from push-broom detectors to CMOS-based matrix detectors due to simpler implementation and reduced costs. The

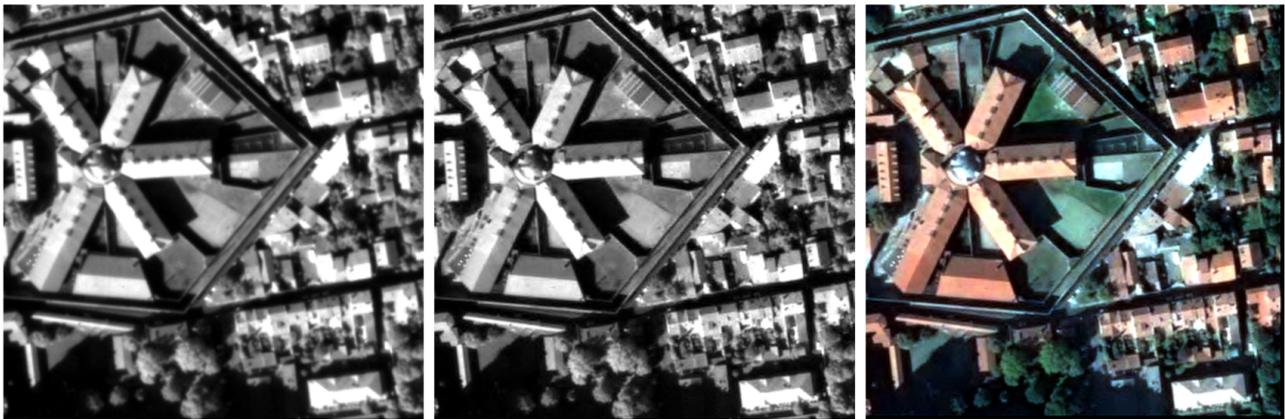


Figure 4. Restoration of a *Pleiades* HR image over Toulouse at a 70cm resolution on the ground. (Left) Raw decompressed image, (Centre) Noise restored, denoised and deconvolved image (Right) Pan-sharpened image using R, G and B channels at 2.8m resolution.



Figure 5. Pan-sharpening of a *Pleiades* image. Zoom over a road crossing restored without (Left) and with (Right) prior noise restoration. Adapted from Figure 6 from Delvit et al. 2018.

use of matrices, in particular chromatic ones such as Bayer detectors (Bayer, 1975), comes with its challenges. In this section, we explain the main differences caused by a Bayer RGB detectors in our high-resolution image processing chain. We assume that these detectors are made of a four-pixel pattern of two quincunx green detectors, completed by blue and red pixels.

4.1 Spectral decorrelating transforms

Unlike on monochromatic detectors, when observing using a Bayer detector each point on the ground is seen either in red, in green or in blue. An accurate compression of the planes as explained in Section 2 can no longer rely on the assumption of redundant information in the nearby spatial pixel, since each colour plane is sparse.

Information present in colour bands remains however correlated. The previous compression method can be adapted using inter-band correlations (notably luminance/chrominance transformations Saghri et al. 2010). To properly decorrelate the information between R, G and B bands, the matrix is truncated into four sub-planes (two green, one red and one blue plane), each shifted by up to one pixel in the vertical and horizontal direction with respect to each other. Knowing the excellent ground sampling resolution of HR pixels (typically lower than 1 m) and assuming a slowly-varying scene, decorrelation relies on linear combinations of colour planes, e.g. luminance/chrominance transformations already used in full-frame compression algorithms (Rao and Yip, 2000). The optimal transformation is found through the covariance matrix of the

planes and by performing a *simile* Karhunen-Loève transform (KLT) of the subsets (Saghri et al., 2010; Rao and Yip, 2000).

In the image processing chain presented here, the transformation matrix is derived from a sample of Bayer images by computing their joint covariance matrix and finding its corresponding eigenvector basis. This operation can be reversed after compression (product of orthogonal matrices). Although not exactly a KLT of each colour plane (as the subsets are shifted), the final matrix is in fact similar to a luminance/chrominance transformation (Saghri et al., 2010) with an average of the four planes and two-by-two plane differences.

This transformation reduces the download bandwidth by an average of 15% for a fixed-quality compression when tested on various sceneries (urban, forest, desert). This approach requires however an accurate knowledge of the noise model for each transformed plane. The use of an Anscombe transform before the KLT is once again necessary (as linear combinations of planes include differences between the red, green or blue planes) and is included in the processing chain.

4.2 Demosaicing and recovery of colour planes

Once the data is received on the ground, the user is provided with a sparse image of each colour plane. Ultimately, these images need to be interpolated to recover the colour planes at the native resolution of the detectors. This step is called demosaicing. This interpolation can be performed using multiple techniques, e.g. Adams and Hamilton 1996; Malvar et al. 2004 (see Menon and Calvagno 2011; Li et al. 2008 for a review). For space applications, computation time is a critical parameter. A good compromise comes from gradient-based techniques such as derivations of the Hamilton-Adams (HA) approach (Adams and Hamilton, 1996). For the chain currently under development green pixels are interpolated as follows (Figure 6, top left). First we compute the horizontal and vertical gradients:

$$\Delta H = |G_4 - G_6| + |R_5 - R_3 + R_5 - R_7| \quad (6)$$

$$\Delta V = |G_2 - G_8| + |R_5 - R_1 + R_5 - R_9| \quad (7)$$

Depending on their local value, we then interpolate as follows:

- If $\Delta H > 1.5\Delta V$, $G_5 = \frac{(G_2+G_8)}{2} + \frac{(R_5-R_1+R_5-R_9)}{4}$
- If $\Delta V > 1.5\Delta H$, $G_5 = \frac{(G_6+G_4)}{2} + \frac{(R_5-R_3+R_5-R_7)}{4}$
- Else, $G_5 = \frac{(G_2+G_4+G_6+G_8)}{4} + \frac{(4R_5-R_1-R_3-R_7-R_9)}{8}$

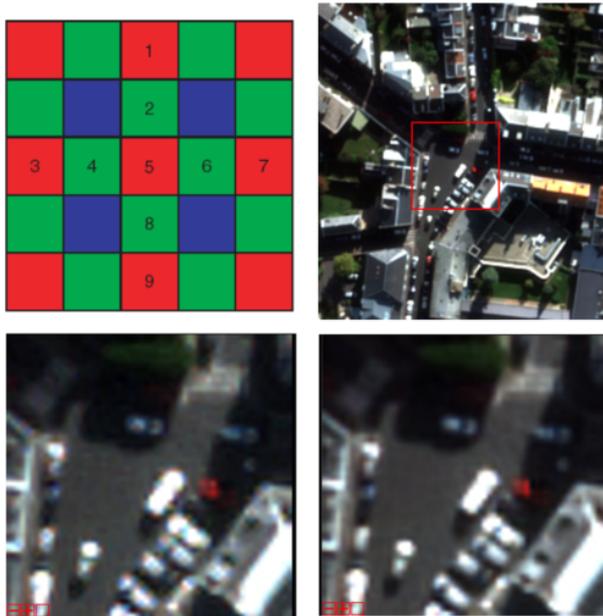


Figure 6. (Top left) Bayer matrix with numbering for HA technique. (Top right) Input simulated image (Amiens, France) for demosaicing with noise. Image processed using HA before (bottom left) and after (bottom right) our denoising approach.

One of HA's weaknesses is the creation of 45° artefacts (Malvar et al., 2004). This was corrected using a 1.5 factor on the gradient, which provided the best results numerically and cancelled most 45° artefacts. The blue and red pixels are in turn interpolated. Using a slowly-varying hue assumption, we first interpolate $R - G$ (resp. $B - G$) through bicubic interpolation, then at a green pixel location i, j we compute $R_{i,j} = (R - G)_{i,j} + G_{i,j}$ (respectively $B_{i,j} = (B - G)_{i,j} + G_{i,j}$). This approach provides reconstruction of the colour planes with few artefacts and with a reasonable computation time. Zipper artefacts are most likely to appear on the image, but are however corrected by the NL denoising (see Figure 6, bottom).

4.3 Denoising

The use of colour matrices provides additional complexities to the denoising step. The noise models between colour planes can indeed differ due to the properties of each colour pixel. Thus, denoising must be performed after a VST to use the same noise model in all the pixels. The nominal chain performs a VST after the HA demosaicing and before denoising the R, G, B plane using NL-Bayes (an inverse VST is performed after denoising, see Figure 7). By doing so however, the interpolated pixels found during demosaicing do not follow the same noise model as the data pixels, since they are obtained by linear combination of noisy pixels of different noise variances. Two approaches were considered to solve this issue: perform the denoising before the demosaicing (when pixels have the same noise model thanks to the VST) or create an additional denoising step after NL-Bayes for interpolated pixels only.

The first solution provides denoising on smaller images ($N/2$ and $N/4$ for green and red/blue pixels respectively for an image of size N) with lower resolutions, thus decreasing the total computational time. However, this approach is incapable to properly take advantage of correlations between colour planes as bands are shifted, and provided systematically less accurate results than the original chain in our tests. Instead, the

	HA only	NL-Bayes	Our chain
R (true)	1	0.87	0.87
R (interpolated)	2.44	1.55	1.21
G (true)	1	0.83	0.83
G (interpolated)	2.11	1.53	1.02
B (true)	1	0.83	0.83
B (interpolated)	1.96	1.30	1.01

Table 1. Root Mean Squared in VST space (1 being the true noise level) comparison between input simulation image and restored image over a pool of 50 images for random noise iterations. Our chain outperforms NL-Bayes for interpolated pixels due to an additional denoising step.

second solution takes advantage of the correlation between colour planes to obtain information of an interpolated pixel of a given patch using the information of measured pixel in another similar patch (similarly to Menon and Calvagno 2011; Li et al. 2008). This additional step improves significantly the quality of the final product on interpolated pixels, with low impacts on the computational time (see Table 1).¹

4.4 First results on Bayer detectors

Colour filter arrays are increasingly used for remote sensing, notably in recent reduced-cost imaging missions such as CubeSat missions, e.g. the *Dove* constellation (Wilson et al., 2017) or *Eyesat* (Carret, 2018). First results of these missions and their detectors show promising leads for these technologies. An example of this restoration chain applied to an uncompressed *Eyesat* image over Australia is provided in Figure 8. Application to real HR Bayer images is expected in the coming months on future CNES satellites. This image processing chain is not exclusively dedicated to satellites, but could also be used on planetary rovers and in theory on any digital cameras based upon Bayer matrix. Example of this image chain restoration applied on MAHLI cameras (Edgett et al., 2012) (*Curiosity* Mars rover) is also provided². Applications on RMI camera aboard *Perseverance* (Maki et al., 2020) and on MMX rover navigation cameras (Ulamec et al., 2019) are also planned.

First colour filter array satellite images also put forward specific effects of these detectors, notably, signal mixing and aliasing related to the use of spectrally correlated bands (e.g. red and green spectral bands overlap and thus are strongly correlated), or the presence of a small level of crosstalk between colour bands, which needs to be included in post-processing techniques. The advent of future large scale missions using these techniques will require a good understanding of these features to ensure the final image quality.

5. CONCLUSION

The growing demand for high resolution satellite images requires to continuously improve both the hardware and software related to image processing. In particular, a deeper and more complex interaction between onboard algorithms and ground image processing chains are used.

In this paper, we have presented our next-generation on-ground/in-flight imaging chain for high resolution satellite applications. In particular, we have described the interest of using fixed-quality compression over typical fixed-bitrate approaches,

¹ This technique is not detailed here for confidentiality issues but will be detailed in forthcoming papers.

² <https://pds-imaging.jpl.nasa.gov/data/msl/>

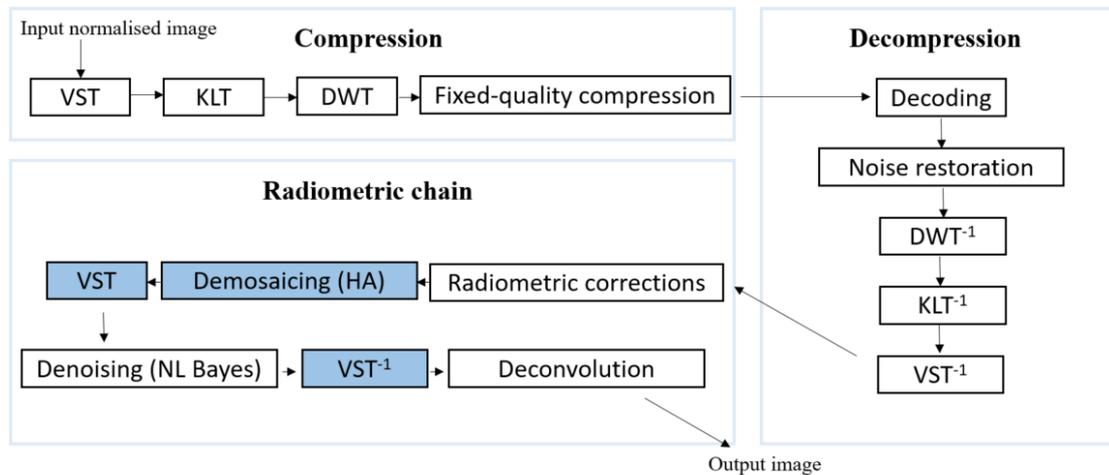


Figure 7. Block diagram of the overall image processing chain. Only the steps discussed in this paper are shown (normalisation is for instance considered for the input image, but not shown). The blocks filled in light blue colour are specific to matrix-based detectors.

as well as the way image restoration is planned to be performed on the ground (denoising, deconvolution, pan-sharpening).

With the advent of low-cost matrices of arrays, the same image processing chain has been adapted for the case of colour filter arrays (in particular Bayer arrays) with only minor modifications. In theory, only an additional demosaicing step is required, but we have underlined how inter-channel correlations can be used efficiently both during compression and restoration stages to improve image quality even further. We have also demonstrated that this image chain can be extended to any other satellite or planetary rover using the same detectors.

The main interest of this image chain is not only to obtain a high-quality visual product, but also to pre-process the image for further treatments such as 3D reconstruction using stereo images, classification and identification of objects, or even deep-learning analysis (e.g. deepzoom, aliasing correction). The image chain presented here is expected to be used on future CNES mission in the upcoming years.

ACKNOWLEDGEMENTS

The authors would like to thank the *Eyesat* team (in particular N. Verdier, F. Viaud and C. Virmontois) for their help and access to raw images, and Q. Douaglin for his post-processing of MAHLI image cameras (data and support were provided by the PDS Cartography and Imaging Sciences Node).

References

Adams, J. E., Hamilton, J. F. J., 1996. Adaptive color plane interpolation in single sensor color electronic camera.

Bayer, B. E., 1975. Color imaging array.

Buades, A., Coll, B., Morel, J., 2005. A non-local algorithm for image denoising. *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, 2, 60–65 vol. 2.

Camarero, R., Delaunay, X., Thiebaut, C., 2012. Fixed-quality/variable bit-rate on-board image compression for future CNES missions. B. Huang, A. J. Plaza, C. Thiebaut (eds), *Satellite Data Compression, Communications, and Processing VIII*, 8514, International Society for Optics and Photonics, SPIE, 9 – 22.

Carret, R., 2018. Mechanical assembly integration and test of the qualification model for the eyesat mission.

Dabov, K., Foi, A., Katkovnik, V., Egiazarian, K., 2009. BM3D Image Denoising with Shape-Adaptive Principal Component Analysis. R. Gribonval (ed.), *SPARS'09 - Signal Processing with Adaptive Sparse Structured Representations*, Inria Rennes - Bretagne Atlantique, Saint Malo, France.

Delvit, J.-M., Thiebaut, C., Latry, C., Blanchet, G., Camarero, R., 2018. A pipeline to improve compressed image quality. Z. Sodnik, N. Karafolas, B. Cugny (eds), *International Conference on Space Optics — ICSSO 2018*, 11180, International Society for Optics and Photonics, SPIE, 2705 – 2713.

Dua, Y., Kumar, V., Singh, R. S., 2020. Comprehensive review of hyperspectral image compression algorithms. *Optical Engineering*, 59(9), 090902.

Edgett, K. S., Yingst, R. A., Ravine, M. A., Caplinger, M. A., Maki, J. N. et al., 2012. Curiosity's Mars Hand Lens Imager (MAHLI) Investigation. *Sp. Sci. Reviews*, 170(1), 259-317.

Gleyzes, M. A., Perret, L., Kubik, P., 2012. Pleiades system architecture and main performances. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 39(1), 537–542.

Latry, C., Fourest, S., Thiebaut, C., 2012. Restoration technique for Pleiades-HR panchromatic images. *Int. Arch. Photogram., Remote Sens. Spatial Inf. Sci.*, 39, 555–560.

Lebrun, M., Buades, A., Morel, J.-M., 2013. Implementation of the "Non-Local Bayes" (NL-Bayes) Image Denoising Algorithm. *Image Processing On Line*, 3, 1–42.

Li, X., Gunturk, B., Zhang, L., 2008. Image demosaicing: a systematic survey. W. A. Pearlman, J. W. Woods, L. Lu (eds), *Visual Communications and Image Processing 2008*, 6822, International Society for Optics and Photonics, SPIE, 489 – 503.

Maki, J. N., Gruel, D., McKinney, C., Ravine, M. A. et al., 2020. The Mars 2020 Engineering Cameras and Microphone on the Perseverance Rover: A Next-Generation Imaging System for Mars Exploration. *Space Science Reviews*, 216(8), 137.

Malvar, H. S., Li-wei He, Cutler, R., 2004. High-quality linear interpolation for demosaicing of bayer-patterned color images. *2004 IEEE International Conference on Acoustics, Speech, and Signal Processing*, 3, iii–485.

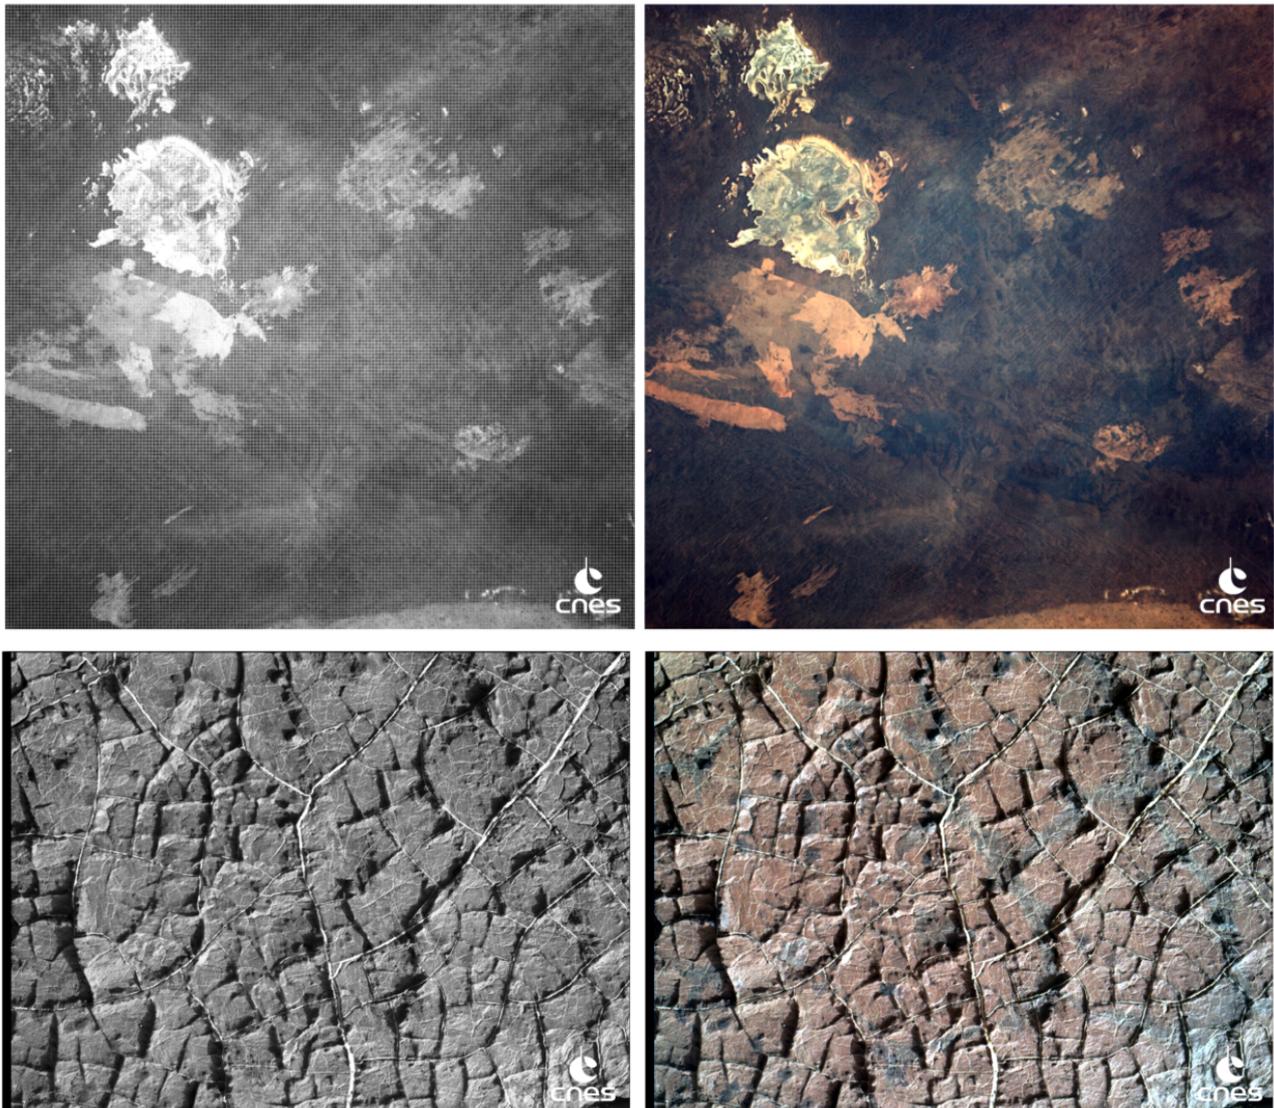


Figure 8. (Top) Eyesat image taken over Lake Maurice, Australia, in December 2019 (local time around 5:30). Copyright: CNES. (Bottom) MAHLI image over Beech Mountain. (Left) Raw Bayer image. (Right) Coloured image post-processed with the chain presented here. White balance is not performed in these images after processing.

Menon, D., Calvagno, G., 2011. Color image demosaicking: An overview. *Signal Processing: Image Communication*, 26(8), 518-533.

Murthy, K., Shearn, M., Smiley, B. D., Chau, A. H., Levine, J., Robinson, M. D., 2014. SkySat-1: very high-resolution imagery from a small satellite. R. Meynart, S. P. Neeck, H. Shimoda (eds), *Sensors, Systems, and Next-Generation Satellites XVIII*, 9241, International Society for Optics and Photonics, SPIE, 367 – 378.

Pritchett, C., Smith, N., Jumpasut, A., Fernandez-Saldivar, J., Zuleta, I., 2020. The Spectral Response of Planet Doves: Pre-launch Method and Results.

Rao, K. R., Yip, P. C., 2000. *The Transform and Data Compression Handbook*. CRC Press.

Saghri, J. A., Schroeder, S., Tescher, A. G., 2010. Adaptive two-stage Karhunen-Loeve-transform scheme for spectral decorrelation in hyperspectral bandwidth compression. *Optical Engineering*, 49(5), 1 – 7.

Ulamec, S., Michel, P., Grott, M., Böttger, U., Hübers, H., Murdoch, N., Vernazza, P., Karatekin, Ö., Knollenberg, J., Willner, K. et al., 2019. A rover for the JAXA MMX mission to phobos. *Proceedings of the International Astronautical Congress, IAC*, 2019-October.

Vivone, G., Alparone, L., Chanussot, J., Dalla Mura, M., Garzelli, A., Licciardi, G. A., Restaino, R., Wald, L., 2014. A critical comparison among pansharpening algorithms. *IEEE Transactions on Geoscience and Remote Sensing*, 53(5), 2565–2586.

Wilson, N., Greenberg, J., Jumpasut, A., Collison, A., 2017. In-Orbit Radiometric Calibration of the Planet Dove Constellation.

Zhou, S., Deng, C., Zhao, B., Xia, Y., Li, Q., Chen, Z., 2015. Remote sensing image compression: A review. *2015 IEEE International Conference on Multimedia Big Data*, IEEE, 406–410.