The Consultative Committee for Space Data Systems

Report Concerning Space Data System Standards

SPECTRAL PRE-PROCESSING TRANSFORM FOR MULTISPECTRAL & HYPERSPECTRAL IMAGE COMPRESSION

INFORMATIONAL REPORT

CCSDS 120.3-G-1

GREEN BOOK

March 2019

The Consultative Committee for Space Data Systems

Report Concerning Space Data System Standards

SPECTRAL PRE-PROCESSING TRANSFORM FOR MULTISPECTRAL & HYPERSPECTRAL IMAGE COMPRESSION

INFORMATIONAL REPORT

CCSDS 120.3-G-1

GREEN BOOK

March 2019

AUTHORITY

Issue:	Informational Report, Issue 1
Date:	March 2019
Location:	Washington, DC, USA

This document has been approved for publication by the Management Council of the Consultative Committee for Space Data Systems (CCSDS) and reflects the consensus of technical panel experts from CCSDS Member Agencies. The procedure for review and authorization of CCSDS Reports is detailed in *Organization and Processes for the Consultative Committee for Space Data Systems* (CCSDS A02.1-Y-4).

This document is published and maintained by:

CCSDS Secretariat
National Aeronautics and Space Administration
Washington, DC, USA
E-mail: secretariat@mailman.ccsds.org

FOREWORD

Through the process of normal evolution, it is expected that expansion, deletion, or modification of this document may occur. This Report is therefore subject to CCSDS document management and change control procedures, which are defined in *Organization and Processes for the Consultative Committee for Space Data Systems* (CCSDS A02.1-Y-4). Current versions of CCSDS documents are maintained at the CCSDS Web site:

http://www.ccsds.org/

Questions relating to the contents or status of this document should be sent to the CCSDS Secretariat at the e-mail address indicated on page i.

At time of publication, the active Member and Observer Agencies of the CCSDS were:

Member Agencies

- Agenzia Spaziale Italiana (ASI)/Italy.
- Canadian Space Agency (CSA)/Canada.
- Centre National d'Etudes Spatiales (CNES)/France.
- China National Space Administration (CNSA)/People's Republic of China.
- Deutsches Zentrum für Luft- und Raumfahrt (DLR)/Germany.
- European Space Agency (ESA)/Europe.
- Federal Space Agency (FSA)/Russian Federation.
- Instituto Nacional de Pesquisas Espaciais (INPE)/Brazil.
- Japan Aerospace Exploration Agency (JAXA)/Japan.
- National Aeronautics and Space Administration (NASA)/USA.
- UK Space Agency/United Kingdom.

Observer Agencies

- Austrian Space Agency (ASA)/Austria.
- Belgian Federal Science Policy Office (BFSPO)/Belgium.
- Central Research Institute of Machine Building (TsNIIMash)/Russian Federation.
- China Satellite Launch and Tracking Control General, Beijing Institute of Tracking and Telecommunications Technology (CLTC/BITTT)/China.
- Chinese Academy of Sciences (CAS)/China.
- China Academy of Space Technology (CAST)/China.
- Commonwealth Scientific and Industrial Research Organization (CSIRO)/Australia.
- Danish National Space Center (DNSC)/Denmark.
- Departamento de Ciência e Tecnologia Aeroespacial (DCTA)/Brazil.
- Electronics and Telecommunications Research Institute (ETRI)/Korea.
- European Organization for the Exploitation of Meteorological Satellites (EUMETSAT)/Europe.
- European Telecommunications Satellite Organization (EUTELSAT)/Europe.
- Geo-Informatics and Space Technology Development Agency (GISTDA)/Thailand.
- Hellenic National Space Committee (HNSC)/Greece.
- Hellenic Space Agency (HSA)/Greece.
- Indian Space Research Organization (ISRO)/India.
- Institute of Space Research (IKI)/Russian Federation.
- Korea Aerospace Research Institute (KARI)/Korea.
- Ministry of Communications (MOC)/Israel.
- Mohammed Bin Rashid Space Centre (MBRSC)/United Arab Emirates.
- National Institute of Information and Communications Technology (NICT)/Japan.
- National Oceanic and Atmospheric Administration (NOAA)/USA.
- National Space Agency of the Republic of Kazakhstan (NSARK)/Kazakhstan.
- National Space Organization (NSPO)/Chinese Taipei.
- Naval Center for Space Technology (NCST)/USA.
- Research Institute for Particle & Nuclear Physics (KFKI)/Hungary.
- Scientific and Technological Research Council of Turkey (TUBITAK)/Turkey.
- South African National Space Agency (SANSA)/Republic of South Africa.
- Space and Upper Atmosphere Research Commission (SUPARCO)/Pakistan.
- Swedish Space Corporation (SSC)/Sweden.
- Swiss Space Office (SSO)/Switzerland.
- United States Geological Survey (USGS)/USA.

DOCUMENT CONTROL

Document	Title	Date	Status
CCSDS 120.3-G-1	Spectral Pre-Processing Transform for Multispectral & Hyperspectral Image Compression, Informational Report, Issue 1	March 2019	Original issue

CONTENTS

CONTENTS (continued)

CONTENTS (continued)

Figure Page

CONTENTS (continued)

Table Page

CONTENTS (continued)

Table Page

1 INTRODUCTION

1.1 PURPOSE

This report presents a summary of the key operational concepts and rationale that underlie the requirements for the CCSDS Recommended Standard, *Spectral Preprocessing Transform for Multispectral and Hyperspectral Image Compression* (reference [1]). Supporting performance information along with illustrations is also included. This report provides a broad tutorial overview of the compression algorithm in reference [1] and is aimed at helping first-time readers to understand the Recommended Standard.

1.2 SCOPE

This document provides supporting and descriptive material only: **it is not part of the Recommended Standard**. In the event of any conflict between the Spectral Pre-Processing Transform for Multispectral & Hyperspectral Image Compression Recommended Standard and the material presented herein, the Recommended Standard is the authoritative document.

1.3 DOCUMENT STRUCTURE

This document is organized as follows.

- Section 2 gives an overview of the compression algorithm formalized in the Recommended Standard and describes some of its features. This section is intended to provide high-level information for a user considering using the Recommended Standard, whether the user intends to produce an implementation or obtain one from a third party.
- Section 3 describes compression options and parameters that may be adjusted by a user or an implementer of the Recommended Standard. Examples are presented to provide an indication of how compression may be affected by parameter selections.
- Section 4 describes rate allocation methods to produce compressed image streams that meet a given bitrate budget or a given quality requirement.
- Section 5 describes issues that would be of concern for a user attempting to produce a software or hardware implementation of the Recommended Standard.
- Section 6 provides results describing lossless and lossy compression performance of the Recommended Standard on test images, along with comparisons to performance results for other compression algorithms.
- Section 7 documents some of the considerations and motivations that influenced the selection of the Recommended Standard and its features.

- Annex A summarizes the corpus of hyperspectral and multispectral images used for compression testing and evaluation in the course of developing the Recommended Standard.
- Annex B provides links to available software implementations of the Recommended Standard, to test data and to look-up tables for different intermediate values described in the Recommended Standard.
- Annex C provides detailed compression results.
- Annex D discusses the compressor selection criteria.
- Annex E provides a derivation of the weights used in the rate allocation methods for the Integer Wavelet Transform and the Pairwise Orthogonal Transform.
- Annex F provides a list of abbreviations and acronyms used in the text of this document.
- Annex G contains a glossary of terms employed throughout this report.
- Annex H provides the default compressor settings employed to generate experimental results.

1.4 CONVENTIONS AND DEFINITIONS

The following convention applies throughout this document:

> 'The Recommended Standard' refers to the *Spectral Pre-Processing Transform for Multispectral & Hyperspectral Image Compression* Recommended Standard described in reference [1].

For the purpose of this document, the mathematical notation and definitions described in subsection 1.7 of reference [1] apply.

1.5 TEST IMAGES AND SOFTWARE IMPLEMENTATIONS

Results and examples in this document make use of the set of test images described in annex A. Available software implementation of the Recommended Standard and a set of verification test data are described in annex B.

1.6 REFERENCES

The following publications are referenced in this document. At the time of publication, the editions indicated were valid. All publications are subject to revision, and users of this document are encouraged to investigate the possibility of applying the most recent editions of the publications indicated below. The CCSDS Secretariat maintains a register of currently valid CCSDS publications.

[1] *Spectral Preprocessing Transform for Multispectral and Hyperspectral Image Compression*. Issue 1. Recommendation for Space Data System Standards (Blue Book), CCSDS 122.1-B-1. Washington, D.C.: CCSDS, September 2017.

[2] *Image Data Compression*. Issue 2. Recommendation for Space Data System Standards (Blue Book), CCSDS 122.0-B-2. Washington, D.C.: CCSDS, September 2017.

[3] *Lossless Multispectral & Hyperspectral Image Compression*. Issue 1-S. Recommendation for Space Data System Standards (Historical), CCSDS 123.0-B-1-S. Washington, D.C.: CCSDS, (May 2012) December 2018.

[4] *Low-Complexity Lossless and Near-Lossless Multispectral and Hyperspectral Image Compression*. Issue 2. Recommendation for Space Data System Standards (Blue Book), CCSDS 123.0-B-2. Washington, D.C.: CCSDS, February 2019.

[5] A. Cohen, I. Daubechies, and J.-C. Feauveau. "Biorthogonal Bases of Compactly Supported Wavelets." *Communications on Pure and Applied Mathematics* 45, no. 5 (June 1992): 485–560.

[6] *Information Technology—JPEG 2000 Image Coding System: Core Coding System*. 2nd ed. International Standard, ISO/IEC 15444-1:2004. Geneva: ISO, 2004.

[7] I. Blanes and J. Serra-Sagrista. "Pairwise Orthogonal Transform for Spectral Image Coding." *IEEE Transactions on Geoscience and Remote Sensing* 49, no. 3 (2011): 961–972.

[8] Ian Blanes, et al. "Isorange Pairwise Orthogonal Transform." *IEEE Transactions on Geoscience and Remote Sensing* 53, no. 6 (2015): 3361–3372.

[9] Shen-En Qian. "Hyperspectral Data Compression Using a Fast Vector Quantization Algorithm." *IEEE Transactions on Geoscience and Remote Sensing* 42, no. 8 (2004): 1791–1798.

[10] C. Thiebaut, et al. "On-Board Compression Algorithm for Satellite Multispectral Images." In *Proceedings of the Data Compression Conference (DCC'06) (28–30 March 2006, Snowbird, Utah)*, 467. Piscataway, NJ: IEEE, 2006.

[11] C. Thiebaut, et al. "CNES Studies of On-Board Compression for Multispectral and Hyperspectral Images." In *Satellite Data Compression, Communication, and Processing III*, 668305-1–668305-15. Edited by Roger W. Heymann, Bormin Huang, and Irina Gladkova. Proceedings of SPIE Vol. 6683. Bellingham, Washington: SPIE, 19 September 2007.

[12] Barbara Penna, et al. "Transform Coding Techniques for Lossy Hyperspectral Data Compression." *IEEE Transactions on Geoscience and Remote Sensing* 45, no. 5 (2007): 1408–1421.

[13] David Taubman and Michael Marcellin. *JPEG2000: Image Compression Fundamentals, Standards and Practice*. Kluwer International Series in Engineering and Computer Science. Norwell, Massachusetts: Kluwer, November 2001.

[14] F. P. Preparata and S. J. Hong. "Convex Hulls of Finite Sets of Points in Two and Three Dimensions." *Communications of the ACM* 20, no. 2 (Feb. 1977): 87–93.

[15] C. Bradford Barber, David P. Dobkin, and Hannu Huhdanpaa. "The Quickhull Algorithm for Convex Hulls." *ACM Transactions on Mathematical Software* 22, no. 4 (Dec. 1996): 469–483.

[16] Donald E. Knuth. *The Art of Computer Programming, Volume 2: Seminumerical Algorithms*. 3rd ed. London: Pearson Education, 1998.

[17] I. Blanes, et al. "Almost Fixed Quality Rate-Allocation under Unequal Scaling Factors for On-Board Remote-Sensing Data Compression." *International Journal of Remote Sensing* 39, no. 7 (2018): 1953–1970.

[18] Michel Barret, et al. "Lossy Hyperspectral Images Coding with Exogenous Quasi Optimal Transforms." In *Proceedings of the 2009 Data Compression Conference*, 411–419. Piscataway, New Jersey: IEEE Conference Publications, 2009.

[19] Monika Stadnicka, et al. "Mitigating Discontinuities in Segmented Karhunen-Loeve Transforms." In *Proceedings of 2016 IEEE International Conference on Image Processing (ICIP) (25–28 Sept. 2016, Phoenix, AZ, USA)*, 2211–2215. N.p.: n.p., 2016.

[20] G. Blanchet, et al. "PLEIADES-HR Innovative Techniques for Radiometric Image Quality Commissioning." *International Archives of the Photogrammetry, Remote Sensing, and Spatial Information Science* 39-B1 (2012): 513–518.

[21] *Lossless Multispectral and Hyperspectral Image Compression*. Issue 1. Report Concerning Space Data System Standards (Green Book), CCSDS 120.2-G-1. Washington, D.C.: CCSDS, December 2015.

[22] *Information Technology—Lossless and Near-Lossless Compression of Continuous-Tone Still Images*. International Standard, ISO/IEC 14495-1:1999. Geneva: ISO, 1999.

[23] *Information Technology—JPEG 2000 Image Coding System: Extensions*. International Standard, ISO/IEC 15444-2:2004. Geneva: ISO, 2004.

2 OVERVIEW

2.1 INTRODUCTION

Recommended Standard CCSDS 122.1-B-1 *Spectral Preprocessing Transform for Multispectral and Hyperspectral Image Compression* (reference [1]) defines a payload data compressor applicable to multispectral and hyperspectral imagers and sounders. The compressor is intended to be suitable for use on board spacecraft; in particular, the algorithm complexity and memory usage are designed to be sufficiently low to make high-speed hardware implementation feasible. The Recommended Standard extends the (two-dimensional) CCSDS 122.0-B-2 *Image Data Compression* Recommended Standard (reference [2]) by providing an effective method of coding three-dimensional data in lossless and lossy regimes.

The purpose of this document is to provide guidance on the usage and implementation of a CCSDS 122.1-B-1 compressor. The three main topics discussed herein are the choice of reasonable parameters in different scenarios, rate-allocation methods to distribute a bit budget, and implementation issues.

This guidance is also applicable to users who want to employ a payload data compressor for multispectral or hyperspectral instruments, for which a near-real-time operation or a fixed data throughput are required, and for which non-exact image recovery is acceptable.

This Recommended practice is applicable as well to hardware manufacturers who are delivering or considering the development of CCSDS 122.1-B-1 integrated circuits and who wish to obtain better understanding of how their devices might be used, or wish to obtain insight on how to maximize the image coding performance of their devices.

Data compression methods may be *lossless* or *lossy*. Under lossless compression, the original image can be reproduced exactly by the decompressor, while under lossy compression, quantization or other approximations used in the compression process result in a reconstructed image that does not exactly match the original. The Recommended Standard is capable of providing both lossless and lossy compression. Notwithstanding, CCSDS 123.0-B-1 (reference [3]) will generally produce better coding performance at lower computational complexity for purely lossless compression. Therefore the CCSDS 122.1-B-1 Recommended Standard is intended for use cases that involve lossy compression. An extension of the CCSDS 123.0-B-1 Recommended Standard (reference [3]), CCSDS 123.0-B-2 (reference [4]) provides a low-complexity approach designed primarily for high fidelity, ‘near-lossless’ compression that allows users to limit the maximum error introduced in the decompressed images.

Metrics for quantifying lossless or lossy compression performance are defined in 3.1.

In this document, it is assumed that the reader is familiar with the contents of reference [1], including terminology defined there.

2.2 IMAGES

This subsection defines parameters and notation pertaining to an image.

The Recommended Standard defines four images in the encoding process:

- An *input* image to be encoded, denoted by I;
- an *upshifted input image*, denoted by $I^{\uparrow}$, resulting from the application of the upshift stage to the input image;
- a *transformed image*, denoted by T, resulting from the application of a spectral transform to the upshifted input image;
- a *downshifted transformed image*, denoted by $T^{\downarrow}$, resulting from the application of the downshift stage to the transformed image.

The properties defined in this subsection apply to all four images (I, $I^{\uparrow}$, T, and $T^{\downarrow}$). For notational convenience, A is used to denote a generic image that could represent any of these four images.

An image A is a three-dimensional array of integer samples $A_{x,y,z}$, where the index z indicates the spectral band, and x and y are indices in the spatial dimensions. For I and $I^{\uparrow}$, indices x, y, and z take on integer values in the ranges $0 \le x < N_X$, $0 \le y < N_Y$, and $0 \le z < N_Z$, where N_X, N_Y, and N_Z are the number of samples along each image dimension. For T and $T^{\downarrow}$, indices x and y take on values in the same ranges, while index z takes on values $0 \le z < N_{TZ}$, where N_{TZ} is the number of components after the spectral transform stage.

The *bit depth* of an image A is the number of bits, n, used to represent each image sample. An image A may be either *unsigned* or *signed*. Samples of unsigned images satisfy $0 \le A_{x,y,z} < 2^n - 1$, while samples of signed images satisfy $-2^{n-1} \le A_{x,y,z} < 2^{n-1} - 1$.

2.3 ALGORITHM OVERVIEW

2.3.1 INTRODUCTION

This subsection provides a high-level overview of the compressor. Detailed descriptions are contained in reference [1].

The input to the compressor is a signed or unsigned image. The compressed image output is an encoded bitstream from which an exact or approximate reconstruction of the input image can be recovered. The compressor, illustrated in figure 2-1, consists of two main functional parts: a spectral transform and a set of 2D encoders. The coded segments produced by the 2D encoders are arranged in groups and interleaved to produce the compressed bit stream. In addition, upshift and downshift processing stages provide potential means of improving the tradeoff between compression ratio and reconstructed image quality, as well as facilitating re-use of existing hardware by limiting the maximum bit depth of the image input to the 2D encoders.

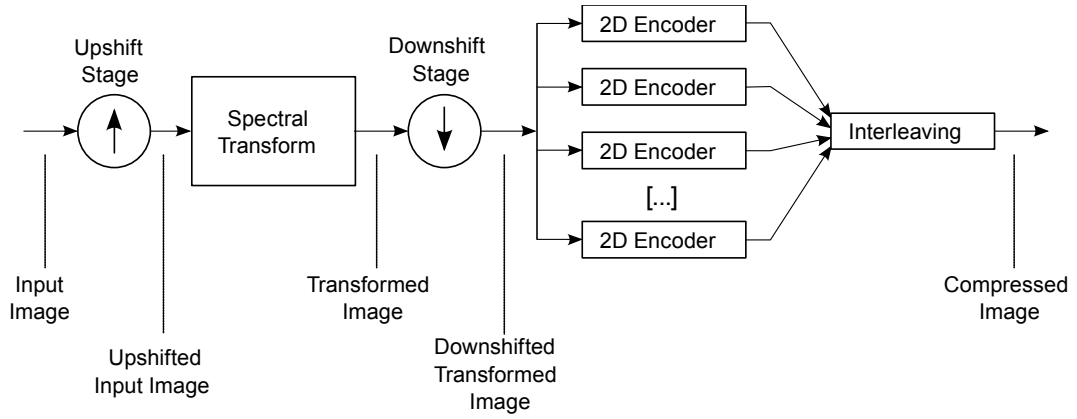

Figure 2-1: Top Level Diagram of the Algorithm

The purpose of the spectral transform is to exploit the similarities between the spectral bands of an image cube, creating transformed images that can be more efficiently compressed by the 2D encoders.

2.3.2 UPSHIFT AND DOWNSHIFT

The upshift stage scales an image by multiplying all image samples by a power of 2, effectively appending additional least significant bits to each sample. This is done to increase the bit depth of an input image to match the maximum bit depth supported by a given implementation of the transform stage, which tends to improve coding performance for high-fidelity lossy compression.

The downshift stage divides each sample in an image by 2^D for a given downshift parameter D and rounds the result to the nearest integer. The downshift stage is used to reduce the bit depth of the transformed image to match the bit depth supported by the 2D encoders. Notwithstanding, because of the rounding operations, downshifting is only guaranteed to reduce bit depth by $D - 1$ bits instead of D.

Lossless compression generally suffers from poor performance when the upshift stage is used (see 3.2.2), and becomes impossible when the downshift stage is used. Therefore, neither upshift nor downshift operations should be applied for lossless compression.

2.3.3 SPECTRAL TRANSFORM

2.3.3.1 General

The Recommended Standard defines four different spectral transforms, to accommodate different combinations of imaging sensors and hardware capabilities. Computational resource requirements vary depending on the transform and the image dimensions. Thus, given limited processing capabilities, a more computationally complex transform might be feasible

for a multispectral imager but not for a hyperspectral imager, because of the differing number of spectral bands to be processed.

Most of the transforms specified are *reversible*, which means that the original input image may be exactly recovered, provided that the transformed image does not suffer any loss of information either from the downshift stage or lossy compression from a 2D encoder.

Transforms are specified with bit-level precision. Thus, two compliant transform implementations, using the same parameters and given the same input image, will produce identical outputs. This facilitates efforts to validate an implementation.

2.3.3.2 Identity Transform

The identity transform performs no operation on the upshifted input image. That is, the transformed image is simply equal to the upshifted input image. This transform is included for the sake of providing a compressed data structure to encode a 3D image without requiring the implementation of a nontrivial transform stage.

2.3.3.3 Integer Wavelet Transform

The Integer Wavelet Transform (IWT) included in the Recommended Standard is a reversible transform that has low computational complexity and provides a noticeable improvement in coding performance over the identity transform.

The IWT is known in the literature as the CDF 5/3 (reference [5]) and is the same as the one used in the JPEG2000 standard (reference [6]). While a variable number of levels of transform decomposition can be used in JPEG2000, in the Recommended Standard, the number is fixed at five. The five-level IWT transforms the spectral vector at each spatial location (i.e., the sequence of samples $(A_{x,y,0}, A_{x,y,1}, \ldots, A_{x,y,N_z-1})$ for a fixed (x, y) pair) independently. Figure 2-2 depicts the IWT processing structure, specified in subsection 4.4 of reference [1].

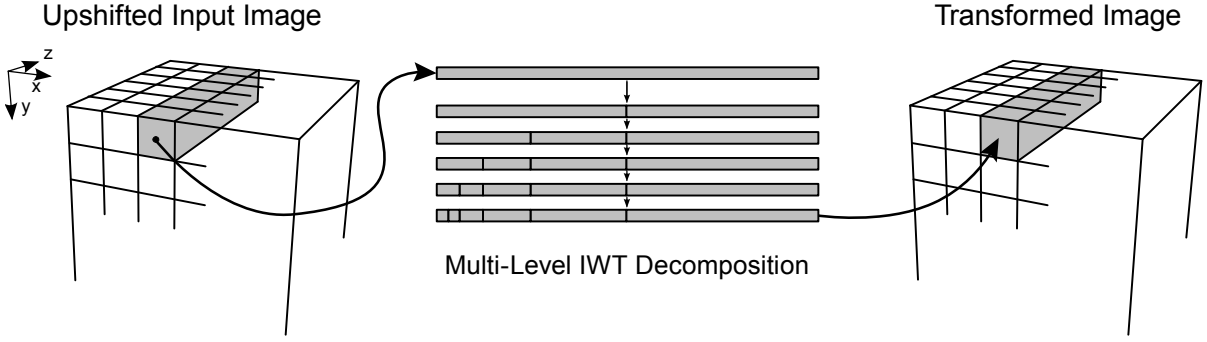

Figure 2-2: Multi-Level IWT Decomposition

Under the IWT, the number of transformed image bands (N_{TZ}) is equal to the number of input image bands (N_z).

The multi-level IWT decomposition is produced by five successive applications of the single-level IWT. L_n and H_n denote the low- and high-frequency wavelet coefficients outputs of the single-level IWT decomposition. The first level of decomposition is applied to the original image, and successive levels are iteratively applied to the low-frequency wavelet coefficients, L_{n-1}, produced by the previous level. The output of the five-level IWT decomposition is defined as the concatenation of the following sequences: L_5, H_5, H_4, H_3, H_2, and H_1, as shown in figure 2-3.

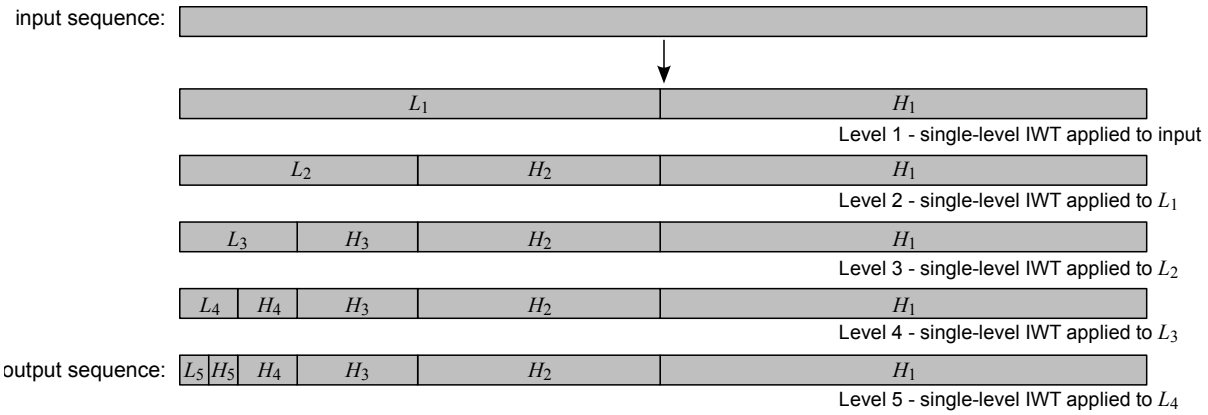

Figure 2-3: Five-Level IWT Decomposition

The inverse 5-level IWT is computed by inverting each successive single-level IWT decomposition in the reverse order in which they were applied.

2.3.3.4 Pairwise Orthogonal Transform

The Pairwise Orthogonal Transform (POT) is an approximation of the transform that provides perfect decorrelation, the Karhunen-Loève Transform (KLT), at a fraction of its computational cost. The POT generally provides better coding performance than the IWT, but with higher implementation complexity.

The POT is described in subsection 4.5 of reference [1], based on the description in references [7] and [8].

The POT is applied separately to each of the image regions produced by partitioning an input image along the *y* axis, as specified in subsection 4.5.6 of reference [1]. Each of these regions has user-defined height *F*, except possibly the last one, which may be shorter. The sequence of samples belonging to the *z*th band and the *i*th region of an image, taken in raster-scan order, is denoted $\Theta_{z,i}$.

The POT is computed by first performing a mean-subtraction operation on each sequence $\Theta_{z,i}$ independently, and then applying pairwise operations organized in a multilevel structure, as depicted in figure 2-4.

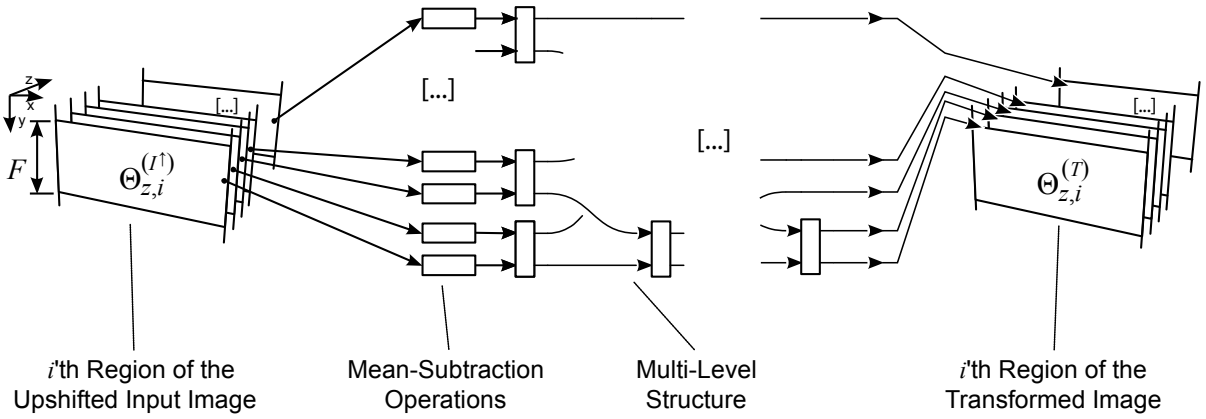

Figure 2-4: Overview of the POT

The mean-subtraction operation is intended to ensure that the mean value of each mean-subtracted sequence is approximately zero. This can be accomplished by subtracting the closest integer m to the actual mean of the sequence. As specified in subsection 4.5.3.3 of reference [1], a user may use a different method of selecting the nominal mean value m to be subtracted.

Mean-subtracted sequences are then processed by multiple levels of pairwise operations. The number of levels is $L = \lceil \log_2(N_Z) \rceil$, unless $N_Z = 1$, in which case there is only one level, and the output is simply equal to the mean-subtracted input sequence with no pairwise operation performed.

As its name suggests, each pairwise operation is applied to a pair of input sequences and produces a pair of output sequences. At any given level, if the number of sequences to be further processed is odd, there is a single unpaired sequence that is not modified by a pairwise operation.

Examples of levels with even and odd numbers of input sequences can be seen in figure 2-5.

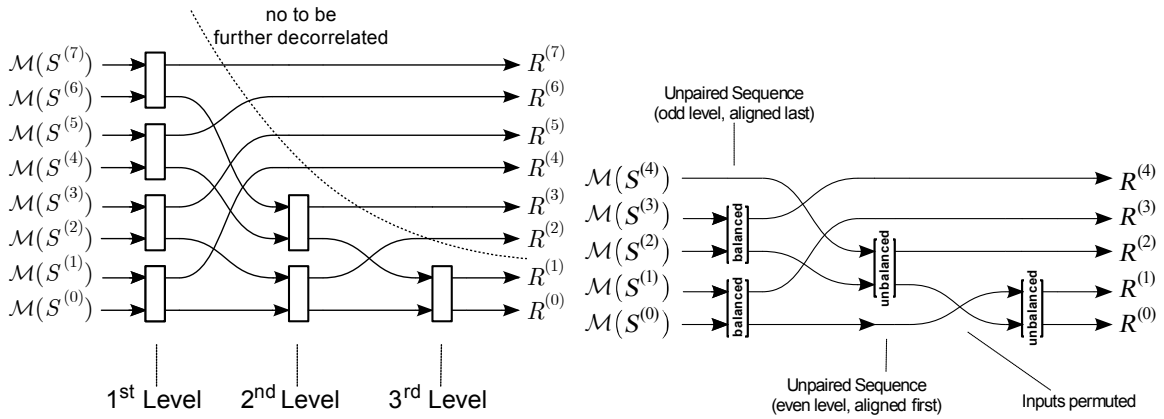

Figure 2-5: Application of the POT Multi-Level Structure for 8 (Left) and 5 (Right) Spectral Bands

As depicted in figure 2-6, each pairwise operation has three inputs: the primary input sequence $\{x_0,x_1,\ldots,x_{N-1}\}$, the secondary input sequence $\{y_0,y_1,\ldots,y_{N-1}\}$, and a Boolean *flip*. The outputs of each operation are the principal output $\{x_0,x_1,\ldots,x_{N-1}\}$, the detail output $\{\phi_0,\phi_1,\ldots,\phi_{N-1}\}$, and the training parameters B and C.

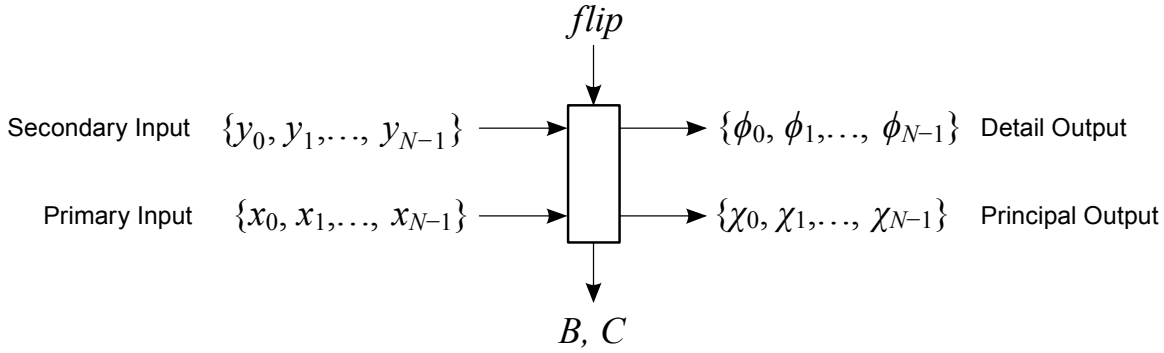

Figure 2-6: The Pairwise Operation

For best compression performance, training parameters B and C should be calculated from the primary and secondary input sequences and the user-specified precision Ω, as specified in subsection 4.5.4 of reference [1]; however, users are permitted to use alternative (e.g., lower complexity) methods of selecting these values. Subsection 4.5.4.4 of reference [1] specifies how the training parameters in turn determine the values of weights used to calculate each POT operation.

The flip input can be used to mitigate discontinuities that can appear at the boundary of image regions. The POT can be inverted even when such discontinuities occur, but compression performance may be reduced.

In bypass mode, each *flip* input is false. Under stable mode, the value of each *flip* input depends on the values of B and C for the current pairwise operation and the corresponding

pairwise operation for the previous region, as specified in subsection 4.5.7 of reference [1] and depicted in figure 2-7.

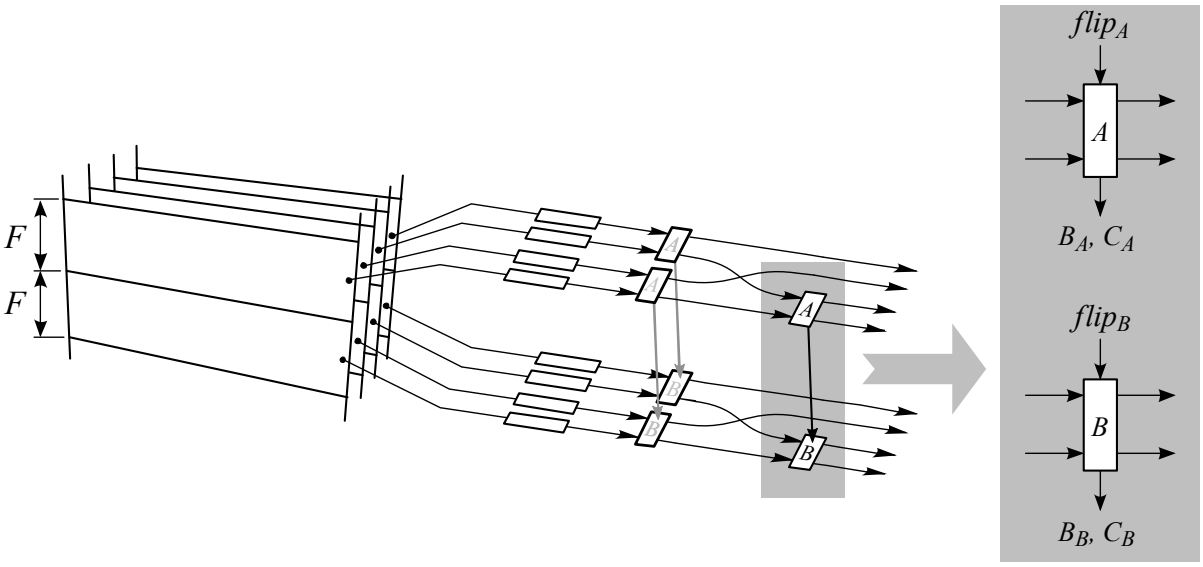

Figure 2-7: Dependencies of the *flip* Parameters

2.3.3.5 Arbitrary Affine Transform

The Arbitrary Affine Transform (AAT) maps an upshifted input image to a transformed image via one or more user-specified one-dimensional affine operations. An affine operation is independently applied to the spectral vector at each spatial location in the upshifted input image.

If X denotes a length-N_Z input sequence and Y the mapped output sequence, then the AAT can be viewed as depicted in figure 2-8.

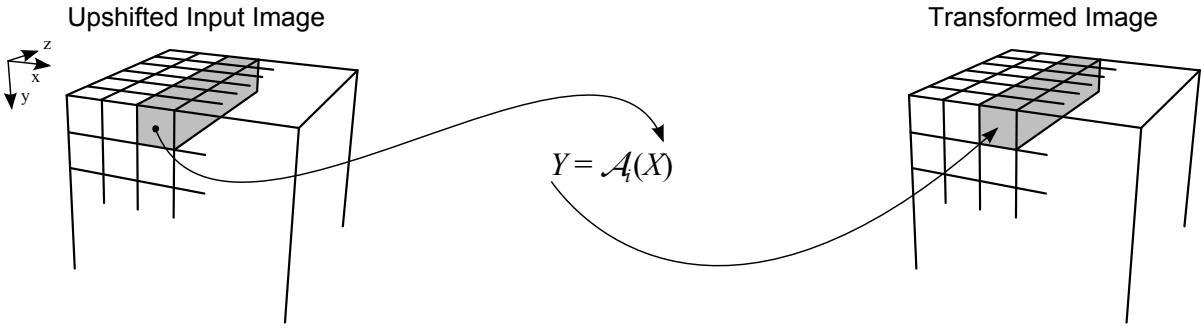

Figure 2-8: Arbitrary Affine Transform Operation

As described in subsection 4.6.3 of reference [1], the AAT calculates an integer approximation of the affine transform $Y = \frac{1}{2^{\Psi}} \cdot Q(X + V)$, where X is the input sequence of

upshifted samples of a given region, Q and V are a matrix and a vector respectively, each having user-defined integer elements, and Y is the output sequence of transformed samples.

The user-defined parameter Ψ is an integer in the range 0,…,31 that determines the arithmetic precision employed in the intermediate calculations. Larger values of Ψ allow for more precision and possibly better compression performance at the cost of higher implementation resources.

2.3.4 ENCODER

In the encoding stage, each band of the downshifted transformed image, $T^{\downarrow}$, is independently coded by a 2D encoder. This produces a bitstream composed of multiple coded segments, which are grouped and interleaved to produce the compressed bit stream.

For each group of coded segments, a header is created to encode information describing the spectral pre-processing transform. The group and its associated header are called a *collection*. The compressed image consists of all collections, as specified in subsection 5.1 of reference [1] and illustrated in figure 2-9.

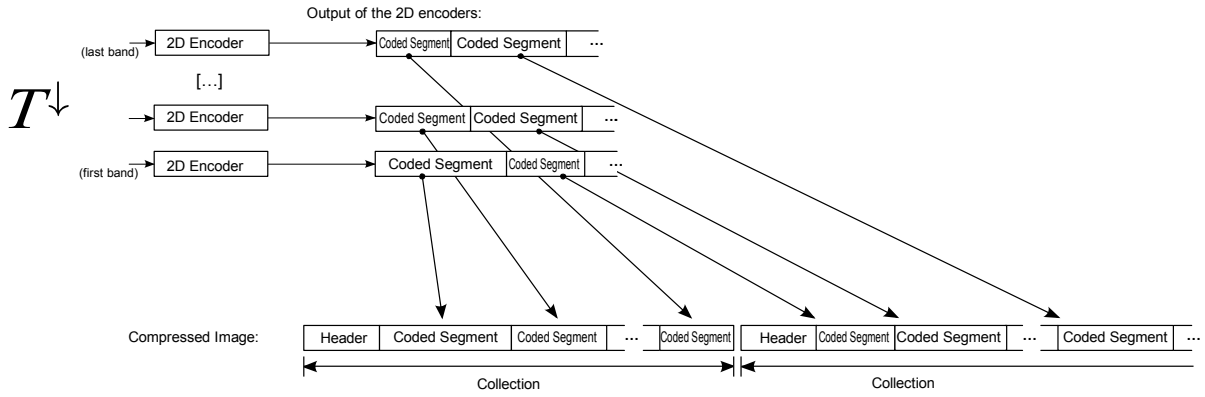

Figure 2-9: Example of the Composition of the Compressed Image

Each collection includes a variable-length header with information required to decompress it, including the relevant side information (if any) produced by the spectral transform stage. The variable-length header encodes mandatory metadata and may also include setup metadata. Both mandatory metadata and setup metadata are necessary to decompress the image, but the latter might be constant for an entire mission and not necessarily produced on board, and thus is optional.

Collections are self-contained pieces of compressed image data in the sense that decompression of the data encoded by a collection does not depend on data encoded in other collections. This provides resilience to missing or corrupted data, limiting the impact of data loss to only the affected collection(s) when an appropriate transport protocol is used (see subsection 2.4 of reference [1]). In the event of a collection not being available after the 2D decoder, decompression can continue by replacing the missing data with arbitrary values. In

this case, since collection boundaries are defined in the Discrete Wavelet Transform (DWT) domain (see subsection 5.1 of reference [1]), the spatial region associated with the affected pixels does not have a sharp boundary. An example of an image reconstructed without one of its collections is shown in figure 2-10.

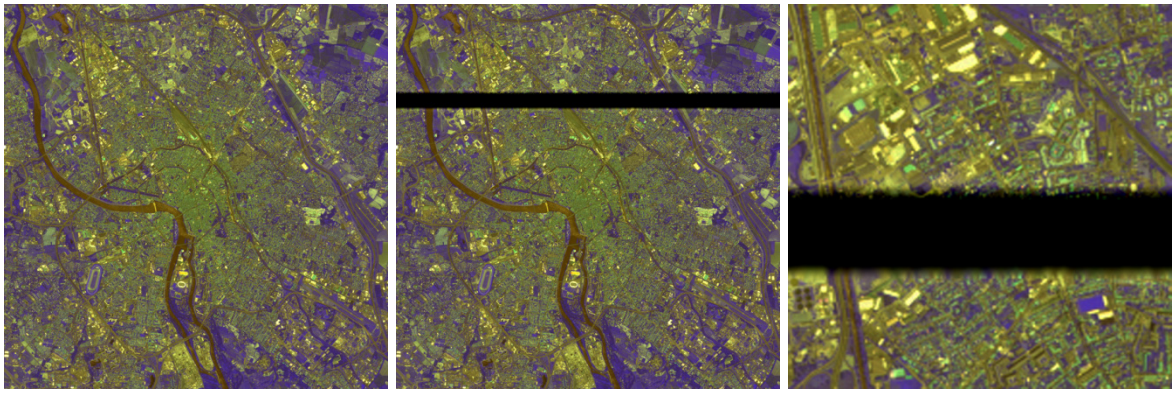

**(left) Original image; (center) Reconstructed using the IWT for compression;
(right) Detail of reconstructed image using IWT for compression.**

Figure 2-10: False Color Reconstruction of Toulouse Image without One of Its Collections

Under lossy compression, the settings of each 2D encoder must be adjusted to achieve an approximate target reconstructed image quality (using the *DCStop*, *BitPlaneStop*, and *StageStop* parameters defined in reference [2]) or to achieve a data rate target (using the *SegByteLimit* parameter). Using the same settings for each 2D encoder does not generally give good performance. In particular, using the same compressed data rate for each transformed band to achieve an overall data rate objective yields poor reconstructed image quality. Instead, a rate-allocation method should be used to assign a variable data rate for each 2D encoder. Section 4 introduces three such rate-allocation methods.

3 COMPRESSION SETTINGS

3.1 INTRODUCTION

The selection of compression parameters affects compression effectiveness and implementation complexity. This section examines the influence of different compression settings on coding performance.

Experimental results are included to illustrate some of the tradeoffs involved. These experiments use the corpus of multispectral and hyperspectral images described in annex A.

To avoid a combinatorial explosion in the number of tests performed, in each experiment one or more interrelated settings are varied while the remaining settings are fixed to the values indicated in annex H. Unless specifically mentioned in the test cases, the default spectral transform stage for both lossless and lossy compression experiments is the IWT. Moreover, the downshift parameter *D* and upshift parameter *U* are both set to 0 for lossless compression. On the other hand, for lossy compression, *U* is set to 2 and D is set to 0.

Lossless compression performance is hereinafter measured by the *compressed data rate* in bits/sample, defined as the size of the compressed image in bits divided by the number of samples in the image. Lower compressed data rates indicate better compression performance.

Lossy compression is evaluated both in terms of compressed data rate and also the fidelity of the reconstructed image, which may be measured using the (energy) *signal-to-noise ratio* (SNR) defined in reference [9],

$$\mathrm{SNR}\left(I,\hat{I}\right) = 10\ \log_{10}\frac{\sum_{x,y,z}\left(I_{x,y,z}\right)^2}{\sum_{x,y,z}\left(\hat{I}_{x,y,z}-I_{x,y,z}\right)^2},$$

where I and $\hat{I}$ denote the original and reconstructed images, respectively. Higher-fidelity reconstructed images yield higher SNR values.

Unless stated otherwise, distortion results for all test images produced by the same instrument are combined into a single value as follows:

$$\text{Combined SNR}\left(\left\{\left(I^i,\hat{I}^i\right)\right\}_{\mathrm{i}}\right) = 10\log_{10}\frac{\sum_i\sum_{x,y,z}\left(I^{\mathrm{i}}_{x,y,z}\right)^2}{\sum_i\sum_{x,y,z}\left(\hat{I}^{\mathrm{i}}_{x,y,z}-I^{\mathrm{i}}_{x,y,z}\right)^2},$$

where $\left\{\left(I^i,\ \hat{I}^i\right)\right\}_{\mathrm{i}}$ is the set of original and reconstructed image pairs.

Moreover, by default, lossy coding performance results are obtained using the Lagrange rate-allocation algorithm described in 4.2.

3.2 UPSHIFT AND DOWNSHIFT

3.2.1 GENERAL

The upshift and downshift operations are defined in section 3 of reference [1]. These operations are controlled by user-specified integer parameters U and D, which are integers in the range 0, …, 15, such that $U - D \leq 5$.

Subsections 3.2.2 and 3.2.3 discuss the effect of U and D in lossless and lossy compression performance, respectively, using the IWT spectral transform. The effect of U and D when other transforms are employed is not analyzed in this document.

3.2.2 LOSSLESS COMPRESSION

This subsection evaluates the lossless compression performance for different values of U. In each case, $D = 0$.

Figure 3-1 provides plots compressed data rates for each instrument at different values of U. Detailed results are provided in table C-1 in annex subsection C2. Upshifting by U bits introduces U redundant (all zeros) bits in each sample, and the subsequent transform and 2D coding stages are not designed to recognize or exploit this redundancy. Consequently, as expected, for lossless compression, setting $U = 0$ is optimum, and nonzero values of U incur a penalty of approximately U bits per sample.

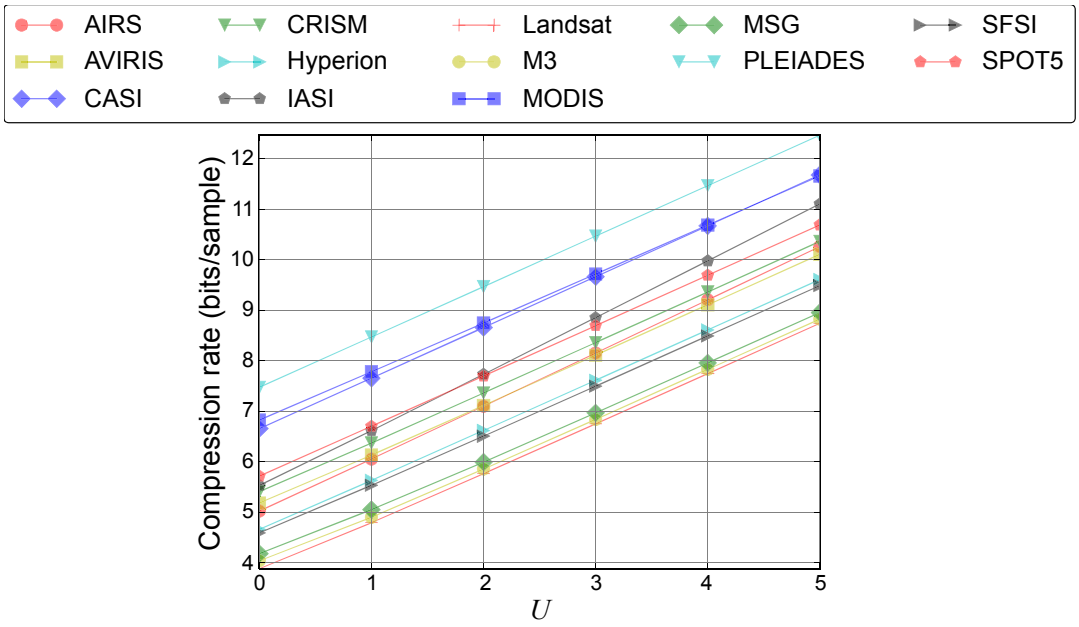

Figure 3-1: Average Lossless Compressed Data Rates for Different Values of *U*

3.2.3 LOSSY COMPRESSION

This subsection analyzes the coding efficiency obtained for different combinations of U and D under lossy compression.

Combined SNR results for the AVIRIS and M3 instruments are provided in figure 3-2. Combined ΔSNR results for some U, D combinations are provided in figure 3-3 using $U = 0$ and $D = 0$ as baseline. ΔSNR is defined for each bit rate r as

$$\Delta\mathrm{SNR}(r) = \mathrm{SNR}_{U,D}(r) - \mathrm{SNR}_{0,0}(r),$$

where $\mathrm{SNR}_{U,D}(r)$ is the combined SNR for U, D. Thus, positive ΔSNR results indicate an image quality improvement over the baseline. Detailed results for each instrument are provided in annex subsection C2.

At bit rates up to 0.5 bits/sample, $(U, D) = (0, 1)$ and $(U, D) = (0, 2)$ improve SNR for most instruments, while other choices of U and D tend to yield worse results than the baseline. At bit rates of 1 bit/sample or higher, all choices satisfying $U - D > 0$ tend to enhance SNR for all instruments except CASI, MODIS, and PLEIADES. For those instruments, $(U, D) = (0, 0)$ is usually preferable. In general, $U - D$ values between 3 and 5 yield the best SNR results, with significant gains at higher rates.

It can also be observed that adding the same constant value to both U and D tends to yield similar compression performance in terms of SNR. Therefore, there is not much motivation to setting $D > 0$ except at very low rates, other than limiting the bit depth, for example, to reduce energy consumption, save hardware resource utilization, or reuse existing hardware implementations of the 2D encoder.

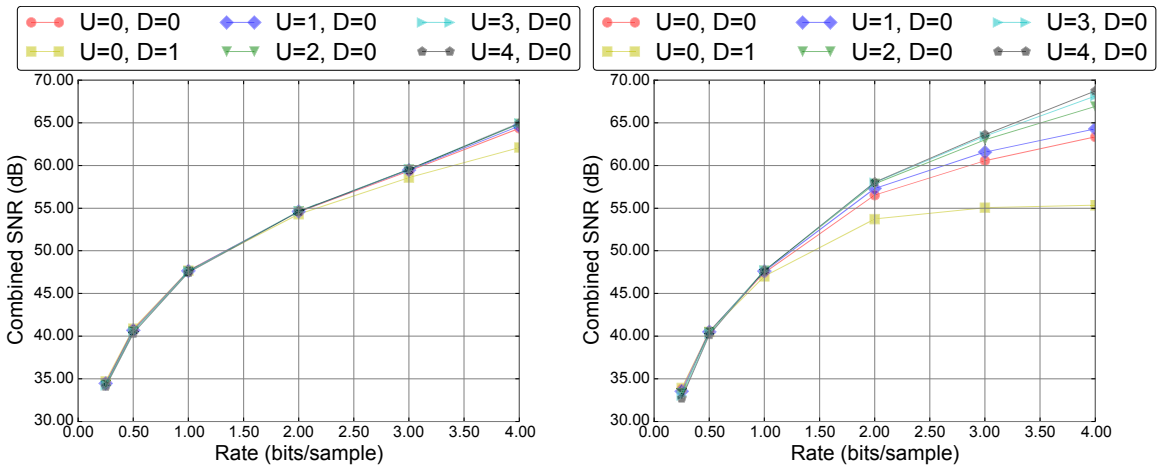

Figure 3-2: Combined SNR Results in dB for the AVIRIS (Left) and M3 (Right) Instruments

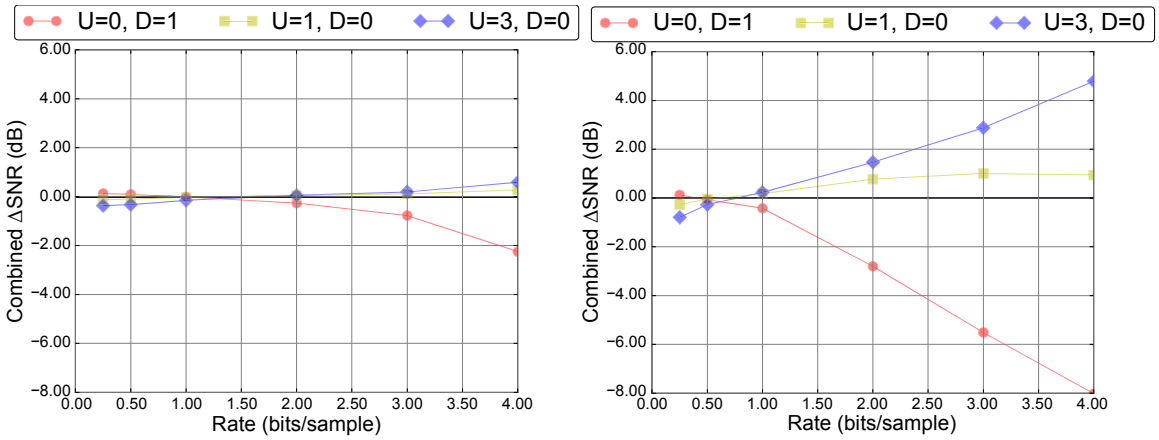

Figure 3-3: Combined ΔSNR Results in dB for the AVIRIS (Left) and M3 (Right) Instruments

3.3 PAIRWISE ORTHOGONAL TRANSFORM

3.3.1 OMEGA SELECTION

The user-specified parameter Ω controls the arithmetic precision of each pairwise operation in the POT. In particular, it affects the calculation of the training parameter C defined in equations (38) and (39) of reference [1], the weights $\widetilde{w}_1$, $\widetilde{w}_2$, and $\widetilde{w}_3$ defined in equations (41) to (49) of reference [1], and the calculation of each lifting stage defined in equation (50) and (55) of reference [1]. A larger value of Ω generally tends to improve image fidelity and reduce the bit rate of the 2D encoders at the cost of larger side information requirements. The reduced bit rate is thus offset by the increased overhead required for the additional side information.

For lossless compression, table 3-1 provides average data rates for all possible values of Ω. Results indicate that this parameter has a small effect on the lossless compression performance. For any instrument, the difference between the best and worst value of Ω is of 0.03 bits/sample or less.

Table 3-1: Combined Lossless Rates in Bits/Sample for All Values of Ω

	Ω = 9	Ω = 10	Ω = 11	Ω = 12	Ω = 13	Ω = 14	Ω = 15	Ω = 16
AIRS	4.744	4.723	4.714	4.711	**4.710**	4.710	4.711	4.711
AVIRIS	4.905	4.886	4.877	4.875	4.874	4.874	4.874	**4.873**
CASI	6.716	6.698	6.691	6.688	6.687	6.687	**6.687**	6.687
CRISM	5.339	5.337	5.336	5.335	**5.335**	5.335	5.335	5.336
Hyperion	4.559	4.557	**4.556**	4.558	4.558	4.557	4.557	4.557
IASI	5.687	5.680	5.678	**5.677**	5.678	5.678	5.679	5.679
Landsat	**3.901**	3.904	3.905	3.905	3.906	3.906	3.906	3.905
M3	4.035	4.033	4.031	4.031	**4.030**	4.030	4.031	4.031
MODIS	**6.472**	6.475	6.479	6.480	6.481	6.481	6.481	6.481
MSG	4.069	4.069	**4.067**	4.068	4.069	4.071	4.070	4.070
PLEIADES	**7.610**	7.610	7.611	7.611	7.612	7.612	7.612	7.612
SFSI	**4.461**	4.462	4.463	4.462	4.463	4.462	4.463	4.462
SPOT5	5.333	5.331	5.331	5.331	**5.331**	5.331	5.332	5.332
VEGETATION	5.428	5.428	5.427	5.427	5.427	5.427	**5.427**	5.427

For lossy compression, figure 3-4 shows combined SNR results for several values of Ω and the AVIRIS and M3 instruments. Similar results are obtained for other instruments. Detailed results for all instruments and values of Ω are provided in annex subsection C3.

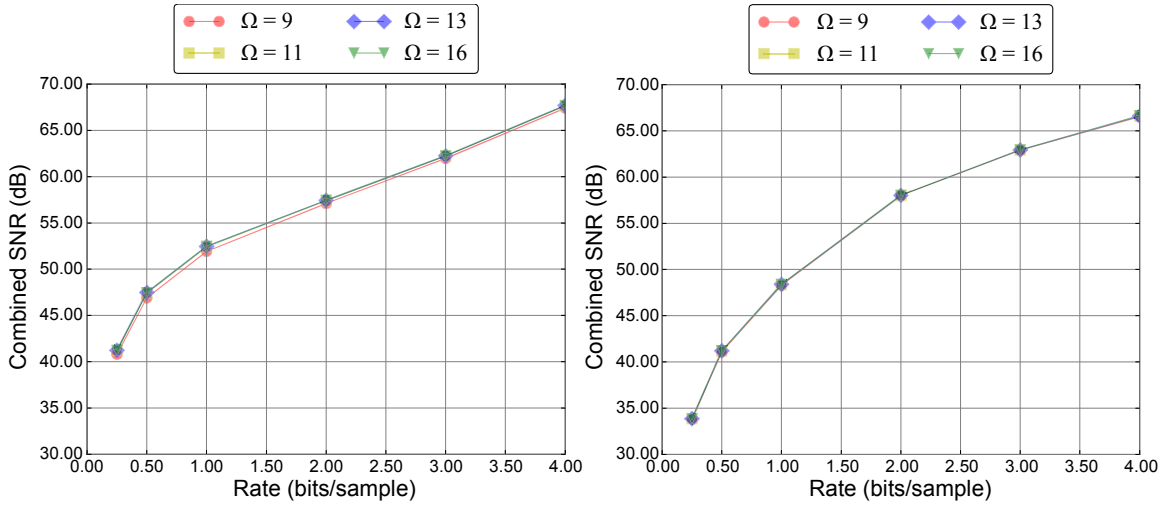

Figure 3-4: Combined SNR Results in dB for the AVIRIS (Left) and M3 (Right) Instruments

In terms of SNR, it can be observed that all values of Ω yield very similar results for most instruments, often with differences smaller than 0.01 dB. Larger differences are observed for AIRS, AVIRIS, and CASI. For these instruments, Ω = 9 generally produces the worst results, while all other values of Ω yield comparable results.

3.3.2 DISCONTINUITY ARTIFACT MITIGATION

The POT calculation for a region depends on the values of several parameters: nominal mean m (one value for each spectral band) and one pair of training parameters (B, C) for each pairwise operation. These sets of parameters generally vary from one region to the next, and this variation can cause discontinuities (sometimes dramatic) at the boundary between two regions of the transformed upshifted image. When compression is lossy, such a discontinuity can lead to an artifact in the reconstructed image, which is referred to as a discontinuity artifact.

Such artifacts are not unique to the POT; they may arise from any transform that is computed separately for each region (references [10] and [11]). Consequently, much of the discussion here also applies to the AAT.

In some cases, variation in POT parameter values can be sufficient to cause a region of the transformed upshifted image to appear to be inverted compared to its neighbor, as illustrated in figure 3-5 (left). Specifically, this occurs when, for a given pairwise operation in the multilevel structure, parameters B or C (but not both) change sign from one region to the next.

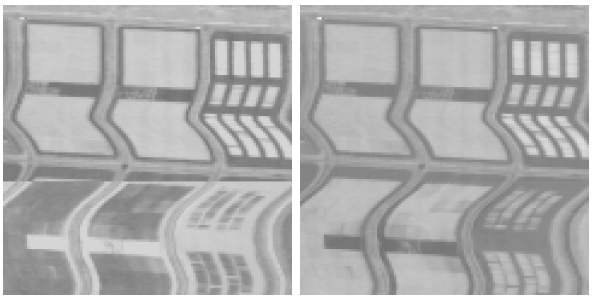

Figure 3-5: Detail of Transformed Band 40 from Image t0477f06-nuc after the POT, with F=32 (Left) Using Bypass Mode and (Right) Using Stable Mode

The use of stable mode prevents these most severe cases. The resulting benefit is illustrated in figure 3-6 (right). Since the type of discontinuity eliminated by stable mode is generally very rare, the use of stable mode tends to have negligible impact on lossless coding performance or lossy rate-distortion curves.

Even when stable mode is used, variations in POT parameters from one region to the next can cause discontinuities in the transformed upshifted image. This can produce discontinuity artifacts in the reconstructed image, particularly at low bitrates, because of the encoder's use of the 2D DWT. This type of artifact can be observed in figure 3-6 (left).

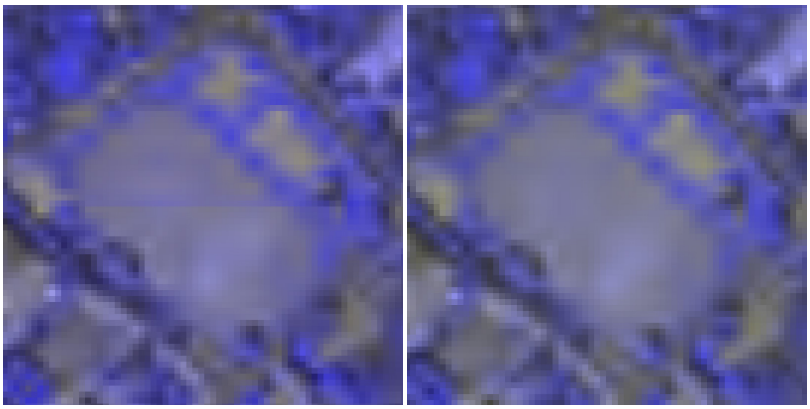

Figure 3-6: Detail of False Color Reconstruction Using the Three First Components (Enhanced Contrast) of montpellier-t at 0.25 bits/sample after the POT, with $F = 32$ (Left) with POT Parameters Calculated Separately for Each Image Region, and (Right) Using the Same POT Parameters for Each Region

When setting POT parameter values, the Recommended Standard allows users to deviate from the calculations specified in 4.5.3.3.2 and 4.5.4.3 of reference [1]. This flexibility permits users to adjust parameter values to reduce discontinuity artifacts.

In particular, by using the same set of POT parameters for all image regions, discontinuities in the upshifted transformed image are eliminated. As can be observed in figure 3-6 (right), this can dramatically reduce the associated artifacts in the reconstructed image. It should be noted that in this situation, the use of stable mode has no effect because, for any given pairwise operation in the multilevel structure, the values of parameters B and C do not change from one region to the next.

When using a fixed set of POT parameters for an image, one would expect the best performance when those parameters are computed from the entire image, but onboard storage constraints might make this impossible. As a more practical alternative, one could compute parameters using the first image region and apply this set of parameter values for all subsequent regions.

Altering the calculation of POT parameters to reduce discontinuity artifacts necessarily sacrifices the ability to use the optimum POT parameters for each region. On the other hand, the reduction of spatial discontinuities in the transformed domain can increase the compression performance of the 2D encoders. The combined effect may help or hurt rate-distortion performance.

SNR results for the AVIRIS and M3 instruments, combined as described in 3.1, are shown in figure 3-7. Detailed results for all instruments in the corpus are provided in tables C-15 and C-16 of annex subsection C3. All results were produced with the POT using the default parameters described in annex H. When a fixed set of POT parameters is used, these parameters are computed using the whole image. In all cases, the region size was set to $F = 32$.

For most images in the corpus, performance is improved by using fixed POT parameters. Only for images collected from MODIS, MSG, and VEGETATION instruments is performance generally worse. The absolute SNR differences are below 0.5 dB in most cases. For the AVIRIS and CASI instruments at 0.25 bits/sample, SNR improvements of 1.25 dB and 0.58 dB are observed, respectively, using fixed parameters. For the AIRS instrument, combined SNR increments between 0.61 dB and 2.41 dB are produced.

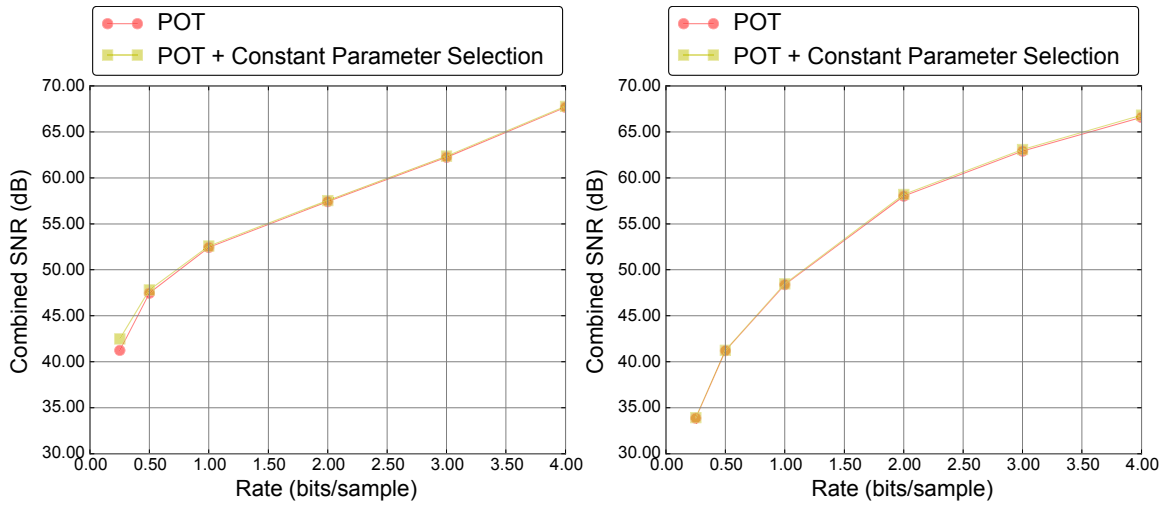

Figure 3-7: Combined SNR Results for the POT, with Fixed and Varying Parameters from Region-to-Region (Left) for the AVIRIS Instrument and (Right) for the M3 Instrument

Finally, even when no discontinuity is present in the transformed upshifted image, the boundary between regions may be evident in the reconstructed image when adjacent regions are reconstructed at different quality levels. The only way to completely avoid this phenomenon is to use the same quality settings for all segments of each 2D image.

3.4 ARBITRARY AFFINE TRANSFORM

3.4.1 GENERAL

The AAT applies an approximation of a user-defined affine transformation, for example, the KLT, to each region of the image. This subsection discusses the main parameter tradeoffs and design decisions for implementing the AAT.

3.4.2 SELECTION OF Ψ

As specified in the Recommended Standard, lossless compression is generally not possible when the AAT is used. Hence, the selection of compression parameters in a lossy compression context is evaluated.

When the downshifted transformed image is intended to be an approximation of $Q' \cdot (X + V')$ for an arbitrary real-valued matrix Q' and a real-valued vector V', the values of Ψ, U, and D should be taken into account in defining Q and V. Specifically, the following equations should be used to derive Q and V from Q' and V':

$$Q = [2^{D-U+\Psi} \cdot Q']$$
$$V = [2^{U} \cdot V'].$$

This guarantees that the $\frac{q_{j,i}}{2^{\Psi}}$ term of equation (87) of reference [1] is within $2^{-\Psi-1}$ of the intended real value $2^{D-U} q'_{j,i}$. Since the upshift stage scales the samples of I by 2^{U}, each term of the sum in equation (86) of reference [1] is $2^{D} \cdot q'_{j,i} \cdot (x_i + v_i) \pm \varepsilon_i$, where $0 \leq \varepsilon \leq 2^{-\Psi+U-1} \cdot (x_i + v_i) + 0.5$. In the downshift stage, samples of T are scaled by 2^{-D}. Therefore, samples of $T^{\downarrow}$ are within $\left(\sum_{i=0}^{N_Z-1} 2^{-\Psi+U-D-1} \cdot (x_i + v_i) \right) + 2^{-D-1}$ of $Q' \cdot (X + V')$. When using these equations, users should carefully choose the values of Ψ, U, and D so that all elements in Q and V fall in the range $-2^{31}, \ldots, 2^{31} - 1$ as stated in subsection 4.6.3.2.3 of reference [1].

In this document, as an illustrative example of the AAT, the real-valued matrix Q' and vector V' approximated by the AAT are those of the KLT, calculated individually for each image as described in reference [12].

Combined SNR for the AVIRIS and M3 instruments for values of a wide range of values of Ψ are shown in figure 3-8. Detailed results for all instruments are provided in tables C-17 to C-22 in annex subsection C4. All results are obtained setting all other parameters to the default values provided in annex H.

As can be observed, too-small values of Ψ result in poor compression performance. This can be explained by the lack of arithmetic precision. Larger values of Ψ, $\Psi \leq 15$, produce better compression performance, although comparable results are obtained for $10 \leq \Psi \leq 15$. Increasing Ψ to values larger than 15 does not improve compression performance. Instead, for sufficiently large values of Ψ, small performance losses are observed, because of the additional side information requirements, and produce significantly worse results than $\Psi \leq 15$,

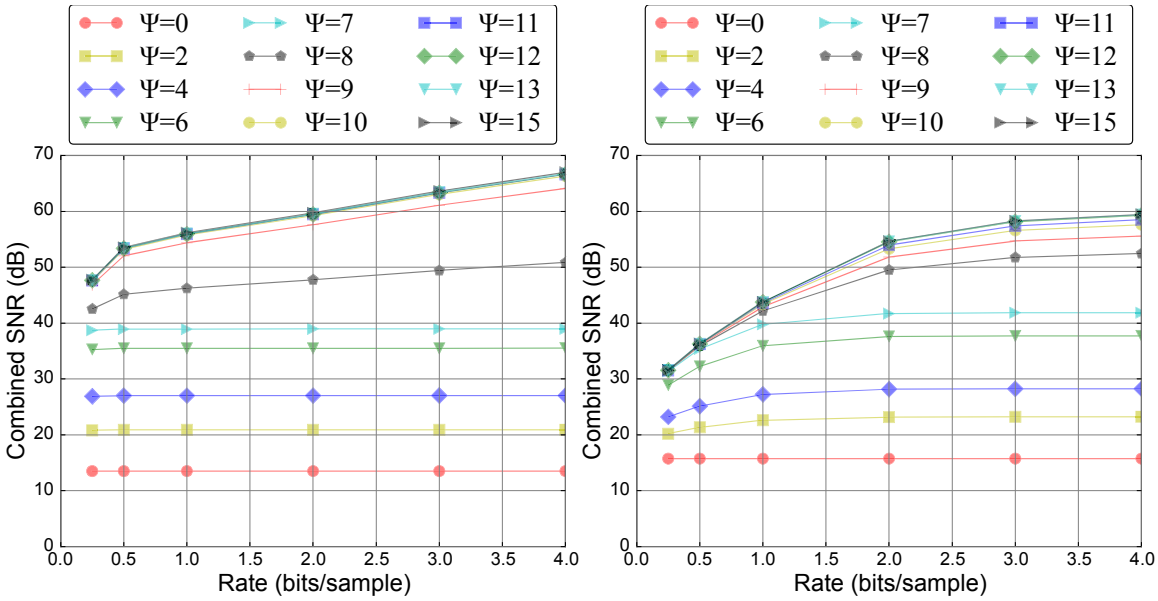

Figure 3-8: Combined SNR Results for the AAT Using Different Values of Ψ (Left) for the AVIRIS Instrument and (Right) for the M3 Instrument

3.4.3 WEIGHT DERIVATION

Rate-allocation strategies, described in section 4, should be used to maximize lossy compression performance. As discussed in annex subsection F1 of reference [1], for a given transform, effective rate allocation makes use of band-dependent constant multiplicative weight factors to account for the approximate relative contribution of MSE distortion in each transformed band to overall MSE distortion in the reconstructed image.

For affine transforms, as described subsection 8.3.5 of reference [13], these weights are the squared norm (energy) of the synthesis basis vectors for each band. Since the AAT applies an integer approximation of an affine transform, the following equation can be used to derive the weight factors $weight_b$ for each band b:

$$weight_b = \sum_{i=0}^{N_{TZ}} (Q_{b,i}^{-1})^2 ,$$

where Q^{-1} is the inverse of Q, and $Q_{b,i}^{-1}$ is the element at row b and column i in Q^{-1}.

3.5 ENCODER

3.5.1 SEGMENTS PER COLLECTION

As mentioned in 2.3.4, segmentation of datasets into collections can improve the robustness of compressed image reconstruction in the event of data loss. When constructing the compressed image, collections are sequentially appended to the compressed bit stream. Each collection contains information to decode a rectangular area of each spectral band, as depicted in figure 3-9.

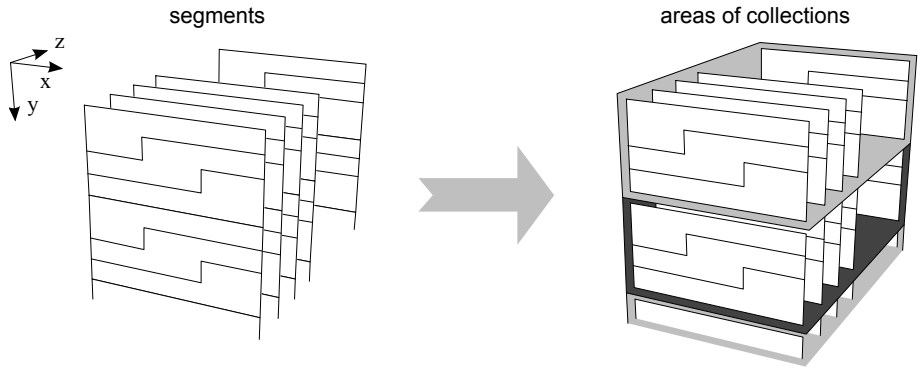

Figure 3-9: Portions of an Image Associated with Segments (Left) and Segments Grouped in Collections (Right) for *R*=3

The number of segments that each band contributes to any collection (except possibly the last one) is controlled via the user-defined parameter R. Given a fixed number of blocks per segment (S'), the value of R determines the height of the rectangular area approximately associated with a collection. This association is only approximate because segments are being defined in the 2D DWT domain instead of in the image domain, as illustrated in figure 5-1 of reference [1]. More specifically, the height associated with a collection is roughly equal to $\frac{8 \cdot S' \cdot R}{\lceil N_X/8 \rceil}$, except possibly for the last collection, which may be smaller. If parameters are selected so that $\frac{8 \cdot S' \cdot R}{\lceil N_X/8 \rceil} \geq N_Y$, then only one collection will be present in the compressed image.

Smaller values of R produce smaller collection heights, which can help reduce buffer memory requirements.

As described in subsection 5.1 of reference [1], collections are self-contained units of information. Therefore, smaller collections can also enhance error resilience in the event of data corruption. It should be noted, however, that error resiliency capabilities require a priori knowledge of collection header locations, which is not directly provided by the Recommended Standard. Instead, this information may be provided by the transport mechanism (see subsection 2.4 of reference [1]).

For a given image, smaller values of R also yield a larger number of collections. Since each header is at least 72 bits long (plus any side information required by the selected spectral

transform), small values of R may result in a significant overhead due to headers, especially when image dimensions N_X and N_Z are small. As an example, illustrating constraints on image size needed to ensure that overhead is very small, ignoring any required side information; to obtain overheads smaller than 0.01 bits/sample, each collection would need to contain more than 7200 samples; that is, $64 \cdot S' \cdot N_Z \cdot R > 7200$.

To evaluate the impact of R on compression performance, here S' is set to $\lceil N_X/8 \rceil$ instead of to the default value of 128. This is done so that any value of R can be selected without violating the parameter constrains described in equation (92) of subsection 5.2.2.3 of reference [1]; that is, $(S' \cdot R) \bmod \lceil N_X/8 \rceil = 0$. Therefore, the following results are not directly comparable to results obtained using the default values of S' and R.

Lossless compression results are provided in table 3-2 using the IWT spectral transform for several values of R. As can be observed, R has a negligible effect on compressed data rate when the IWT is employed, with a peak difference of less than 0.008 bits/sample for the entire corpus and all tested values of R between 1 and 65535.

Combined SNR results are provided for the AVIRIS and M3 instruments using the IWT transform are provided in tables C-23 to C-28 in annex subsection C5. Results consistently show that R has a negligible effect also on lossy compression performance.

Table 3-2: Combined Lossless Compressed Data Rates in Bits/Sample for Different Values of *R*

	***R* = 1**	***R* = 10**	***R* = 100**	***R* = 1000**	***R* = 65535**
AIRS	5.124	5.124	5.124	5.124	5.124
AVIRIS	5.195	5.195	5.195	5.195	5.195
CASI	6.692	6.692	6.692	6.692	6.692
CRISM	5.452	5.451	5.451	5.451	5.451
Hyperion	4.718	4.718	4.718	4.718	4.718
IASI	5.664	5.664	5.664	5.664	5.664
Landsat	3.882	3.881	3.881	3.881	3.881
M3	4.068	4.068	4.068	4.068	4.068
MODIS	6.818	6.817	6.817	6.817	6.817
MSG	4.169	4.168	4.168	4.168	4.168
PLEIADES	7.543	7.536	7.535	7.535	7.535
SFSI	4.612	4.612	4.612	4.612	4.612
SPOT5	5.719	5.717	5.716	5.716	5.716
VEGETATION	5.430	5.429	5.428	5.428	5.428

3.5.2 REGION SIZES

When the POT and AAT are used, bands are divided across the *Y* axis and the different regions are transformed separately. The parameter *F* controls the height of these regions.

Larger values of *F* create fewer regions, each with larger size, which reduces the volume of side information coded, but increases the required memory buffer.

Table 3-3 provides lossless compressed data results for the POT, expressed as the bitrate difference

$$\Delta\text{rate} = \text{rate}_F - \text{rate}_{32},$$

where rate_F is the average lossless compression rate for the POT and a given value of the region size *F*. As can be observed, $F = 32$ (the largest allowed) produces the best lossless compression results.

Table 3-3: Average Δ*rate* in Bits/Sample (Using *F* = 32 As Baseline) for All Possible Values of *F*

	***F* = 1**	***F* = 2**	***F* = 4**	***F* = 8**	***F* = 16**	***F* = 32**
AIRS	0.611	0.373	0.226	0.133	0.063	0.000
AVIRIS	0.301	0.212	0.139	0.076	0.028	0.000
CASI	0.223	0.140	0.085	0.040	0.014	0.000
CRISM	0.271	0.158	0.092	0.046	0.017	0.000
Hyperion	0.568	0.371	0.234	0.121	0.044	0.000
IASI	0.961	0.590	0.348	0.217	0.090	0.000
Landsat	0.641	0.499	0.353	0.196	0.070	0.000
M3	0.098	0.056	0.034	0.018	0.008	0.000
MODIS	0.095	0.099	0.079	0.042	0.012	0.000
MSG	0.185	0.190	0.173	0.134	0.060	0.000
PLEIADES	0.125	0.071	0.033	0.012	-0.002	0.000
SFSI	0.223	0.117	0.072	0.042	0.014	0.000
SPOT5	0.348	0.269	0.167	0.096	0.026	0.000
VEGETATION	0.117	0.130	0.118	0.084	0.047	0.000

Combined ΔSNR results are provided for the AVIRIS and M3 instruments in figure 3-10. Here, ΔSNR is defined as the difference between the SNR obtained for a given value of *F* at a target bitrate *rate* and the SNR obtained at the same rate for $F = 32$; that is,

$$\Delta\text{SNR}_{rate,F} = \text{SNR}_{rate,F} - \text{SNR}_{rate,32}.$$

Detailed SNR results for all instruments are provided in tables C-29 to C-34 of annex subsection C5. Results indicate that $F = 32$ yields the best SNR results also for lossy compression. At low bit rates, the differences between $F = 32$ and the other allowed values of *F* are significant for most instruments. At high bit rates, these differences tend to be smaller.

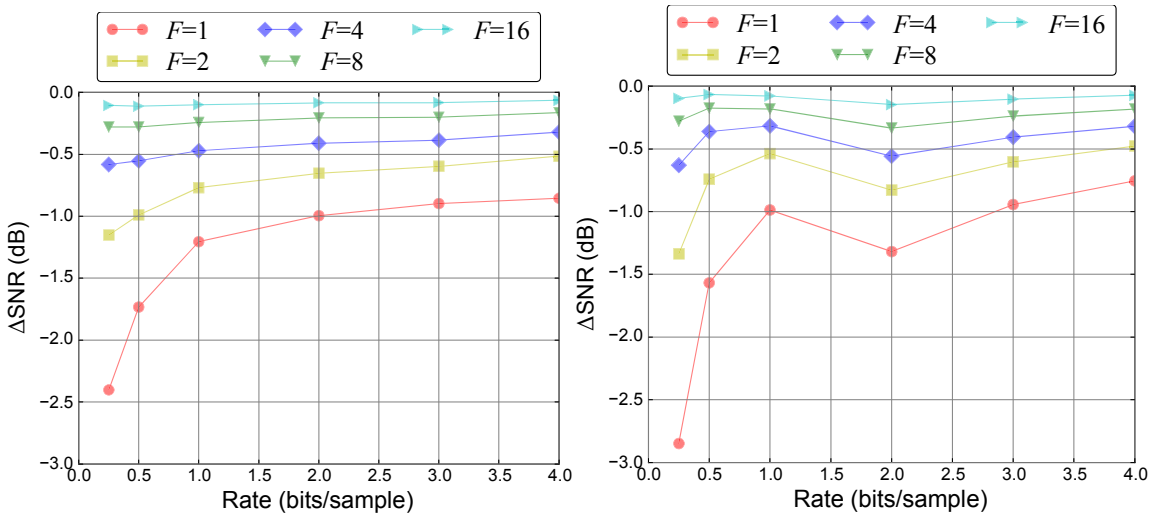

Figure 3-10: Combined ΔSNR Results (Using $F = 32$ As Baseline) for the POT Using Different Values of F (Left) for the AVIRIS Instrument and (Right) for the M3 Instrument

4 RATE ALLOCATION

4.1 INTRODUCTION

The relative impact of MSE distortion in transformed bands on overall reconstructed image distortion is not uniform over all transformed bands. Furthermore, the relationship between the compressed data rate for a transformed spectral band and the reconstructed fidelity of that band varies from band-to-band and image-to-image. For these reasons, compressing each transformed spectral band to the same bit rate usually produces poor rate-distortion results. Performance is improved by instead using a *rate-allocation* strategy that better allots a limited data volume among the transformed bands.

A decompressor can reconstruct an image without knowing what strategy was used by the compressor to perform rate allocation. Consequently, the Recommended Standard does not mandate any particular rate-allocation method, and users are free to choose any desired strategy.

This section describes three *fixed-rate allocation* methods that may be used with the Recommended Standard to produce a compressed image that meets a given rate budget. These methods can be used with any spectral transform including the identity transform. The methods make use of band-dependent weight factors described in annex subsection F1 of the Recommended Standard that account for the approximate relative contribution of MSE distortion in transformed spectral bands to overall MSE in the reconstructed image for the given transform.

Compressed segments produced by the 2D encoders are embedded, that is, they can be truncated at different lengths, with increasing compressed segment length (nearly always) yielding improved reconstruction fidelity for the segment. Rate allocation can thus be performed by truncating compressed segments in a way that meets a user's upper bound on overall compressed data rate for the image.

Truncation points are selected individually for each segment by setting the *SegByteLimit* parameter of the 2D encoder. Compressed segment size can range from including the entire compressed segment to keeping only the header. Headers are mandatory for each segment, and thus impose a lower bound on achievable compressed data rate.

In the context of the Recommended Standard, rate allocation is the problem of assigning truncation points (or, equivalently, compressed size or data rate values) to each compressed segment of the image. The best rate-distortion performance would be obtained by making this assignment based on an analysis of the entire image. But the computational complexity of such a global approach may be prohibitive, especially for a large image. Therefore, in practice, one might perform rate allocation independently over groups of segments, in which each group contains a matching set of segments from each spectral band (e.g., rate allocation might be performed independently over each collection; refer to Fig. 3-9 for a diagram of a collection).

To simplify the discussion, in the remainder of this section, it is assumed that rate allocation is performed over a group of segments containing one segment from each transformed spectral band; generalization to a larger number of segments per band is straightforward. The group of segments is referred to as a Rate Allocation Subset (RAS). It is assumed that some pre-determined rate budget is allocated to each RAS; the most obvious choice would be to distribute an overall image rate budget among a set of RASes proportionally based on the number of DWT blocks in each RAS.

Subsections 4.2, 4.3, and 4.4 discuss the Lagrange, the Reverse-Waterfill, and the Near-Lossless-Rate (NLR) algorithms, respectively, to select truncation points for the segments of an RAS to meet a given rate budget for the RAS. These methods vary in complexity and performance. Hence, their use is application dependent. A performance comparison of these methods is provided in subsection 4.5.

4.2 LAGRANGE POST-COMPRESSION RATE ALLOCATION

The Lagrange optimization algorithm performs a quasi-optimal rate assignment for each segment in a given RAS. This approach generally produces good rate-distortion performance for all valid rates, at the cost of greater implementation complexity than other methods. The algorithm is applied after the 2D encoding stage.

For each segment s, the algorithm begins with some set of candidate truncation points, each one yielding some segment compressed bit rate $T_{s,i}$ and distortion $D_{s,i}$, with truncation points arranged in order of increasing rate and the first point $i = 0$ corresponding to the inclusion of only segment headers. It is assumed that all candidate $(T_{s,i}, D_{s,i})$ pairs lie on the lower convex hull of the set of $(T_{s,i}, D_{s,i})$ pairs; that is, candidates that do not lie on the lower convex hull have been eliminated. Numerous implementations of the lower convex hull algorithm can be found in literature, including Gift Wrapping (reference [14]) and QuickHull (reference [15]). An illustration of points on a lower convex hull is illustrated in figure 4-1.

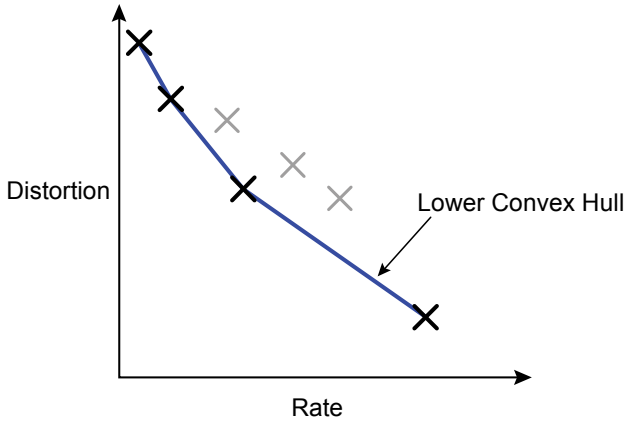

Figure 4-1: Points on the Lower Convex Hull of a Set of Points

Eliminating truncation points that are not in the lower convex hull for each segment reduces the number of candidate truncation points and guarantees that distortion is reduced each time the rate allocated to a compressed segment is increased.

The algorithm obtains quasi-optimal truncation points for each segment, as follows:

a) Each segment is initially assigned truncation point index $t_s = 0$, that is, the truncation point that includes only the segment headers.

b) Segment truncation points are iteratively refined. At each iteration, the truncation point index of a single segment is incremented by one. To select which segment is updated, for each segment s having truncation point index satisfying $t_s < N_s - 1$, the following criterion is calculated

$$\frac{-\Delta D}{\Delta R} weight_s = \frac{-(D_{s,t_s+1} - D_{s,t_s})}{T_{s,t_s+1} - T_{s,t_s}} weight_s,$$

where $weight_s$ is the weight assigned to band s, as described in annex subsection F1 of reference [1], and ΔD and ΔR are the increments in distortion and rate corresponding to changing from the current segment truncation point to the next one. The segment that maximizes this metric provides the largest distortion reduction per compressed bit. This segment is selected, and its associated truncation point index is incremented by one.

c) The selection process is repeated until the combined rate of the selected truncation points reaches the rate target for the RAS or the last truncation point of each segment is selected. After the last iteration, every segment is assigned a quasi-optimal rate.

4.3 REVERSE WATERFILLING

The Reverse Waterfilling method makes use of a simple rate-distortion model to distribute a budget of B bytes among the compressed segments of a RAS based on the variance of data within each segment. A detailed justification of this method can be found in reference [13]. The Reverse Waterfilling method is applied to the RAS as follows:

a) The variance $\sigma_k'^2$ of each segment k, $0 \leq k < N_{TZ}$, is calculated. Variances should be computed after the spectral transform, either before or after the spatial wavelet transform. Measuring variances after the spatial wavelet transform may provide better performance, although this may not be suitable for all hardware implementations. Safe one-pass methods for variance calculation can be found in reference [16].

b) Variances are scaled using the appropriate band weight as described in annex subsection F1 of the Recommended Standard. For example, the following formula can be used:

$$\sigma_k^2 = weight_k \cdot \sigma_k'^2,$$

where $weight_k$ is the weight associated with the transformed band to which segment k belongs for the spectral transform used.

c) The geometric mean $GM = \sqrt[N_{TZ}]{\prod_{k=0}^{N_{TZ}-1}\sigma_k^2}$ of the aforementioned variances is calculated. It may be necessary to perform such a task in the logarithm domain to avoid overflow, and it also may be desirable to use stable accumulation methods if floating point operations are involved. Detail description of these methods can be found in reference [16].

d) An initial data volume allocation for each segment in the spectral segment is found using the following formula:

$$b_k = \begin{cases} \dfrac{B - B_{header}}{N_{TZ}} + \dfrac{1}{8 \cdot 2} \cdot \log_2\left(\dfrac{\sigma_{k,j}^2}{GM_j}\right) & \text{if } \sigma_{k,j}^2 \neq 0, \\ 0 & \text{otherwise} \end{cases}$$

where B_{header} denotes the total header data volume for all segments. Segments with larger scaled variances are assigned larger initial data volumes. It should be noted that the factor $1/8$ is used above to produce results in bytes instead of bits.

e) The initial data volume allocation is adjusted using

$$b_k = \max\{\min\{b_k, B_{max}\}, B_{min}\},$$

where B_{max} is the maximum compressed data volume allowed by CCSDS 122.0-B-2 (reference [2]) for any segment, that is, $R_{max} = 2^{27}$ bytes, and B_{min} is the minimum data volume that can be assigned to any segment. Sufficiently large values of B_{min} allow the reconstruction of most DC coefficients of the 8×8 blocks at the CCSDS 122.0-B-2 2D encoder, which usually improves rate-distortion performance. Excessively large values of B_{min} typically degrade performance since spectral segments with low variances are assigned more bits than necessary. B_{min} values between 2 and 4 bytes per 2D block have been experimentally determined to provide relatively better results for all bit rates and spectral transforms and the parameters specified in annex H.

f) The data volume allocated to each segment is scaled linearly so that the total data volume is the target $B - B_{header}$. First, the scaling factor γ is calculated as

$$\gamma = \frac{B - B_{header}}{\sum_{k=0}^{N_{TZ}-1} b_k}.$$

and the target data volume of each segment is scaled as follows:

$$b_k = \lfloor \gamma b_k \rfloor.$$

g) S_{under} is the set of segments assigned a data volume lower than B_{min} after step f), and S_{over} is the set of remaining segments. The data volume assigned to all segments in

S_{under} is increased to B_{min}, and ε is defined as the increment in data volume due to this operation. Then the data volume of each segment in S_{over} is updated as

$$b_k = b_k - \left\lceil \frac{\varepsilon}{|S_{over}|} \right\rceil,$$

where $|S_{over}|$ is the cardinality of S_{over}.

4.4 NEAR-LOSSLESS RATE

The Near-Lossless Rate (NLR) rate-distortion model (reference [17]), expresses the rate required to code an image region as a function of the NLR for that region and the MSE of the reconstructed image region.

In this context, the term 'near-lossless' is not related to bounded absolute error compression but to the use of the float DWT in the original work (which does not provide lossless reconstruction). This term was originally used in reference [17] and is kept here for consistency.

The NLR method yields similar coding performance as the Lagrange method at medium and high rates, at a lower computational cost.

To apply the NLR rate-allocation method, the following algorithm is applied to a RAS with an allocated compressed size B, expressed in bytes.

a) The near-lossless compressed rate for each segment s, NLR_s, is computed as the size of the CCSDS 122.0-B-2 (reference [2]) compressed segment when all bit planes are coded, also expressed in bytes.

b) The compressed rate assigned to each segment s is defined as

$$B_s(C') = \begin{cases} 0 & \text{if } C' < 8 \cdot P_s \cdot \eta_s \\ C' - 8 \cdot P_s \cdot \eta_s & \text{if } 8 \cdot P_s \cdot \eta_s \le C' \le \mathrm{NLR}_s + 8 \cdot P_s \cdot \eta_s \\ NLR_s & \text{otherwise} \end{cases}$$

for some C' (defined in step c), where $\eta_s = \log_2\left(\sqrt{weight_s}\right)/P_s$ is a bit rate value (in bits per pixel), and $weight_s$ is the weight of the component to which segment s belongs for the spectral transform used. P_s is the number of pixels in segments, and the factor $8 \cdot P_s$ is used to transform bitrates into sizes in bytes.

c) Parameter C' is selected as follows so that $\Sigma_s B_s(C') = B$ Since $B_s(C')$ is a monotonically nondecreasing function of C' for each segment s, $\Sigma_s B_s(C')$ is also monotonically nondecreasing. The user can therefore calculate C' by

1) initializing i=0 and $C'_0 = \max_s\{NLR_s + 8 \cdot P_s \cdot \eta_s\}$;

2) setting $\Delta_i = \frac{B - \sum_s B_s(C_i')}{N_{seg}}$, where N_{seg} is the number of segments in the RAS;

3) setting $C'_{i+1} = \lfloor C'_i + \Delta_i \rfloor$; and,

4) if $\Delta_i > 0$, setting $i = i + 1$ and going to (b), or otherwise, ending the algorithm.

The algorithm is then guaranteed to provide at most a difference of $N_{seg} - 1$ bytes between B and $\Sigma_s B_s(C')$. This difference can be addressed by adding one extra byte to B_s for s = 0 … $B_{RAS} - \Sigma_s B_s(C') - 1$.

4.5 PERFORMANCE COMPARISON

The performance of the rate allocation algorithms is data dependent. Therefore, the results present in this section are for reference only. With that said, combined SNR values for all test images using the IWT as spectral transform are provided in figure 4-2. Default parameters described in annex H are used for the Lagrange and Reverse-Waterfill rate-allocation methods. Variance calculation for the latter is calculated after the spectral and spatial transforms. In addition, the NLR algorithm requires a larger number of blocks per segment to produce competitive compression performance. Hence, results for NLR are obtained for $S = 2 \cdot N_X$.

In general, the Lagrange algorithm is a more effective rate-allocation method compared to Reverse-Waterfill and the NLR methods, as evident from its higher SNR performance over all bit rates tested. Moreover, the Reverse-Waterfill yields better results than NLR methods by 3 to 5 dB SNR under rate of 2 bits/sample. On the contrary, at rates higher than 2 bits/sample, the performance of the NLR is better than the Reverse-Waterfill and yields almost identical performance to the Lagrange algorithm at 4 bits/sample. Additional results can be found in reference [17].

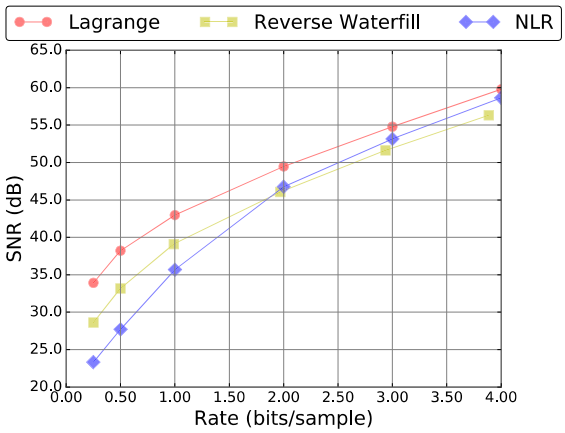

Figure 4-2: Combined SNR Results for All Images Using Different Rate-Allocation Algorithms

5 IMPLEMENTATION ISSUES

5.1 INTRODUCTION

This section discusses some practical issues that may be of interest to implementers of the Recommended Standard. Subsection 5.2 provides dynamic range expansion values as a function of the input bit depth. Subsection 5.3 briefly examines a method for efficient implementation of the IWT filters. Subsection 5.4 discusses some aspects of the POT and how they can be efficiently implemented. Subsection 5.5 references several training methods to approximate the KLT using the AAT with reduced computational complexity. Subsection 5.6 discusses the impact of streaking artifacts on compression performance.

5.2 DYNAMIC RANGE EXPANSION

5.2.1 GENERAL

The dynamic range output by a spectral transform is usually larger than the dynamic range of input data. Such an increase is referred to as *dynamic range expansion*. When implementing the Recommended Standard, it is necessary to consider the maximum image bit depths of the downshifted transformed image $T^{\downarrow}$ and the maximum bit depths that can be accepted by the CCSDS 122.0-B-2 (reference [2]) 2D encoders. Partial or total data loss can occur if the bit depth of $T^{\downarrow}$ exceeds the maximum accepted by the 2D encoders. This subsection discusses the maximum bit depth of T as a function of the bit depth of the upshifted image $I^{\uparrow}$, denoted $n + U$, where n is the bit depth of the input image, and U is the upshift parameter.

When the identity transform is used, the input data are not modified, that is, $T = I^{\uparrow}$. Therefore, the output bit depth is $n + U$.

All other spectral transforms defined in the Recommended Standard may increase the output bit depth. Subsections 5.2.2 to 5.2.4 discuss the maximum dynamic range expansion of the IWT, the POT, and the AAT.

5.2.2 INTEGER WAVELET TRANSFORM

As described in subsection 4.4 of the Recommended Standard, the IWT consists of five levels. The first level is applied to the original data samples X to produce the L_1 and H_1 subbands. The second is applied to L_1 to produce the L_2 and H_2 subbands, and this process is iterated until the L_5 and H_5 subbands are obtained. Each level of the IWT can increase the dynamic range of the transformed data. To calculate the maximum possible dynamic range of T, it suffices to analyze the extrema of L_5 and H_5. Lower and upper bounds for L_5 and H_5 can be calculated as described in this subsection.

Using equations (22) and (23) of reference [1], it is possible to obtain an analytical expression of any element of L_1 and H_1 as a function of the input samples in X. For example, for an input sequence of length 256, elements l_{64} from L_1 and h_{64} from H_1 can be obtained as

$$l_{64} = x_{128} + \left\lfloor \frac{x_{127}}{4} - \frac{1}{4}\left\lfloor \frac{x_{126} + x_{128}}{2} \right\rfloor + \frac{x_{129}}{4} - \frac{1}{4}\left\lfloor \frac{x_{128} + x_{130}}{2} \right\rfloor + \frac{1}{2} \right\rfloor$$

$$h_{64} = x_{129} - \left\lfloor \frac{x_{128} + x_{130}}{2} \right\rfloor .$$

Using the same equation, it is possible to express any element of L_2 and H_2 as a function of the elements of L_1, which in turn can be expressed as a function of the input sequence X. This procedure can be iterated to obtain analytical expressions for any element in L_5 and H_5 as a function of the input samples $x_0, \ldots, x_{N-1}$.

If l and h are the elements in the central position of L_5 and H_5, respectively, after a five-level IWT is applied to a sequence of sufficient length (e.g., $N = 256$), without loss of generality, the extrema of L_5 and H_5 can be obtained by analyzing the extrema of l and h. The analytical expressions of l and h contain a number of floor operations. The ith floor operation of each expression can be substituted as follows:

$$\lfloor a \rfloor = a - \varepsilon_i,$$

Where a is the expression on which the floor operation is applied and ε_i is the rounding error, which verifies $0 \leq \varepsilon_i < 1$.

After this substitution, l and h can be expressed as a linear combination of the input sequence and the rounding errors; that is,

$$l = \sum_i \alpha_i x_i + \sum_i \beta_i \varepsilon_i$$

$$h = \sum_i \gamma_i x_i + \sum_i \delta_i \varepsilon_i ,$$

where $\alpha_i, \beta_i, \gamma_i$, and δ_i are real rational numbers, possibly 0. Since these expressions are linear, lower bounds for l and h can be produced by

a) setting x_i to the minimum sample value allowed in $I^{\uparrow}$ for all i;

b) setting ε_i to 0 for all i;

c) evaluating the corresponding expression and store it in *current_evaluation*;

d) for each input sample x_i:

1) setting x_i to the maximum sample value allowed in $I^{\uparrow}$;

2) evaluating the expression and save it to *candidate_evaluation*;

3) if *current_evaluation* < *candidate_evaluation*, setting *current_evaluation* to *candidate_evaluation*; else setting x_i to the minimum sample value allowed in $I^{\uparrow}$;

e) for each error ε_i:

1) setting ε_i to 1;

2) evaluating the expression and saving it to *candidate_evaluation*;

3) if *current_evaluation* < *candidate_evaluation*, setting *current_evaluation* to *candidate_evaluation*; else setting ε_i to 0;

f) the lower bound is the value stored in *current_evaluation*.

Upper bounds for *l* and *h* can be obtained using the same method, but substituting < for > in steps d)3) and e)3).

Table 5-1 provides the maximum bit depth required to store any sample of *T* as a function of the input bit depth of $I^{\uparrow}$. As can be observed, the maximum dynamic range expansion is 2 bits. This expansion is not dependent on the input data being signed or unsigned. Input sequences for which the IWT introduces a 2-bit dynamic range expansion were found for bit depths of 8 bits or higher.

Table 5-1: Bit Depth of *T* in Bits As a Function of the Bit Depth of $I^{\uparrow}$, $n + U$, for the IWT

$n + U$	Maximum Output Bit Depth
8 or lower	10
9	11
10	12
11	13
12	14
13	15
14	16
15	17
16	18
17	19
18	20
19	21
20	22
21	23
22	24
23	25
24	26

5.2.3 PAIRWISE ORTHOGONAL TRANSFORM

As described in subsection 4.5 of the Recommended Standard, the POT first applies a mean subtraction operation and then a multi-level structure of pairwise operations. The dynamic range expansion of the POT is the result of adding the dynamic range expansion of each of these two stages.

The expansion due to the nominal mean subtraction is 1 bit.

The dynamic range expansion due to the application of the pairwise orthogonal operations depends on the number of spectral bands, N_Z, which determines the number of levels of the multi-level structure, $\lceil \log_2 N_Z \rceil$, and the type of pairwise operations involved in each level, for example, balanced or unbalanced.

Brute-force approaches can be used to compute the minimum and maximum output values of the POT for any number of levels and pairwise operation types. Table 5-2 provides the worst case (highest) bit depth of T as a function of Ω. These bounds are valid for images such that $1 \leq N_Z \leq 2^{16}$, as a function of Ω. The dynamic range expansion due to the mean subtraction stage is included in these values. It should be noted that the dynamic range expansion of the POT does not depend on I being signed or unsigned.

Table 5-2: Maximum Dynamic Range of T in Bits As a Function of the Bit Depth of $\overset{\uparrow}{I}$, $n + U$, and Ω

$n + U$	Ω=9	Ω=10	Ω=11	Ω=12	Ω=13	Ω=14	Ω=15	Ω=16
1	4	4	4	4	4	4	4	4
2	5	7	7	7	7	7	7	7
3	7	7	7	7	7	7	7	7
4	7	8	8	8	8	8	8	8
5	8	8	8	8	8	8	8	8
6	9	9	9	9	9	9	9	9
7	10	10	10	10	10	10	10	10
8	11	11	11	11	11	11	11	11
9	12	12	12	12	12	12	12	12
10	13	13	13	13	13	13	13	13
11	14	14	14	14	14	14	14	14
12	15	15	15	15	15	15	15	15
13	16	16	16	16	16	16	16	16
14	17	17	17	17	17	17	17	17
15	18	18	18	18	18	18	18	18
16	19	19	19	19	19	19	19	19

5.2.4 ARBITRARY AFFINE TRANSFORM

The output dynamic range of the AAT depends on the user-defined transform matrix Q and the vector V. Therefore, it is not possible to produce a general table with the dynamic range expansion for the AAT. Notwithstanding, given Q,V and the extrema of $I^{\uparrow}$, it is possible to bound the minimum and maximum output value as follows.

If $\mu_z = \max\left\{\left|\max_{x,y}\{I^{\uparrow}_{x,y,z}\} + v_z\right|, \left|\min_{x,y}\{I^{\uparrow}_{x,y,z}\} + v_z\right|\right\}$, then, substituting $(x_i + v_i)$ for μ_z in equation (87) of reference [1], samples of the zth component of the transformed image are bounded by

$$\min_{x,y}\{T_{x,y,z}\} \geq \alpha_z = \left\lfloor -2^{-\Psi} \cdot \mu_z \cdot \sum_{i=0}^{N_Z-1} \left|q_{z,i}\right| - \frac{N_Z}{2} \right\rfloor$$

$$\max_{x,y}\{T_{x,y,z}\} \leq \beta_z = \left\lceil 2^{-\Psi} \cdot \mu_z \cdot \sum_{i=0}^{N_Z-1} \left|q_{z,i}\right| + \frac{N_Z}{2} \right\rceil .$$

The term N_Z / 2, not present in equation (87) of reference [1], is used to take into account the N_Z rounding operations of that equation, one for each term in the sum, each of which adds or subtracts at most 1/2.

Using these values, it is possible to calculate upper and lower bounds for the transformed image,

$$T_{x,y,z} \in \min_z\{\alpha_z\}, \ldots, \max_z\{\beta_z\},$$

and for the downshifted transformed image,

$$T^{\downarrow}_{x,y,z} \in [2^{-D} \min_z\{\alpha_z\}], \ldots, [2^{-D} \max_z\{\beta_z\}].$$

5.3 IWT FILTER IMPLEMENTATION

The normative definition of the IWT filter operations is given in equations (22) and (23) of the Recommended Standard. In these equations, two and three cases are employed to define the output of the kth high-pass and low-pass elements, respectively. In this subsection, it is explained how to reduce these definitions to a single expression.

The additional cases are employed to avoid invalid memory accesses, that is, indices lower than 0 or larger than the largest element index. Notwithstanding, it is possible to reduce the definitions to a single case as follows.

For the high-pass elements, the special case for $k = q - 1$ and N even can be avoided by mapping accesses from x_N (invalid index) to x_{N-2} (valid index). Then h_k can be defined as follows:

$$h_k = x_{2k+1} - \left\lfloor \frac{1}{2} \cdot (x_{2k} + x_{2k+2}) \right\rfloor.$$

For the low-pass elements, the special cases for $k = 0$, and for $k = p - 1$ and N odd can be avoided by mapping accesses from h_{-1} (invalid index) to h_0 (valid index) and from h_{p-1} (invalid index) to h_{p-2} (valid index). Then l_k can be defined as follows:

$$l_k = x_{2k} + \left\lfloor \frac{1}{4} \cdot (h_{k-1} + h_k + 2) \right\rfloor.$$

5.4 EFFICIENT PAIRWISE ORTHOGONAL TRANSFORM

5.4.1 INTRODUCTION

This subsection discusses some aspects that can be considered to create efficient implementations of the POT. Subsection 5.4.2 deals with the calculation of individual pairwise operations. Subsection 5.4.3 addresses efficient methods for computing the multi-level structure.

5.4.2 EFFICIENT PAIRWISE OPERATIONS

The computation of the principal and detail output of the pairwise operations requires the weights $\widetilde{w_1}$, $\widetilde{w_2}$, and $\widetilde{w_3}$ to be available. These weights are a function of B, C, and Ω, which in turn are a function of the scaled variances $\widetilde{\sigma_x^2}$ and $\widetilde{\sigma_y^2}$ and the covariance $\widetilde{\sigma_{x,y}}$.

Precise variance and covariance calculation may demand a complex hardware implementation. Notwithstanding, the definition of B and C uses the word ‘should’ and not ‘shall’. Therefore, approximate variance and covariance calculation methods can be used to compute B and C. In particular, it may be sufficient to use a single-pass method, which, because of catastrophic cancellation, can yield negative values, followed by an abs(·) function for the purpose of calculating the POT. For example, the following equations may be employed:

$$\widetilde{\sigma_x^2} = \left| \sum_{i=0}^{N-1} x_i^2 - N \cdot \left\lfloor \frac{1}{N} \cdot \sum_{i=0}^{N-1} x_i \right\rfloor^2 \right|;$$

$$\widetilde{\sigma_y^2} = \left| \sum_{i=0}^{N-1} y_i^2 - N \cdot \left\lfloor \frac{1}{N} \cdot \sum_{i=0}^{N-1} y_i \right\rfloor^2 \right|;$$

$$\widetilde{\sigma_{x,y}} = \sum_{i=0}^{N-1} (x_i \cdot y_i) - N \cdot \left\lfloor \frac{1}{N} \cdot \sum_{i=0}^{N-1} x_i \right\rfloor \cdot \left\lfloor \frac{1}{N} \cdot \sum_{i=0}^{N-1} y_i \right\rfloor.$$

Alternative variance and covariance calculation methods, that is, different from equation (36) in reference [1], may produce different values of *B* and *C*. In turn, this can result in different transformed values in *T* and a different compressed bit stream. Notwithstanding, since *B* and *C* are stored as side information, the image decompression process is not affected by the variance calculation method, and lossless reconstruction is not compromised.

The intermediate weights $\widetilde{w_1}$, $\widetilde{w_2}$, and $\widetilde{w_3}$ are a function of *B*, *C*, and Ω. These weights need not be calculated for each pairwise operation. Instead, for any value of Ω, look-up tables can be created for every possible value of *B* and *C*.

The *B* and *C* parameters have a bit depth of 1 and $\Omega - 1$ bits, respectively. Therefore, an implementation supporting one value of Ω would require two tables with 2^{Ω} entries each: one table for balanced operations and the other for unbalanced operations. Weights $\widetilde{w_1}$, $\widetilde{w_2}$, and $\widetilde{w_3}$ have a bit depth of at most $\Omega + 2$ bits, and the amount of information to be stored in each table would be $3 \cdot (\Omega + 2) \cdot 2^{\Omega}$ bits. When look-up tables are used, the intermediate parameters $\tilde{p}$, $\tilde{t}$, and $\widetilde{\sqrt{2}}$ need not be calculated nor stored.

Reference look-up table values are provided in annex B. Possible table implementation methods include synthesis, storage on an external ROM, and generation on startup, for example, using a soft-core processor.

The logic associated with the flip input of each pairwise operation can be implemented with a two-bit flip-flop, where the first bit is the sign bit of *B* and the second bit is the sign bit of *C*.

5.4.3 EFFICIENT MULTI-LEVEL OPERATIONS

A pairwise operation at level *l* of the multi-level structure takes two input sequences and produces a detail output sequence and a principal output sequence. In all levels except the last one, the principal output becomes a level-($l + 1$) input or a level-($l + 2$) input in case it is an unpaired sequence. The detail sequence is not further processed and becomes part of one of the bands of the transformed image *T*. In the last level, both output sequences become part of *T* without further processing.

This property can be exploited to produce an efficient implementation of the multi-level structure of the POT using a stack. Mean-subtracted sequences are sequentially read and pushed to the top of the stack. Each sequence is assigned a structure level, initially 1 or 0, depending on whether the sequence is unpaired or not. Whenever a sequence is pushed, if the two sequences at the top of the stack belong to the same level, a pairwise operation is applied to these sequences. The operation is unbalanced if and only if one of the inputs is unpaired. The detail output sequence is saved, and the principal output sequence is pushed to the top of the stack. The structure level of the principal output sequence is 2 or 1 units higher than the level of the input sequences, depending on whether the principal output sequence is unpaired or not. Principal output sequences pushed to the stack may also trigger further pairwise operations whenever the two sequences at the top of the stack belong to the same level.

By construction, there can be at most one sequence per level in the stack, that is, at most $\lceil \log_2 N_Z \rceil$ sequences. In comparison, a naive algorithm that processes all pairwise operations of one level before proceeding to the next one would require to store $\left\lceil \frac{N_Z}{2} \right\rceil$ sequences.

If $N_Z = 2^k$, then all pairwise operations are balanced.

5.5 ARBITRARY AFFINE TRANSFORM

The KLT provides perfect spectral decorrelation. Even though perfect decorrelation does not necessarily imply best coding efficiency, it is well motivated to define the vector and matrix of the AAT so that a spectral transform that approximates the KLT is applied, as described in 3.4. However, calculating the KLT matrix has a high computational cost. Several training strategies can be used to mitigate this cost and enhance the overall compression performance.

Instead of using all input samples to calculate the KLT matrix, subsampling may be used to accelerate the computation of the covariance matrix, as described in reference [12]. This reduces the computational complexity of obtaining the covariance matrix by a factor equal to the subsampling factor employed.

Fixed approximations to the KLT matrix may be computed offline using appropriate training data, as described in references [10], [11], and [18]. This approach eliminates any onboard training costs, but may reduce compression performance if the training data are not representative of the input image.

When different Q and V parameters are used in adjacent regions, discontinuities across region edges may appear, penalizing compression performance. A method to mitigate these discontinuities based on optimal eigenvector and sign selection is described in reference [19].

5.6 STREAKING ARTIFACTS

Pushbroom hyperspectral imagers employ two-dimensional detector arrays, in which each element corresponds to a spectral band and spatial position. Detector elements generally present varying characteristics, which cause discontinuities in adjacent samples of a given spectral band. These discontinuities, which generally appear as along-track lines, are usually referred to as *streaking artifacts*. Figure 5-1 depicts a false-color image exhibiting this type of discontinuity.

As discussed in subsection 2.3 of reference [1], streaking artifacts tend to have a negative impact on compression performance. This impact can be reduced by applying nonuniformity correction prior to compressing the images. Complete radiometric calibration can remove streaking artifacts, but it may not be practical to perform such calibration on board (reference [20]). On the other hand, a simple reversible pre-processing stage may enhance compression performance.

Subsection 5.5 of reference [21] discusses simple nonuniformity correction algorithms for this purpose.

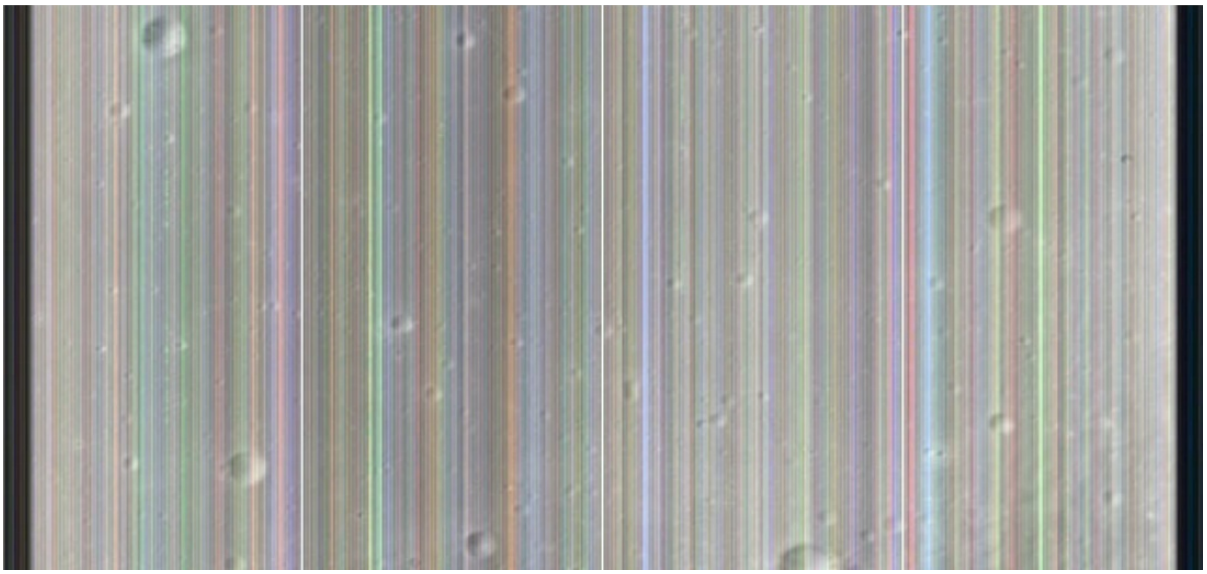

Figure 5-1: False-Color Image Derived from Spectral Channels 200, 201, and 202 from a Portion of an M3 Image

6 PERFORMANCE

6.1 INTRODUCTION

This section presents lossless and lossy compression performance results for the Recommended Standard.

Results included here are intended to provide some indication of the relative compression performance of the Recommended Standard and other standard compression algorithms. In particular, the Recommended Standard is compared to JPEG-LS and JPEG 2000. However, the different compressors vary significantly in complexity, and it is not straightforward to provide a meaningful complexity comparison between different compressors. In particular, some of the compression methods used here, for example, JPEG 2000, may be challenging to implement on board a spacecraft.

Subsection 6.2 describes the compression methods used. Subsection 6.3 provides lossy compression results in terms of SNR. Subsection 6.4 provides lossless compressed data rates in bits/sample.

6.2 EXPERIMENTAL SETUP

6.2.1 OVERVIEW

This subsection lists the compressors used and their parameters and describes some of the methodology employed to obtain the compression results.

The 47 hyperspectral and multispectral images described in annex A are used to evaluate these compressors.

6.2.2 CCSDS 122.1-B-1

For the Recommended Standard, lossy compression results are obtained using the IWT, the POT, and the AAT approximating the KLT in the spectral transform stage and the Float DWT in the 2D encoders. It should be noted that the Integer DWT can also be used to produce lossy compression results, at the cost of lower coding efficiency. Rate allocation is obtained with the Lagrange algorithm using RASs equal to collections.

Lossless compression results are obtained using the IWT and the POT in the spectral transform stage and the Integer DWT in the 2D encoders.

Lossy results for the AAT are obtained by approximating the KLT. The KLT matrix and vector for each image is first calculated as described in reference [12], and then approximated as described in 3.4, for $\Psi = 15$ and $D = 3$.

All other compression settings for lossy and lossless compression use the values indicated in annex H.

6.2.3 JPEG-LS

The JPEG-LS image compression standard (reference [22]) supports compression of three-dimensional ('multi-component') images such as multispectral and hyperspectral images. However, JPEG-LS is primarily designed to exploit two-dimensional image structure and takes almost no advantage of similarities between adjacent spectral bands. Moreover, because of a lack of readily available JPEG-LS software implementations that support multi-component images having more than a few spectral bands, no attempt to exploit spectral band redundancy was made.

The independent application of JPEG-LS to each spectral band would produce poor results for images with small spatial dimensions compared to the number of bands, because of image header overhead and because the compressor would not have much time to adapt to each band. Therefore, spectral bands were concatenated vertically to produce a flat 2D image, as depicted in figure 6-1. To circumvent software limitations on the input image height supported, flattened images with height exceeding $2^{16} - 1$ were split along the vertical axis as needed.

The JPEG-LS standard can only be applied to unsigned samples. Therefore, a constant offset equal to the minimum image sample value was subtracted from each sample of signed images.

Results for JPEG-LS were obtained using Release 1.51 of the libjpeg implementation by Thomas Richter, University of Stuttgart, and Accusoft (https://jpeg.org/jpegls/software.html).

For lossless compression, the following parameters were employed:

`-ls -c`.

For lossy compression, the following parameters were used:

`-ls -c -m` *m*,

where *m* is the desired maximum absolute error. JPEG-LS was evaluated for all $m \in \{1,8,32,64,128,254\}$. These values of *m* were selected to produce compressed bit rates similar to those employed for the other compression algorithms.

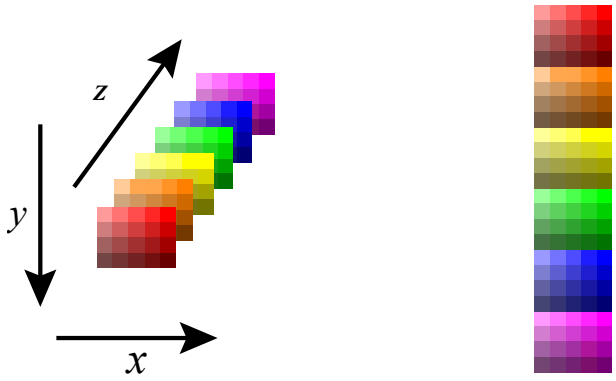

Figure 6-1: Flattening a 3D Image (Left) into a 2D Image (Right)

6.2.4 JPEG2000

The JPEG2000 image compression standard (reference [6]) defines a two-dimensional image coder based on wavelet transforms and bit plane encoding.

Two different compression modes are available: reversible and irreversible. The reversible mode enables lossless compression and employs the 2D LeGall 5/3 IWT. The irreversible mode can only be used for lossy compression, and employs the 2D CDF 9/7 DWT.

Spectral band decorrelation is supported in the Part 2 extension of the JPEG2000 standard (reference [23]). This extension enables the application of one 1D transforms, including wavelet transforms. In particular, the 1D reversible IWT and the 1D irreversible CDF 9/7 DWT are available for lossless and lossy compression, respectively.

Results for JPEG2000 were obtained using version 7.5 of the Kakadu implementation by David Taubman and NewSouth Innovations Ltd (http://kakadusoftware.com).

For lossless compression, 5 levels of 2D IWT and 5 levels of 1D IWT were employed. The portion of the command line that controls lossless settings is

```
Clayers=1 Clevels=5 Creversible=yes Cycc=no -full -no_weights
Mstage_xforms:I1={DWT,1,0,5,0}.
```

For lossy compression, 5 levels of 2D CDF 9/7 DWT and 5 levels of 1D CDF 9/7 DWT were used. It should be noted that mixing the reversible IWT and the irreversible CDF 9/7 DWT as spectral and spatial transformations is not defined in the JPEG 2000 standard. Therefore it is not possible to instruct JPEG 2000 coders to employ the same spatial and spectral transforms as the Recommended Standard. The following portion of the command line controls these settings:

```
Clayers=1 Clevels=5 Creversible=no Cycc=no -full -no_weights
Mstage_xforms:I1={DWT,0,0,5,0}-rate r,
```

where $r = r' \cdot N_Z$ and r' is the desired compressed data rate in bits/sample.

Even though the KLT can be used within the JPEG 2000 standard, the computational complexity and side information requirements of this approach make it impractical for many practical scenarios. Therefore, in this document JPEG 2000 is used only with the spectral IWT.

6.3 LOSSY COMPRESSION

Combined SNR results for each instrument are plotted in figures 6-3–6-9. In all plots, the legend shown in figure 6-2 is employed. Individual image results are provided in table C-35 of annex subsection C6. Results for the Recommended Standard using the AAT are not provided for the AIRS, IASI, or CRISM instruments, because the required side information exceeds most or all of the target bit rates.

As expected, the Recommended Standard generally produces better compression results when any of the spectral transforms are employed. The only exceptions are instruments MSG and Spot5, for which the IWT yields worse results than not using any transform. This could be explained by the low number of spectral bands captured by these instruments, which limits the decorrelation efficiency of the IWT but not its dynamic range expansion.

When it can be used, the AAT in the Recommended Standard produces the best results, except for the M3 and Vegetation instruments. For these instruments, the POT yields the best results. Only for the M3 corpus, the IWT attains better coding performance than the AAT. This can be explained by the large number of components of this set, since the overhead resulting from storing Q and V grows quadratically with N_Z.

For the CRISM, Hyperion, IASI, M3, MODIS, SFSI, and Vegetation instruments, JPEG2000 provides better coding performance than all tested CCSDS coders in terms of SNR, especially at high bit rates. These differences result from JPEG2000 use of larger blocks and a more complex entropy coder that uses multiple coding passes and contexts.

At low and medium bit rates, JPEG-LS generally produces the worst SNR results of all tested coders. For the CRISM, M3, MODIS, MSG, PLEIADES, SFSI, SPOT5, and Vegetation instruments, JPEG-LS outperforms all CCSDS coders at high or very high bit rates.

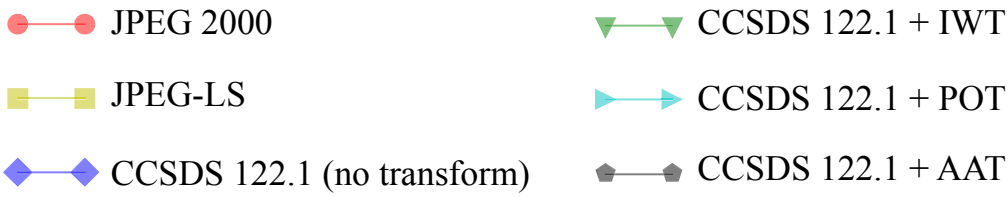

Figure 6-2: Legend for the Lossy Compression Performance Plots

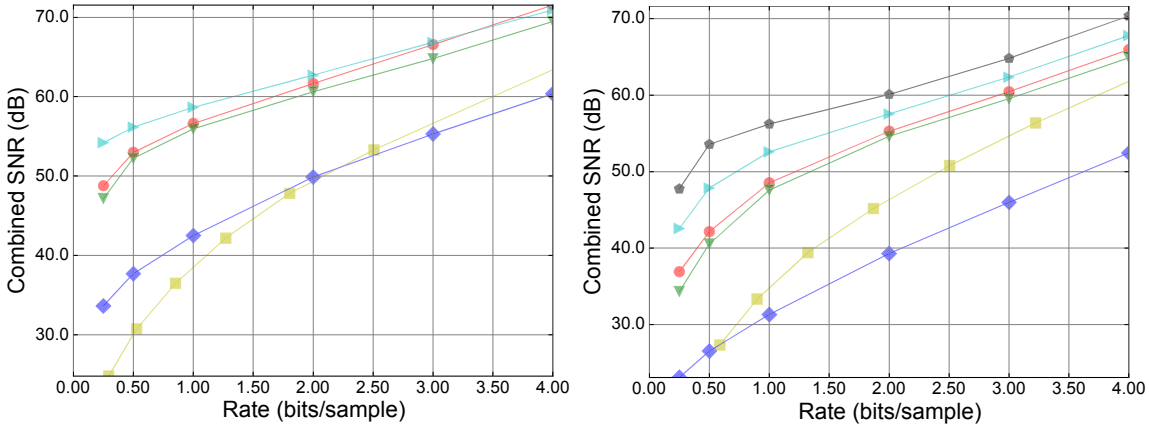

Figure 6-3: SNR Results for the AIRS (Left) and AVIRIS (Right) Instruments

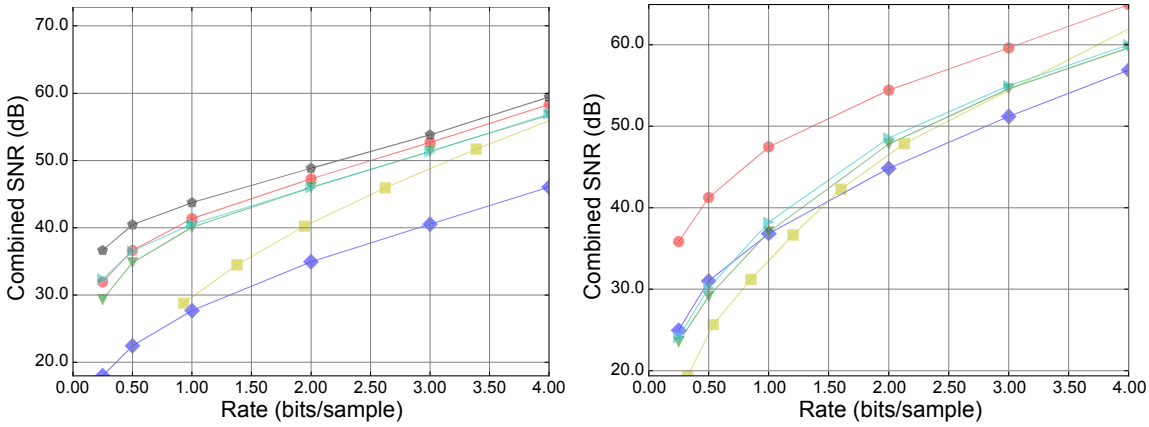

Figure 6-4: SNR Results for the CASI (Left) and CRISM (Right) Instruments

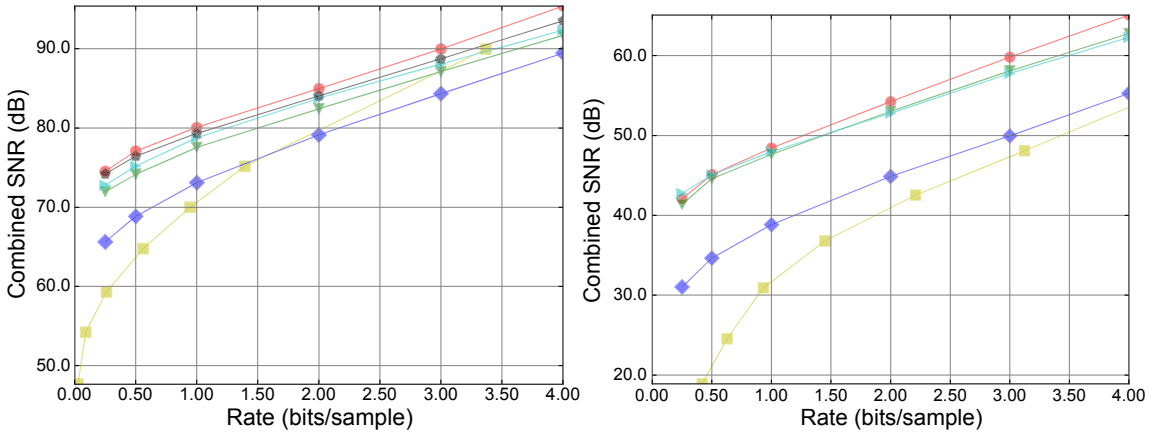

Figure 6-5: SNR Results for the Hyperion (Left) and IASI (Right) Instruments

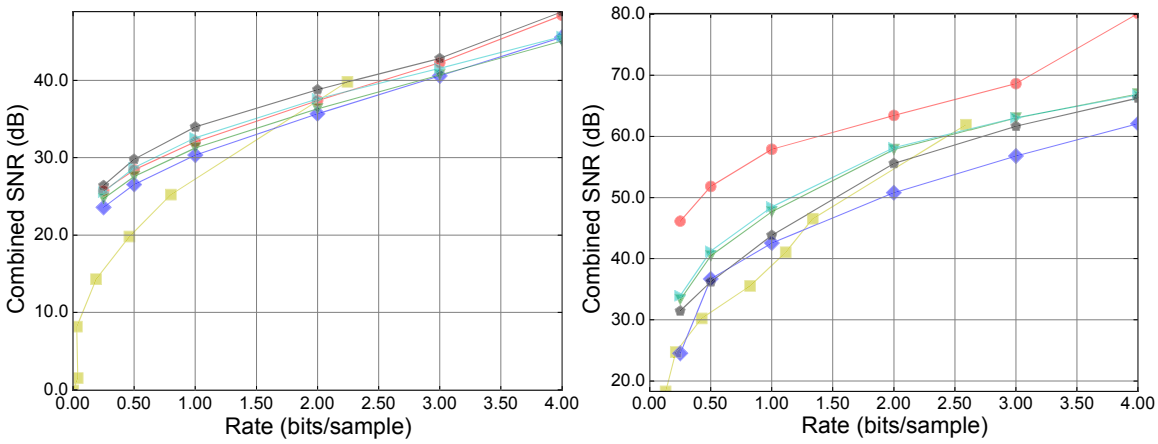

Figure 6-6: SNR Results for the Landsat (Left) and M3 (Right) Instruments

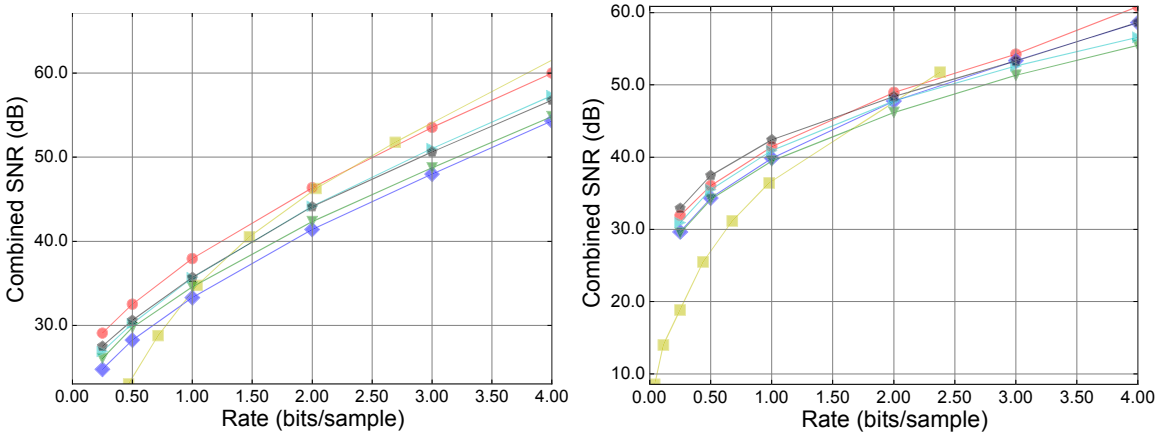

Figure 6-7: SNR Results for the MODIS (Left) and MSG (Right) Instruments

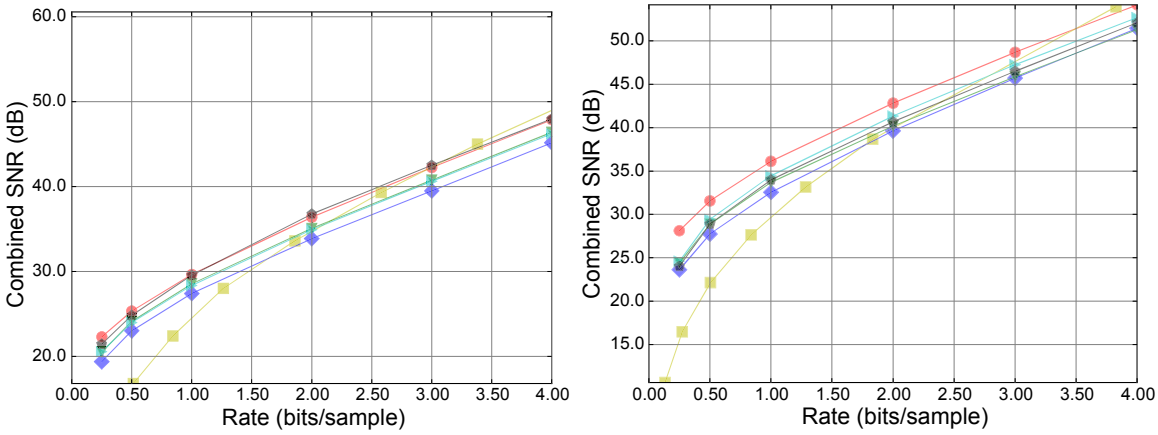

Figure 6-8: SNR Results for the PLEIADES (Left) and SFSI (Right) Instruments

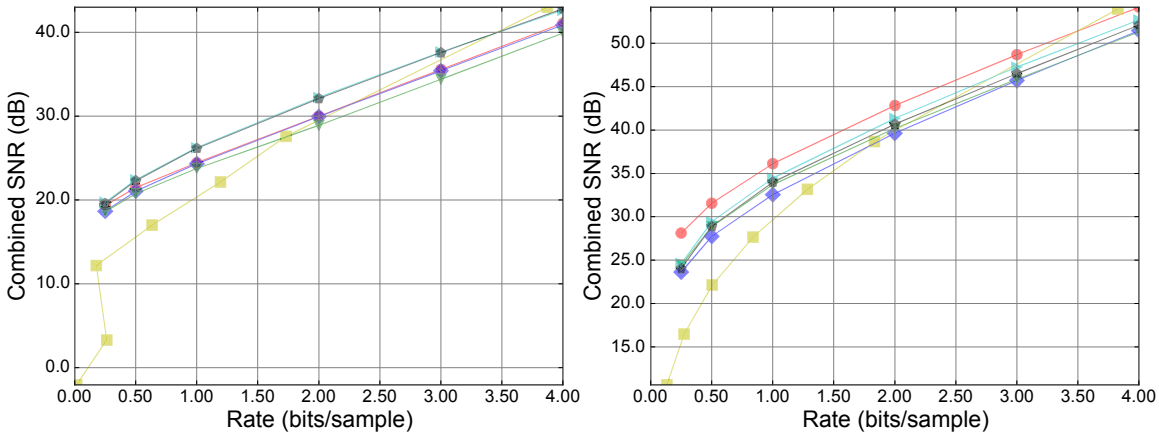

Figure 6-9: SNR Results for the SPOT5 (Left) and Vegetation (Right) Instruments

6.4 LOSSLESS COMPRESSION

Lossless compressed data rates averaged for all images of each instrument are provided in table 6-1. Individual image results are provided in table C-36 of annex subsection C6.

As expected, the use of any spectral transform improves upon not using any transform for most instruments. The only exceptions are data collected from the MODIS, MSG, and Spot5 instruments, for which the IWT produces worse results. This is explained by the low number of spectral bands produced by these instruments, for example, between 2 and 17, which hinders the decorrelation performance of that transform but not its dynamic range expansion.

For most instruments, the Recommended Standard attains the best lossless compression efficiency when the POT is used. Notwithstanding, the differences observed when using the IWT and the POT are generally small.

JPEG2000 provides the best results for most instruments, with small or moderate gains as compared to CCSDS 122.1. These gains can be explained by the use of larger coding blocks and a more complex arithmetic entropy coder.

JPEG-LS generally yields the worst results of all tested coders. This can be explained by the fact that JPEG-LS does not exploit the correlation among spectral bands. For instruments that produce only few bands, for example, the MODIS, MSG, and PLEIADES instruments, JPEG-LS produces the best lossless compression results, with moderate gains over CCSDS 122.1.

Table 6-1: Average Lossless Compressed Data Rates in Bits/Sample

	CCSDS 122.0 (No Transform)	CCSDS 122.1 + IWT	CCSDS 122.1 + POT	JPEG-LS	JPEG2000
AIRS	6.923	5.025	**4.711**	6.365	4.754
AVIRIS	7.079	5.183	**4.875**	6.888	5.039
CASI	8.465	6.658	6.688	8.168	**6.523**
CRISM	6.167	5.408	5.335	5.450	**4.640**
Hyperion	5.281	4.661	4.558	4.910	**4.301**
IASI	7.119	5.532	5.677	7.122	**5.128**
Landsat	4.042	3.881	3.906	3.695	**3.668**
M3	5.074	4.039	4.031	4.385	**3.192**
MODIS	6.668	6.827	6.480	**6.027**	6.258
MSG	3.899	4.178	4.068	**3.789**	3.959
PLEIADES	7.664	7.476	7.612	**7.431**	7.516
SFSI	5.020	4.591	4.462	4.558	**4.148**
SPOT5	5.656	5.716	**5.331**	5.443	5.678
Vegetation	5.516	5.436	5.453	5.335	**5.319**

7 STANDARD DEVELOPMENT CONSIDERATIONS

7.1 OVERVIEW

This section describes some of the decisions made in defining the Recommended Standard along with some of the motivation behind those decisions. Subsection 7.2 provides justification for the general design of the compressor. Subsection 7.3 motivates individual design choices.

7.2 GENERAL DESIGN MOTIVATION

The Recommended Standard has been designed taking into account the limitations of onboard systems.

As a result, several decisions have been made to trade compression efficiency for lower algorithm complexity and memory/buffering requirements.

7.3 DESIGN CHOICES

7.3.1 MAXIMUM U-D DIFFERENCE

Subsection 3.2.1 of the Recommended Standard restricts the upshift and downshift parameters to ensure that $U - D \leq 5$.

This constraint is imposed to limit the maximum dynamic range expansion before the application of the 2D encoders. (See 5.2 for a description of the maximum expansion for each spectral transform.)

As discussed in 3.2, larger values of $U - D$ do not provide any significant coding performance gain.

7.3.2 IWT LEVELS

Subsection 4.4.1 of the Recommended Standard fixes the number of decomposition levels of the IWT to five.

Experimental results indicate that for most settings, using five decomposition levels yields better average compression performance than any other number between three and twelve levels. In the remaining cases, five levels produce results similar to those of the best choice of transform levels.

Therefore, to maximize average compression effectiveness and user-friendliness while minimizing implementation complexity, the number of decomposition levels is fixed to five.

7.3.3 ARBITRARY AFFINE TRANSFORM PRECISION

The range of parameter Ψ has been chosen so that the rounding errors introduced in equation (87) of the Recommended Standard can be made sufficiently small if desired.

Error bounds as a function of Ψ can be found in 3.4.2.

7.3.4 ROUNDING OPERATIONS

7.3.4.1 General

Rounding operations are employed in different parts of the Recommended Standard to transform a real value *x* into an integer value *v*. The following rounding modes were considered:

a) rounding down; that is, $v = \lfloor x \rfloor$;

b) rounding half down; that is, $v = \lceil x - 0.5 \rceil$;

c) rounding half to zero; that is, $v = \text{sign}(x) \cdot \lceil \text{abs}(x) - 0.5 \rceil$, where sign(x) returns -1 and 1 when x is signed and unsigned, respectively (this operation is equivalent to equation (7) in reference [1]).

d) rounding half to even; that is,

$$v = \begin{cases} \lceil x - 0.5 \rceil & \text{if } x - \lfloor x \rfloor \neq 0.5 \\ \lfloor x \rfloor & \text{if } x - \lfloor x \rfloor = 0.5 \text{ and } \lceil x \rceil \text{ is even} \\ \lfloor x \rfloor & \text{if } x - \lfloor x \rfloor = 0.5 \text{ and } \lceil x \rceil \text{ is odd} \end{cases};$$

e) rounding half up; that is, $v = \lfloor x + 0.5 \rfloor$.

Subsections 7.3.4.2 to 7.3.4.5 motivate the choice of rounding operation in each part of the Standard.

7.3.4.2 Upshift and Downshift

Rounding is used in the downshift operation defined in equation (13) of reference [1] and is recommended for the inverse upshift operation defined in equation (12) of reference [1].

Because of image sample distributions, fractional values $x - \lfloor x \rfloor$ are more likely to be lower than 0.5. Moreover, rounding towards zero may reduce the number of bit planes that need to be processed by the 2D encoders. Therefore, rounding modes (b)-(c) defined in the previous section were evaluated experimentally.

Results indicate modest compression performance gains when using (c). In particular, differences smaller than 0.2 dB were observed for an early version of the Recommended Standard. Therefore, the round half to zero mode was included in equations (12) and (13) of the Recommended Standard.

7.3.4.3 IWT

Rounding is employed in the definition of the spectral IWT in equations (22), (23), (25), and (26) of reference [1].

The IWT is a well-known filter, thoroughly studied in the literature. In particular, it is employed in the JPEG2000 compression standard (see reference [13]).

The JPEG2000 implementation, which combines rounding modes (a) and (e), is adopted without modification in the Recommended Standard.

7.3.4.4 POT

Rounding is employed in several parts of the POT, including the calculation of the nominal mean and the pairwise operation weights.

In the nominal mean calculation defined in equation (34) of reference [1], all modes yield similar results. Therefore the computationally inexpensive mode (a) is recommended. Notwithstanding, since the word 'should' is used in the Recommended Standard, implementers may opt to use any other implementation.

The direct and inverse lifting operations defined in equations (50) and (55) of reference [1] produce almost identical coding performance results for modes (b)-(e), while using mode (a) yields worse results. Therefore, mode (e) is employed for its simplicity.

In the calculation of the pairwise operation weights $\widetilde{w_1}$, $\widetilde{w_2}$, and $\widetilde{w_3}$ and the intermediate quantities $\tilde{t}$, $\tilde{p}$, and $\widetilde{\sqrt{2}}$ (equations (41) to (49) of reference [1]), experiments indicate that mode (a) yields slightly worse results than modes (b)-(e). Since a fractional part of exactly 0.5 is unlikely to happen, the latter modes are essentially equivalent. Experimental results suggest almost identical performance for modes (b)-(e). Therefore mode (e) is also consistently chosen for its simplicity.

7.3.4.5 AAT

Rounding is specified in the forward calculation of the AAT in equation (86) of reference [1]. In order to provide rounding operation that is unbiased towards negative cases, mode (c) is used in the Recommended Standard.

ANNEX A

TEST IMAGES

Table A-1 summarizes the corpus of 47 hyperspectral and multispectral images used for compression testing and evaluation in the course of developing the Recommended Standard.

The corpus includes images from the following instruments:

- Atmospheric Infrared Sounder (AIRS);
- Airborne Visible/Infrared Imaging Spectrometer (AVIRIS);
- Compact Airborne Spectrographic Imager (CASI);
- Compact Reconnaissance Imaging Spectrometer for Mars (CRISM);
- Hyperion;
- Infrared Atmospheric Sounding Interferometer (IASI);
- Landsat;
- Moon Mineralogy Mapper (M3);
- Moderate Resolution Imaging Spectroradiometer (MODIS);
- Meteosat Second Generation (MSG);
- PLEIADES High Resolution;
- SWIR Full Spectrum Imager (SFSI);
- Système Pour l'Observation de la Terre 5 High Resolution Geometric (SPOT5);
- Vegetation.

IASI and MSG images are not available for public distribution. All other images can be downloaded at http://cwe.ccsds.org/sls/docs/sls-dc/123.0-B-Info/TestData.

For a few of these instruments, more than one type of image is included in the standard test set. The CRISM data includes Full-Resolution Target (FRT) images as well as Half-Resolution Long (HRL) and Multispectral Survey (MSP). The HRL and MSP images are produced on board the Mars Reconnaissance Orbiter spacecraft by spatial and spectral binning of the samples from the imager. Similarly, M3 'global' images are produced on board the Chandrayaan-1 spacecraft by spatial and spectral binning of the full resolution 'target' images. The 16-bit AVIRIS images were produced in 2006 from a newer version of the AVIRIS instrument than the 12-bit images, which are from 2001 and 2003.

Table A-1: Summary of the Corpus of Hyperspectral and Multispectral Test Images

Instrument	Image Type	Bit Depth, D	N_Z	N_X	N_Y (# images)
AIRS	raw	12−14	1501	90	135 (1)
AVIRIS	16-bit, raw	16	224	680	512 (1)
AVIRIS	12-bit, raw	12	224	614	512 (1)
AVIRIS	16-bit, calibrated	16	224	677	512 (1)
CASI	raw	12	72	406	1225 (3)
CRISM	FRT, raw	12	545	640	510 (6)
CRISM	HRL, raw	12	545	320	450 (2)
CRISM	MSP, raw	12	74	64	2700 (2)
Hyperion	raw	12	242	256	1024 (3)
IASI	calibrated	12	8461	66	60 (1)
Landsat	raw	8	6	1024	1024 (3)
M3	target, raw	12	260	640	512 (2)
M3	global, raw	12	86	320	512 (2)
MODIS	night, raw	12	17	1354	2030 (2)
MODIS	day, raw	12	14	1354	2030 (2)
MODIS	500m, raw	12	5	2708	4060 (2)
MODIS	250m, raw	12	2	5416	8120 (2)
MSG	calibrated	10	11	3712	3712 (1)
PLEIADES	HR, simulated	12	4	224	2456 (1), 2448 (3)
SFSI	raw	12	240	496	140 (3)
SPOT5	HRG, processed	8	3	1024	1024 (1)
Vegetation	raw	10	4	1728	10080 (2)

ANNEX B

AVAILABLE SOFTWARE, TEST DATA, AND LOOK-UP TABLES

B1 AVAILABLE SOFTWARE

Software implementations of the Recommended Standard may be available via links provided at the 'Implementations' section of https://www.ccsds.org. Not that some of these implementations may have only been developed to demonstrate compression performance; no optimization in speed or programming is guaranteed. The software and data are provided 'as is' to users as a courtesy. In no event will the CCSDS, its member Agencies, or the author of the software be liable for any consequential, incidental, or indirect damages arising out of the use of or inability to use the software.

B2 TEST DATA

Compressed versions of the test images as well as intermediate results of the compression process are available at http://cwe.ccsds.org/sls/docs/sls-dc/122.1-B-Info/TestData.

B3 LOOK-UP TABLES

The intermediate weights $\widetilde{w_1}$, $\widetilde{w_2}$, and $\widetilde{w_3}$ for the balanced and unbalanced pairwise operations of the POT can be stored in look-up tables as a function of the Ω, B, and C parameters to reduce computational complexity, as described in 5.4.2.

Reference look-up tables for these weights are provided at http://cwe.ccsds.org/sls/docs/sls-dc/122.1-B-Info/Tables.

ANNEX C

COMPRESSION RESULTS

C1 INTRODUCTION

This annex provides detailed compression results for the Recommended Standard, complementing the data presented in sections 3 and 6.

Results supporting the discussion on the Upshift and Downshift parameters in 3.2.3 are shown in C2. Results for the discussion on the POT in 2.3.3.4 are provided in C3. Results supporting 3.4.2 on the AAT are provided in C4. Results regarding the discussion on the encoder parameters in 3.5 are provided in C5. Results complementing the aggregated performance data shown in 6.3 and 6.4 are provided in C6.

SNR results are computed for each image and, when applicable, combined for each instrument, as described in 3.1. Unless stated otherwise, default parameters described in annex H are used to obtain all results.

C2 UPSHIFT AND DOWNSHIFT RESULTS

Average lossless compression results for $D = 0$ and several values of U are provided for each instrument in table C-1.

Table C-1: Average Lossless Compressed Data Rates in Bits/Sample for Different Values of *U*

	***U* = 0**	***U* = 1**	***U* = 2**	***U* = 3**	***U* = 4**	***U* = 5**
AIRS	5.025	6.054	7.101	8.153	9.207	10.262
AVIRIS	5.183	6.131	7.115	8.111	9.110	10.111
CASI	6.658	7.656	8.660	9.666	10.674	11.682
CRISM	5.408	6.371	7.362	8.361	9.364	10.367
Hyperion	4.661	5.630	6.619	7.614	8.612	9.611
IASI	5.532	6.617	7.732	8.855	9.981	11.108
Landsat	3.881	4.793	5.762	6.750	7.745	8.743
M3	4.038	4.901	5.855	6.838	7.831	8.828
MODIS	6.826	7.783	8.749	9.719	10.690	11.662
MSG	4.178	5.051	5.993	6.970	7.959	8.955
PLEIADES	7.476	8.472	9.470	10.470	11.470	12.470
SFSI	4.591	5.533	6.508	7.498	8.493	9.490
SPOT5	5.716	6.703	7.698	8.695	9.694	10.693

Combined SNR results for $U = 0$ and $D = 0$ for all instruments are provided in table C-2. Combined SNR results for different combinations of U and D are provided using $U = 0$, $D = 0$ as baseline in terms of ΔSNR in tables C-3 to C-8. ΔSNR is defined as

$$\Delta\text{SNR} = \text{SNR}_{U,D} - \text{SNR}_{0,0},$$

where $\text{SNR}_{U,D}$ is the combined SNR for upshift and downshift parameters U and D. Positive ΔSNR results indicate an image quality improvement over the baseline. In the tables, better performance is highlighted in green and worse performance in red.

Table C-2: Combined SNR Results in dB for $U = 0$ and $D = 0$ for Different Bit Rates in Bits/Sample

	Rate = 0.25	Rate = 0.5	Rate = 1	Rate = 2	Rate = 3	Rate = 4
AIRS	47.564	52.271	55.733	59.927	63.474	67.414
AVIRIS	34.583	40.771	47.679	54.567	59.355	64.352
CASI	29.571	35.004	40.166	46.025	51.353	56.644
CRISM	23.834	29.372	37.113	47.695	54.083	58.501
Hyperion	72.012	74.084	77.318	81.668	85.506	89.263
IASI	41.515	44.611	47.573	52.675	56.980	60.690
Landsat	24.756	27.414	30.793	35.107	38.862	42.635
M3	33.806	40.560	47.438	56.533	60.566	63.358
MODIS	26.298	29.887	34.652	42.346	48.683	54.642
MSG	29.708	34.160	39.289	44.965	49.065	52.808
PLEIADES	20.797	24.269	28.602	35.081	40.743	46.338
SFSI	50.380	53.961	56.830	60.823	65.489	69.557
SPOT5	18.725	20.784	23.733	28.807	34.160	39.088
VEGETATION	24.768	29.004	33.696	40.141	45.638	50.504

Table C-3: Combined ΔSNR Results in dB at 0.25 Bits/Sample for Different Choices of *U* and *D*, Using 0, 0 As Baseline

	0,1	0,2	1,1	3,3	1,0	2,0	3,0	6,3	4,1	4,0	5,0	6,1	8,3
AIRS	0.05	**0.13**	-0.02	0.02	-0.19	-0.41	-0.67	-0.67	-0.68	-0.91	-1.16	-1.16	-1.16
AVIRIS	0.13	**0.26**	0.00	0.00	-0.13	-0.26	-0.37	-0.37	-0.37	-0.48	-0.54	-0.54	-0.54
CASI	0.12	**0.25**	0.00	0.00	-0.15	-0.26	-0.37	-0.37	-0.37	-0.44	-0.49	-0.49	-0.49
CRISM	0.17	**0.31**	0.00	0.00	-0.15	-0.32	-0.47	-0.47	-0.47	-0.64	-0.74	-0.74	-0.74
Hyperion	-0.13	-0.28	-0.04	**0.04**	-0.02	-0.07	-0.14	-0.13	-0.14	-0.22	-0.30	-0.30	-0.30
IASI	0.01	**0.05**	-0.01	0.02	-0.07	-0.15	-0.23	-0.23	-0.23	-0.32	-0.41	-0.41	-0.41
Landsat	-0.08	-0.52	0.02	**0.04**	-0.01	-0.05	-0.12	-0.12	-0.12	-0.23	-0.34	-0.34	-0.34
M3	0.13	**0.26**	0.00	0.00	-0.27	-0.49	-0.79	-0.79	-0.79	-1.08	-1.40	-1.40	-1.40
MODIS	0.10	**0.18**	0.00	0.00	-0.11	-0.22	-0.31	-0.31	-0.31	-0.41	-0.45	-0.45	-0.45
MSG	**0.06**	-4.76	0.01	0.01	-0.11	-0.24	-0.42	-0.42	-0.42	-0.57	-0.70	-0.70	-0.70
PLEIADES	0.13	**0.26**	0.00	0.00	-0.16	-0.29	-0.43	-0.44	-0.44	-0.55	-0.64	-0.64	-0.64
SFSI	-0.07	-0.08	-0.01	**0.03**	-0.06	-0.12	-0.25	-0.25	-0.25	-0.36	-0.52	-0.52	-0.52
SPOT5	0.02	**0.04**	0.00	0.00	-0.12	-0.14	-0.32	-0.32	-0.32	-0.36	-0.53	-0.54	-0.54
VEGETATION	0.14	**0.21**	0.00	0.00	-0.14	-0.35	-0.51	-0.51	-0.51	-0.73	-0.85	-0.85	-0.85

Table C-4: Combined ΔSNR Results in dB at 0.5 Bits/Sample for Different Choices of *U* and *D*, Using 0, 0 As Baseline

	0,1	0,2	1,1	3,3	1,0	2,0	3,0	6,3	4,1	4,0	5,0	6,1	8,3
AIRS	-0.30	-0.41	-0.06	**0.06**	0.02	-0.03	-0.11	-0.11	-0.11	-0.20	-0.29	-0.29	-0.29
AVIRIS	0.09	**0.17**	0.00	0.00	-0.10	-0.21	-0.33	-0.33	-0.33	-0.44	-0.53	-0.53	-0.53
CASI	0.09	**0.17**	0.00	0.00	-0.11	-0.20	-0.28	-0.28	-0.28	-0.36	-0.41	-0.41	-0.41
CRISM	0.11	**0.19**	0.00	0.00	-0.11	-0.25	-0.39	-0.39	-0.39	-0.56	-0.69	-0.69	-0.69
Hyperion	-0.26	-0.47	-0.04	**0.08**	0.06	0.07	0.04	0.05	0.04	0.00	-0.05	-0.05	-0.05
IASI	-0.10	-0.17	-0.02	**0.04**	0.01	-0.02	-0.07	-0.07	-0.07	-0.13	-0.19	-0.19	-0.19
Landsat	-0.20	-0.88	0.06	0.09	0.10	**0.15**	0.12	0.13	0.13	0.09	0.03	0.03	0.03
M3	-0.07	-0.18	0.00	**0.03**	-0.05	-0.11	-0.27	-0.27	-0.27	-0.41	-0.60	-0.60	-0.60
MODIS	0.06	**0.10**	0.00	0.00	-0.06	-0.13	-0.19	-0.19	-0.19	-0.25	-0.28	-0.28	-0.28
MSG	-0.08	-8.24	0.03	**0.05**	0.02	0.00	-0.07	-0.07	-0.07	-0.15	-0.26	-0.26	-0.26
PLEIADES	0.05	**0.09**	0.00	0.00	-0.08	-0.15	-0.23	-0.23	-0.23	-0.29	-0.34	-0.34	-0.34
SFSI	-0.27	-0.47	-0.02	**0.08**	0.06	0.06	0.01	0.01	0.01	-0.03	-0.09	-0.09	-0.09
SPOT5	-0.01	-0.10	**0.00**	**0.00**	-0.04	-0.05	-0.12	-0.12	-0.12	-0.15	-0.24	-0.24	-0.24
VEGETATION	**0.02**	-0.02	0.00	0.00	-0.04	-0.12	-0.19	-0.19	-0.19	-0.29	-0.35	-0.35	-0.35

Table C-5: Combined ΔSNR Results in dB at 1 Bit/Sample for Different Choices of *U* and *D*, Using 0, 0 As Baseline

	0,1	0,2	1,1	3,3	1,0	2,0	3,0	6,3	4,1	4,0	5,0	6,1	8,3
AIRS	-0.68	-1.00	-0.13	0.16	0.15	**0.21**	0.21	0.21	0.21	0.18	0.15	0.16	0.16
AVIRIS	-0.01	-0.03	-0.01	**0.01**	-0.03	-0.09	-0.15	-0.15	-0.15	-0.22	-0.27	-0.27	-0.27
CASI	0.02	**0.04**	0.00	0.00	-0.04	-0.08	-0.13	-0.13	-0.13	-0.17	-0.20	-0.20	-0.20
CRISM	-0.01	-0.05	0.00	**0.01**	-0.04	-0.11	-0.19	-0.19	-0.19	-0.29	-0.38	-0.38	-0.38
Hyperion	-0.56	-1.04	-0.08	0.18	0.16	0.24	0.24	**0.25**	0.25	0.22	0.19	0.19	0.19
IASI	-0.27	-0.39	-0.04	**0.08**	0.05	0.06	0.03	0.04	0.04	0.00	-0.04	-0.04	-0.04
Landsat	-0.35	-1.77	0.15	0.25	0.29	0.45	0.47	**0.48**	0.48	0.46	0.44	0.44	0.44
M3	-0.42	-0.71	-0.06	0.11	0.18	0.24	0.24	**0.25**	0.24	0.22	0.14	0.14	0.14
MODIS	**0.02**	0.01	0.00	0.00	-0.03	-0.07	-0.11	-0.11	-0.11	-0.16	-0.18	-0.18	-0.18
MSG	-0.20	-12.88	0.13	**0.20**	0.12	0.15	0.16	0.17	0.17	0.15	0.14	0.14	0.14
PLEIADES	0.02	**0.03**	0.00	0.00	-0.03	-0.08	-0.12	-0.12	-0.12	-0.16	-0.20	-0.20	-0.20
SFSI	-0.54	-0.96	-0.05	0.19	0.18	**0.24**	0.24	0.24	0.24	0.22	0.19	0.19	0.19
SPOT5	-0.02	-0.19	0.00	0.00	0.00	**0.01**	-0.01	-0.01	-0.01	-0.02	-0.08	-0.08	-0.08
VEGETATION	-0.06	-0.25	-0.01	0.00	**0.01**	-0.01	-0.03	-0.03	-0.03	-0.07	-0.10	-0.10	-0.10

Table C-6: Combined ΔSNR Results in dB at 2 Bits/Sample for Different Choices of *U* and *D*, Using 0, 0 As Baseline

	0,1	0,2	1,1	3,3	1,0	2,0	3,0	6,3	4,1	4,0	5,0	6,1	8,3
AIRS	-1.60	-2.46	-0.35	0.44	0.46	0.69	0.73	**0.75**	0.73	0.74	0.72	0.72	0.72
AVIRIS	-0.27	-0.50	-0.06	0.05	0.05	**0.06**	0.06	0.06	0.06	0.03	0.01	0.01	0.01
CASI	-0.01	-0.03	0.00	**0.01**	-0.01	-0.04	-0.07	-0.07	-0.07	-0.09	-0.12	-0.12	-0.12
CRISM	-0.34	-0.63	-0.05	0.09	0.08	0.12	0.11	**0.12**	0.12	0.10	0.08	0.08	0.08
Hyperion	-1.46	-2.66	-0.19	0.54	0.52	0.76	0.81	**0.83**	0.81	0.81	0.81	0.81	0.81
IASI	-0.93	-1.39	-0.20	0.28	0.23	0.34	0.34	**0.35**	0.34	0.32	0.28	0.29	0.29
Landsat	-1.07	-4.24	0.34	0.63	0.69	1.21	1.34	1.38	1.37	1.37	1.38	1.38	**1.38**
M3	-2.80	-4.29	-0.73	0.78	0.78	1.33	1.47	1.51	1.47	1.51	1.53	1.53	**1.53**
MODIS	-0.03	-0.12	0.00	**0.01**	0.00	-0.01	-0.03	-0.03	-0.03	-0.06	-0.08	-0.08	-0.08
MSG	-0.82	-3.02	0.47	0.84	0.81	1.20	1.33	1.37	1.36	1.38	1.39	1.39	**1.39**
PLEIADES	-0.01	-0.05	0.00	**0.00**	0.00	-0.01	-0.03	-0.03	-0.03	-0.05	-0.08	-0.08	-0.08
SFSI	-1.25	-2.12	-0.22	0.43	0.36	0.57	0.59	**0.61**	0.60	0.58	0.56	0.56	0.56
SPOT5	-0.11	-0.69	0.02	0.02	0.05	0.09	0.11	**0.12**	**0.12**	0.12	0.08	0.08	0.08
VEGETATION	-0.22	-0.93	-0.03	0.01	0.02	0.04	0.04	**0.04**	0.04	0.03	0.02	0.02	0.02

Table C-7: Combined ΔSNR Results in dB at 3 Bits/Sample for Different Choices of *U* and *D*, Using 0, 0 As Baseline

	0,1	0,2	1,1	3,3	1,0	2,0	3,0	6,3	4,1	4,0	5,0	6,1	8,3
AIRS	-3.00	-4.05	-0.91	0.85	0.76	1.32	1.44	1.49	1.44	1.48	1.49	1.49	**1.49**
AVIRIS	-0.77	-1.33	-0.20	0.12	0.11	0.18	0.19	**0.20**	0.19	0.18	0.17	0.17	0.17
CASI	-0.07	-0.19	0.00	**0.03**	0.01	0.01	-0.01	-0.01	-0.01	-0.03	-0.06	-0.06	-0.06
CRISM	-1.36	-2.31	-0.31	0.33	0.31	0.50	0.54	**0.55**	0.54	0.54	0.54	0.54	0.55
Hyperion	-3.15	-5.09	-0.74	1.02	0.94	1.62	1.78	1.85	1.80	1.84	**1.86**	1.85	1.85
IASI	-2.32	-3.37	-0.57	0.76	0.73	1.11	1.18	**1.21**	1.18	1.18	1.16	1.16	1.16
Landsat	-3.07	-7.52	0.07	0.75	0.67	1.83	2.22	2.36	2.32	2.35	2.40	2.42	**2.42**
M3	-5.51	-7.64	-2.31	0.86	1.01	2.45	2.88	3.06	2.89	3.06	3.11	3.11	**3.13**
MODIS	-0.07	-0.32	0.01	0.03	0.03	**0.05**	0.04	0.04	0.04	0.03	0.02	0.02	0.02
MSG	-2.50	-5.96	0.29	1.29	1.13	2.26	2.66	2.80	2.77	2.81	2.85	2.86	**2.86**
PLEIADES	-0.04	-0.22	0.00	**0.01**	0.00	0.00	0.00	0.00	0.00	-0.02	-0.04	-0.04	-0.04
SFSI	-3.27	-4.90	-1.03	0.86	0.70	1.27	1.38	**1.43**	1.38	1.40	1.41	1.41	1.41
SPOT5	-0.42	-1.30	0.06	0.06	0.10	0.23	0.29	0.30	0.30	0.30	0.31	**0.31**	**0.31**
VEGETATION	-0.78	-2.51	-0.10	0.03	0.12	0.22	0.24	**0.24**	0.24	0.24	0.23	0.23	0.23

Table C-8: Combined ΔSNR Results in dB at 4 Bits/Sample for Different Choices of *U* and *D*, Using 0, 0 As Baseline

	0,1	0,2	1,1	3,3	1,0	2,0	3,0	6,3	4,1	4,0	5,0	6,1	8,3
AIRS	-5.60	-7.99	-2.59	0.92	0.83	2.07	2.35	**2.50**	2.35	2.45	2.48	2.47	2.49
AVIRIS	-2.25	-3.69	-0.74	0.23	0.28	0.54	0.59	**0.62**	0.59	0.60	0.59	0.59	0.59
CASI	-0.27	-0.62	-0.04	0.06	0.06	0.09	0.10	**0.10**	0.10	0.08	0.06	0.06	0.06
CRISM	-3.22	-4.97	-1.06	0.60	0.57	1.11	1.24	1.30	1.25	1.29	1.30	1.30	**1.30**
Hyperion	-5.98	-8.83	-2.58	0.71	0.87	2.44	2.99	3.18	3.00	3.18	3.21	3.21	**3.22**
IASI	-4.49	-6.37	-1.70	1.17	1.04	2.09	2.32	2.42	2.32	2.40	2.41	2.41	**2.42**
Landsat	-6.36	-11.29	-2.04	-0.98	0.20	2.44	3.66	4.14	4.02	4.06	4.20	4.27	**4.28**
M3	-8.02	-10.43	-4.28	-0.04	0.94	3.56	4.79	5.47	4.81	5.38	5.59	5.58	**5.63**
MODIS	-0.29	-1.13	0.03	0.08	0.11	0.18	0.20	**0.21**	0.21	0.20	0.20	0.20	0.20
MSG	-5.39	-9.70	-1.60	0.09	0.53	2.67	3.91	4.43	4.30	4.40	4.55	4.58	**4.59**
PLEIADES	-0.20	-0.85	-0.02	0.02	0.02	0.04	0.04	**0.04**	0.04	0.03	0.02	0.02	0.02
SFSI	-6.38	-8.97	-3.03	0.65	0.78	2.38	2.83	3.06	2.83	3.02	3.08	3.07	**3.09**
SPOT5	-0.59	-3.92	0.19	0.19	0.37	0.81	1.00	1.07	1.07	1.05	1.07	**1.07**	**1.07**
VEGETATION	-1.87	-4.97	-0.39	0.01	0.42	0.80	0.89	**0.92**	0.91	0.91	0.91	0.92	0.92

C3 POT RESULTS

Combined SNR results for the POT and all possible values of Ω are provided in tables C-9 to C-14, supporting 3.3.1. The best choice of Ω for each instrument is highlighted in bold font, while the worst choice is highlighted in red font.

Combined SNR results with and without discontinuity artifact mitigation strategies are provided in tables C-15 and C-16.

Table C-9: Combined SNR Results in dB at 0.25 Bits/Sample for All Possible Values of Ω

	Ω = 9	Ω = 10	Ω = 11	Ω = 12	Ω = 13	Ω = 14	Ω = 15	Ω = 16
AIRS	51.16	51.39	51.56	51.63	**51.63**	51.63	51.62	51.61
AVIRIS	40.80	41.06	41.22	41.24	**41.24**	41.23	41.23	41.23
CASI	31.23	31.48	31.54	**31.56**	31.55	31.55	31.54	31.54
CRISM	23.91	23.92	**23.92**	23.92	23.92	23.91	23.91	23.91
Hyperion	72.30	72.32	72.33	**72.33**	72.33	72.33	72.33	72.33
IASI	42.28	42.34	42.37	42.38	42.38	42.38	42.37	42.37
Landsat	25.21	25.21	25.21	**25.21**	25.21	25.21	25.21	25.20
M3	33.81	33.85	**33.87**	33.87	33.86	33.86	33.85	33.85
MODIS	26.82	**26.82**	26.82	26.82	26.82	26.82	26.82	26.82
MSG	31.14	31.14	31.14	31.15	31.15	31.15	**31.15**	31.15
PLEIADES	20.35	20.35	20.36	20.36	**20.36**	20.36	20.36	20.36
SFSI	53.21	53.23	53.24	53.24	**53.24**	53.24	53.24	53.24
SPOT5	19.58	19.59	19.59	19.59	**19.59**	19.59	19.59	19.59
VEGETATION	24.46	24.46	24.47	24.47	**24.47**	24.47	24.47	24.47

Table C-10: Combined SNR Results in dB at 0.5 Bits/Sample for All Possible Values of Ω

	Ω = 9	Ω = 10	Ω = 11	Ω = 12	Ω = 13	Ω = 14	Ω = 15	Ω = 16
AIRS	54.99	55.12	55.19	55.22	55.23	**55.23**	55.23	55.23
AVIRIS	46.90	47.25	47.42	47.47	47.48	**47.48**	47.48	47.48
CASI	35.94	36.12	36.19	**36.21**	36.21	36.20	36.20	36.20
CRISM	30.07	30.08	30.09	**30.09**	30.09	30.09	30.08	30.08
Hyperion	74.88	74.90	74.91	**74.91**	74.91	74.91	74.91	74.91
IASI	44.84	44.88	44.90	44.91	**44.91**	44.91	44.91	44.91
Landsat	28.48	28.48	28.48	**28.48**	28.48	28.48	28.48	28.48
M3	41.09	41.15	41.20	**41.20**	41.20	41.20	41.20	41.20
MODIS	30.14	**30.14**	30.14	30.14	30.14	30.14	30.14	30.14
MSG	35.73	35.74	35.75	35.75	35.75	35.75	**35.75**	35.75
PLEIADES	23.84	23.84	23.85	**23.85**	23.85	23.84	23.84	23.84
SFSI	55.54	55.55	55.56	55.56	**55.56**	55.56	55.56	55.56
SPOT5	22.38	22.38	22.38	22.38	22.38	22.38	22.38	**22.38**
VEGETATION	29.41	29.42	29.42	29.42	29.42	**29.42**	29.42	29.42

Table C-11: Combined SNR Results in dB at 1 Bit/Sample for All Possible Values of Ω

	Ω = 9	Ω = 10	Ω = 11	Ω = 12	Ω = 13	Ω = 14	Ω = 15	Ω = 16
AIRS	57.81	57.94	57.99	58.01	58.02	58.02	**58.02**	58.02
AVIRIS	51.92	52.24	52.39	52.44	52.46	52.46	**52.46**	52.46
CASI	40.20	40.36	40.42	40.44	40.44	40.45	**40.45**	40.45
CRISM	38.02	38.04	38.05	38.05	**38.05**	38.05	38.05	38.05
Hyperion	78.50	78.52	78.53	**78.53**	78.53	78.53	78.53	78.53
IASI	47.78	47.82	47.84	**47.85**	47.85	47.85	47.85	47.85
Landsat	32.39	32.39	32.39	32.40	32.40	**32.40**	32.40	32.39
M3	48.31	48.34	48.38	48.39	48.40	**48.40**	48.40	48.40
MODIS	**35.61**	35.61	35.61	35.60	35.60	35.60	35.60	35.60
MSG	41.03	41.04	41.04	41.04	41.04	41.04	41.04	**41.04**
PLEIADES	**28.30**	28.29	28.29	28.29	28.29	28.29	28.29	28.29
SFSI	58.00	58.01	58.01	58.01	**58.01**	58.01	58.01	58.01
SPOT5	26.21	26.21	26.21	26.21	26.21	**26.21**	26.21	26.21
VEGETATION	34.45	34.45	34.46	34.46	34.46	34.46	34.46	**34.46**

Table C-12: Combined SNR Results in dB at 2 Bits/Sample for All Possible Values of Ω

	Ω = 9	Ω = 10	Ω = 11	Ω = 12	Ω = 13	Ω = 14	Ω = 15	Ω = 16
AIRS	62.06	62.17	62.22	62.23	62.23	**62.24**	62.24	62.23
AVIRIS	57.10	57.32	57.41	57.43	57.44	57.44	**57.44**	57.44
CASI	45.69	45.82	45.87	45.89	45.90	45.90	45.90	**45.90**
CRISM	48.25	48.26	48.27	48.28	**48.28**	48.28	48.28	48.28
Hyperion	83.59	83.61	83.61	**83.61**	83.61	83.61	83.61	83.61
IASI	52.65	52.69	52.71	52.72	**52.72**	52.72	52.72	52.72
Landsat	37.47	37.47	37.47	**37.48**	37.48	37.48	37.48	37.47
M3	58.01	58.01	58.01	58.02	58.03	58.04	58.04	**58.04**
MODIS	**44.12**	44.11	44.10	44.09	44.09	44.09	44.09	44.09
MSG	47.88	47.88	47.89	47.89	47.89	47.89	47.89	**47.89**
PLEIADES	**34.83**	34.83	34.83	34.83	34.82	34.82	34.82	34.82
SFSI	62.44	62.44	**62.44**	62.44	62.44	62.44	62.44	62.44
SPOT5	32.19	32.19	32.20	32.20	**32.20**	32.20	32.20	32.20
VEGETATION	41.31	41.31	41.32	41.32	41.32	41.32	41.32	**41.32**

Table C-13: Combined SNR Results in dB at 3 Bits/Sample for All Possible Values of Ω

	Ω = 9	Ω = 10	Ω = 11	Ω = 12	Ω = 13	Ω = 14	Ω = 15	Ω = 16
AIRS	66.19	66.28	66.32	66.34	66.34	**66.35**	66.34	66.34
AVIRIS	61.97	62.15	62.22	62.25	62.26	62.26	62.26	**62.26**
CASI	51.07	51.18	51.22	51.24	51.25	51.25	51.25	**51.25**
CRISM	54.74	54.74	54.75	54.75	54.75	**54.75**	54.75	54.75
Hyperion	87.96	87.96	87.98	87.99	**87.99**	87.98	87.96	87.98
IASI	57.63	57.68	57.70	**57.70**	57.70	57.70	57.70	57.70
Landsat	41.41	41.42	41.41	**41.43**	41.42	41.42	41.42	41.42
M3	62.92	62.92	62.93	62.93	62.93	62.94	62.94	**62.95**
MODIS	**51.02**	50.99	50.97	50.96	50.96	50.96	50.95	50.95
MSG	52.67	52.66	52.67	52.66	52.66	52.67	52.67	**52.67**
PLEIADES	**40.52**	40.51	40.51	40.51	40.50	40.50	40.50	40.50
SFSI	67.75	67.75	67.75	67.75	67.75	67.75	67.76	**67.76**
SPOT5	37.61	37.62	37.63	37.63	37.63	**37.63**	37.63	37.63
VEGETATION	47.23	47.23	47.24	47.24	47.24	47.24	**47.24**	47.24

Table C-14: Combined SNR Results in dB at 4 Bits/Sample for All Possible Values of Ω

	Ω = 9	Ω = 10	Ω = 11	Ω = 12	Ω = 13	Ω = 14	Ω = 15	Ω = 16
AIRS	70.43	70.48	70.52	70.52	70.52	**70.53**	70.53	70.52
AVIRIS	67.41	67.59	67.67	67.69	67.70	67.70	**67.70**	67.70
CASI	56.71	56.82	56.85	56.87	56.87	56.88	**56.88**	56.88
CRISM	59.80	59.80	59.81	59.81	59.81	59.81	59.82	**59.82**
Hyperion	92.14	92.19	**92.21**	92.21	92.18	92.19	92.20	92.21
IASI	62.13	62.18	62.20	62.21	62.20	62.20	62.21	**62.22**
Landsat	45.56	45.56	45.55	**45.58**	45.57	45.57	45.57	45.57
M3	66.54	66.57	66.60	66.60	66.61	66.63	66.64	**66.66**
MODIS	**57.33**	57.29	57.27	57.25	57.25	57.25	57.24	57.24
MSG	56.58	56.58	56.58	56.57	56.57	56.57	56.58	**56.58**
PLEIADES	**46.11**	46.11	46.10	46.10	46.10	46.10	46.10	46.10
SFSI	72.56	72.56	72.57	72.57	72.57	72.57	72.58	**72.59**
SPOT5	42.60	42.61	42.61	42.61	42.61	**42.61**	42.61	42.61
VEGETATION	52.66	52.66	52.66	52.66	52.66	52.66	52.66	**52.66**

Table C-15: Combined SNR Results for the POT with and without Constant Parameter Selection at Bitrates between 0.25 and 1 Bit/Sample

	Rate = 0.25 bit / sample		Rate = 0.5 bits / sample		Rate =1 bit / sample	
	POT	**POT + CPS**	**POT**	**POT + CPS**	**POT**	**POT + CPS**
AIRS	51.63	54.04	55.22	56.12	58.01	58.62
AVIRIS	41.24	42.49	47.47	47.82	52.44	52.59
CASI	31.57	32.14	36.21	36.41	40.44	40.46
CRISM	23.92	24.09	30.09	30.18	38.06	38.24
Hyperion	72.34	72.83	74.91	75.27	78.53	78.87
IASI	42.37	42.78	44.91	45.12	47.85	47.99
Landsat	25.21	25.53	28.48	28.67	32.40	32.51
M3	33.87	33.93	41.20	41.24	48.39	48.48
MODIS	26.82	26.75	30.14	30.00	35.60	35.26
MSG	31.15	30.84	35.75	35.44	41.04	40.86
PLEIADES	20.36	20.54	23.85	23.95	28.29	28.27
SFSI	53.24	53.29	55.56	55.58	58.01	58.03
SPOT5	19.59	19.70	22.38	22.41	26.21	26.24
VEGETATION	24.47	24.02	29.42	29.20	34.46	34.21

Table C-16: Combined SNR Results for the POT with and without Constant Parameter Selection at Bitrates between 2 and 4 Bit/Sample

	Rate = 2 bits / sample		Rate = 3 bits / sample		Rate =4 bits / sample	
	POT	POT + CPS	POT	POT + CPS	POT	POT + CPS
AIRS	62.23	62.70	66.33	66.83	70.52	70.92
AVIRIS	57.43	57.55	62.25	62.37	67.69	67.79
CASI	45.89	45.89	51.24	51.23	56.87	56.86
CRISM	48.28	48.63	54.75	55.05	59.81	60.06
Hyperion	83.61	83.85	87.99	88.14	92.20	92.37
IASI	52.72	52.82	57.70	57.81	62.21	62.32
Landsat	37.48	37.56	41.43	41.51	45.58	45.61
M3	58.02	58.21	62.93	63.10	66.60	66.82
MODIS	44.10	43.87	50.97	50.78	57.26	57.09
MSG	47.89	47.77	52.66	52.69	56.57	56.72
PLEIADES	34.83	34.74	40.51	40.43	46.10	46.03
SFSI	62.44	62.46	67.75	67.79	72.57	72.69
SPOT5	32.20	32.23	37.63	37.67	42.61	42.65
VEGETATION	41.32	41.05	47.24	47.00	52.66	52.49

C4 AAT RESULTS

Combined SNR results for the AAT using different values of Ψ are provided in tables C-17 to C-22.

For instruments having a large number of spectral bands, the use of the AAT may not be computationally feasible, and side information may become larger than the uncompressed image, especially for small values of N_X and N_Y. Results for the CRISM and IASI instruments are not provided in the tables. For other instruments, side information requirements higher than the target bit rate are identified with a dash.

Table C-17: Combined SNR Results in dB at 0.25 Bits/Sample for Different Values of Ψ

	Ψ = 0	Ψ = 7	Ψ = 15	Ψ = 23	Ψ = 28	Ψ = 30
AIRS	-	-	-	-	-	-
AVIRIS	13.506	38.715	**47.761**	47.554	47.418	47.362
CASI	10.765	35.358	**36.656**	36.594	36.555	36.539
Hyperion	54.530	73.339	**74.174**	74.103	74.057	74.039
Landsat	10.781	26.447	**26.466**	26.466	26.465	26.465
M3	15.752	31.375	**31.548**	31.386	31.272	31.224
MODIS	9.513	27.586	**27.607**	27.607	27.606	27.606
MSG	4.722	32.752	**33.039**	33.039	33.038	33.038
PLEIADES	10.295	**21.525**	21.525	21.521	21.520	21.520
SFSI	29.920	51.568	**53.144**	52.624	52.220	52.023
SPOT5	13.230	19.652	**19.653**	19.653	19.652	19.652
VEGETATION	0.712	24.158	**24.177**	24.177	24.176	24.176

Table C-18: Combined SNR Results in dB at 0.5 Bits/Sample for Different Values of Ψ

	Ψ = 0	Ψ = 7	Ψ = 15	Ψ = 23	Ψ = 28	Ψ = 30
AIRS	**19.460**	-	-	-	-	-
AVIRIS	13.519	38.949	**53.559**	53.510	53.478	53.466
CASI	10.859	37.041	**40.464**	40.444	40.432	40.426
Hyperion	54.531	74.915	**76.328**	76.286	76.259	76.248
Landsat	10.750	29.757	**29.806**	29.805	29.805	29.805
M3	15.756	35.385	**36.277**	36.189	36.134	36.113
MODIS	9.484	30.612	**30.663**	30.662	30.662	30.662
MSG	4.719	36.739	37.493	**37.493**	37.493	37.493
PLEIADES	10.325	24.898	**24.906**	24.904	24.903	24.902
SFSI	29.920	52.769	**55.687**	55.467	55.321	55.259
SPOT5	13.231	22.355	22.357	**22.357**	22.357	22.357
VEGETATION	0.694	28.906	**28.962**	28.962	28.962	28.962

Table C-19: Combined SNR Results in dB at 1 Bit/Sample for Different Values of Ψ

	Ψ = 0	Ψ = 7	Ψ = 15	Ψ = 23	Ψ = 28	Ψ = 30
AIRS	**19.460**	-	-	-	-	-
AVIRIS	13.522	38.938	**56.167**	56.150	56.137	56.132
CASI	10.888	38.086	**43.743**	43.730	43.721	43.718
Hyperion	54.531	76.534	**78.916**	78.880	78.856	78.847
Landsat	10.711	33.790	**33.936**	33.935	33.935	33.935
M3	15.757	39.780	**43.789**	43.718	43.673	43.655
MODIS	9.469	35.610	**35.769**	35.769	35.768	35.768
MSG	4.719	40.386	42.363	**42.364**	42.364	42.364
PLEIADES	10.327	29.589	**29.611**	29.610	29.610	29.609
SFSI	29.920	53.675	**57.833**	57.750	57.698	57.676
SPOT5	13.541	26.190	26.193	**26.193**	26.193	26.193
VEGETATION	0.688	33.824	**34.003**	34.002	34.002	34.002

Table C-20: Combined SNR Results in dB at 2 Bits/Sample for Different Values of Ψ

	Ψ = 0	Ψ = 7	Ψ = 15	Ψ = 23	Ψ = 28	Ψ = 30
AIRS	19.460	**37.851**	-	-	-	-
AVIRIS	13.522	38.957	**59.734**	59.720	59.707	59.701
CASI	10.895	38.926	**48.862**	48.850	48.842	48.839
Hyperion	54.531	78.166	**82.514**	82.487	82.472	82.466
Landsat	10.711	38.168	38.622	**38.623**	38.622	38.622
M3	15.758	41.727	**54.657**	54.614	54.587	54.577
MODIS	9.472	43.150	**44.143**	44.143	44.143	44.142
MSG	4.720	43.155	48.211	**48.216**	48.216	48.216
PLEIADES	10.346	36.665	**36.779**	36.778	36.778	36.777
SFSI	29.920	54.660	**61.178**	61.064	60.994	60.966
SPOT5	13.826	32.100	**32.104**	32.104	32.104	32.104
VEGETATION	0.687	39.907	**40.664**	40.664	40.664	40.664

Table C-21: Combined SNR Results in dB at 3 Bits/Sample for Different Values of Ψ

	Ψ = 0	Ψ = 7	Ψ = 15	Ψ = 23	Ψ = 28	Ψ = 30
AIRS	19.460	37.840	**58.054**	50.559	-	-
AVIRIS	13.522	38.973	63.649	**63.649**	63.637	63.632
CASI	10.896	39.227	**53.805**	53.795	53.788	53.785
Hyperion	54.531	78.898	**84.980**	84.969	84.960	84.957
Landsat	10.706	41.081	42.242	**42.242**	42.242	42.242
M3	15.758	41.854	**58.307**	58.302	58.297	58.295
MODIS	9.473	47.305	50.623	**50.623**	50.623	50.623
MSG	4.720	44.124	52.679	**52.694**	52.694	52.694
PLEIADES	10.352	42.100	**42.504**	42.503	42.503	42.502
SFSI	29.920	55.194	**64.499**	64.420	64.370	64.350
SPOT5	13.905	37.563	**37.568**	37.568	37.567	37.567
VEGETATION	0.687	44.109	**46.479**	46.479	46.479	46.479

Table C-22: Combined SNR Results in dB at 4 Bits/Sample for Different Values of Ψ

	Ψ = 0	Ψ = 7	Ψ = 15	Ψ = 23	Ψ = 28	Ψ = 30
AIRS	19.460	37.842	57.720	**58.063**	57.800	57.152
AVIRIS	13.522	38.982	67.006	**67.037**	67.030	67.027
CASI	10.896	39.333	**59.357**	59.348	59.339	59.335
Hyperion	54.531	79.180	**86.291**	86.281	86.272	86.270
Landsat	10.709	43.377	46.473	**46.476**	46.476	46.476
M3	15.758	41.877	59.451	**59.451**	59.450	59.450
MODIS	9.473	49.173	56.740	**56.741**	56.741	56.741
MSG	4.721	44.443	56.064	**56.110**	56.109	56.109
PLEIADES	10.354	46.684	**47.942**	47.941	47.941	47.941
SFSI	29.920	55.370	**66.356**	66.327	66.307	66.300
SPOT5	13.934	42.722	**42.739**	42.738	42.738	42.738
VEGETATION	0.687	46.451	52.007	**52.007**	52.007	52.007

C5 ENCODER RESULTS

Combined SNR results for the IWT for different values of R are provided in tables C-23 to C-28. Results in these tables are obtained using $S' = \lceil N_X/8 \rceil$.

Combined SNR results for the POT for all allowed values of F are provided in tables C-29 to C-34. In these tables, dashes indicate that the required side information is larger than the target bit rate.

Table C-23: Combined SNR Results in dB at 0.25 Bits/Sample for Different Values of *R*

	R=1	*R*=10	*R*=100	*R*=1000	*R*=65535
AIRS	42.50	**42.89**	**42.89**	**42.89**	**42.89**
AVIRIS	34.01	34.04	**34.07**	**34.07**	**34.07**
CASI	26.94	27.32	27.91	**28.31**	**28.31**
CRISM	23.00	23.00	23.00	**23.00**	**23.00**
Hyperion	70.78	70.83	**70.96**	70.96	70.96
IASI	37.32	**37.63**	**37.63**	**37.63**	**37.63**
Landsat	24.65	24.69	**24.71**	24.71	24.71
M3	32.35	32.35	**32.35**	**32.35**	**32.35**
MODIS	26.04	26.14	26.21	**26.27**	26.27
MSG	29.47	29.49	29.74	**30.61**	**30.61**
PLEIADES	18.60	18.97	19.06	**19.43**	**19.43**
SFSI	49.57	49.60	**49.60**	**49.60**	**49.60**
SPOT5	18.48	18.56	18.58	**18.58**	**18.58**
VEGETATION	24.34	24.42	24.58	24.88	**24.98**

Table C-24: Combined SNR Results in dB at 0.5 Bits/Sample for Different Values of *R*

	***R* = 1**	***R* = 10**	***R* = 100**	***R* = 1000**	***R* = 65535**
AIRS	50.46	**50.93**	**50.93**	**50.93**	**50.93**
AVIRIS	40.25	40.31	**40.36**	**40.36**	**40.36**
CASI	33.01	33.40	34.00	**34.30**	**34.30**
CRISM	28.68	28.69	28.69	**28.69**	**28.69**
Hyperion	73.66	73.68	73.72	**73.72**	**73.72**
IASI	43.29	**43.36**	**43.36**	**43.36**	**43.36**
Landsat	27.51	27.54	27.55	**27.56**	**27.56**
M3	39.76	**39.76**	39.76	39.76	39.76
MODIS	29.62	29.74	29.84	30.05	**30.05**
MSG	33.87	33.90	34.21	**34.70**	**34.70**
PLEIADES	23.08	23.22	23.30	**23.81**	**23.81**
SFSI	53.82	**53.84**	53.83	53.83	53.83
SPOT5	20.71	20.74	**20.74**	20.74	20.74
VEGETATION	28.69	28.75	28.99	29.50	**29.81**

Table C-25: Combined SNR Results in dB at 1 Bit/Sample for Different Values of *R*

	***R* = 1**	***R* = 10**	***R* = 100**	***R* = 1000**	***R* = 65535**
AIRS	55.28	**55.55**	**55.55**	**55.55**	**55.55**
AVIRIS	47.40	47.46	**47.49**	**47.49**	**47.49**
CASI	39.01	39.30	39.63	**39.89**	**39.89**
CRISM	36.71	36.72	36.72	**36.72**	**36.72**
Hyperion	77.20	77.22	**77.26**	77.26	77.26
IASI	46.90	**46.95**	**46.95**	**46.95**	**46.95**
Landsat	31.21	31.23	31.24	**31.24**	**31.24**
M3	47.30	47.30	**47.30**	**47.30**	**47.30**
MODIS	34.31	34.45	34.68	34.79	**34.79**
MSG	39.02	39.05	39.27	**40.24**	**40.24**
PLEIADES	27.65	27.78	27.85	**28.65**	**28.65**
SFSI	56.98	56.99	**56.99**	**56.99**	**56.99**
SPOT5	23.73	**23.76**	23.74	23.75	23.75
VEGETATION	33.38	33.44	33.71	34.35	**34.54**

Table C-26: Combined SNR Results in dB at 2 Bits/Sample for Different Values of *R*

	R = 1	*R* = 10	*R* = 100	*R* = 1000	*R* = 65535
AIRS	60.19	**60.31**	**60.31**	**60.31**	**60.31**
AVIRIS	54.54	54.57	**54.59**	**54.59**	**54.59**
CASI	45.39	45.58	45.74	**45.82**	**45.82**
CRISM	47.62	47.63	47.64	**47.64**	**47.64**
Hyperion	82.13	82.14	82.18	**82.18**	**82.18**
IASI	52.32	**52.43**	**52.43**	**52.43**	**52.43**
Landsat	36.29	36.30	36.31	**36.32**	**36.32**
M3	57.64	57.64	**57.64**	**57.64**	**57.64**
MODIS	41.69	42.09	42.34	42.80	**42.80**
MSG	45.75	45.77	46.05	**47.62**	**47.62**
PLEIADES	34.21	34.33	34.45	**35.44**	**35.44**
SFSI	61.29	61.35	**61.35**	**61.35**	**61.35**
SPOT5	28.84	28.87	**28.90**	28.90	28.90
VEGETATION	39.76	39.81	40.18	40.77	**41.31**

Table C-27: Combined SNR Results in dB at 3 Bits/Sample for Different Values of *R*

	R = 1	*R* = 10	*R* = 100	*R* = 1000	*R* = 65535
AIRS	64.44	**64.50**	**64.50**	**64.50**	**64.50**
AVIRIS	59.41	59.46	**59.49**	**59.49**	**59.49**
CASI	50.73	50.89	51.06	**51.20**	**51.20**
CRISM	54.44	54.45	54.46	**54.46**	**54.46**
Hyperion	86.84	86.84	86.86	**86.86**	**86.86**
IASI	57.26	**57.42**	**57.42**	**57.42**	**57.42**
Landsat	40.65	40.68	40.69	**40.69**	**40.69**
M3	62.92	62.93	**62.93**	**62.93**	**62.93**
MODIS	48.16	48.49	48.72	**49.25**	49.25
MSG	51.01	51.02	51.22	**52.45**	**52.45**
PLEIADES	39.86	39.98	40.08	**41.17**	**41.17**
SFSI	66.58	66.68	**66.69**	**66.69**	**66.69**
SPOT5	34.27	34.32	**34.39**	34.39	34.39
VEGETATION	45.39	45.44	45.84	46.65	**47.24**

Table C-28: Combined SNR Results in dB at 4 Bits/Sample for Different Values of *R*

	R = 1	*R* = 10	*R* = 100	*R* = 1000	*R* = 65535
AIRS	69.03	**69.14**	**69.14**	**69.14**	**69.14**
AVIRIS	64.74	64.77	**64.84**	**64.84**	**64.84**
CASI	56.12	56.27	56.51	**56.56**	**56.56**
CRISM	59.48	59.50	59.51	**59.51**	**59.51**
Hyperion	91.41	91.43	91.47	**91.48**	**91.48**
IASI	62.09	**62.25**	**62.25**	**62.25**	**62.25**
Landsat	45.03	45.05	45.06	**45.07**	**45.07**
M3	66.80	66.80	**66.80**	**66.80**	**66.80**
MODIS	54.27	54.58	54.80	55.42	**55.42**
MSG	55.26	55.28	55.43	**55.92**	**55.92**
PLEIADES	45.49	45.62	45.72	**46.70**	**46.70**
SFSI	71.71	71.84	**71.87**	**71.87**	**71.87**
SPOT5	39.82	39.88	**39.90**	39.90	39.90
VEGETATION	50.89	50.94	51.29	52.17	**52.68**

Table C-29: Combined SNR Results in dB at 0.25 Bits/Sample for Different Values of *F*

	F = 1	*F* = 2	*F* = 4	*F* = 8	*F* = 16	*F* = 32
AIRS	-	32.50	43.34	47.15	49.34	**51.63**
AVIRIS	38.83	40.08	40.65	40.96	41.13	**41.24**
CASI	28.60	30.03	30.71	31.15	31.39	**31.56**
CRISM	-	-	23.43	23.70	23.85	**23.92**
Hyperion	66.96	69.69	70.88	71.62	72.05	**72.33**
IASI	-	-	37.30	39.59	41.16	**42.38**
Landsat	23.18	23.85	24.31	24.66	25.01	**25.21**
M3	31.02	32.53	33.24	33.59	33.77	**33.87**
MODIS	26.29	26.37	26.55	26.71	26.80	**26.82**
MSG	30.43	30.60	30.73	30.83	31.00	**31.15**
PLEIADES	17.42	19.16	19.81	20.11	20.27	**20.36**
SFSI	51.98	52.69	52.97	53.12	53.21	**53.24**
SPOT5	18.98	19.21	19.42	19.51	19.57	**19.59**
VEGETATION	24.12	24.35	24.48	**24.52**	24.49	24.47

Table C-30: Combined SNR Results in dB at 0.5 Bits/Sample for Different Values of *F*

	***F* = 1**	***F* = 2**	***F* = 4**	***F* = 8**	***F* = 16**	***F* = 32**
AIRS	36.80	48.06	51.28	52.95	54.13	**55.22**
AVIRIS	45.73	46.48	46.91	47.19	47.36	**47.47**
CASI	34.69	35.31	35.68	35.92	36.08	**36.21**
CRISM	24.90	29.05	29.57	29.84	30.00	**30.09**
Hyperion	71.53	72.94	73.74	74.32	74.68	**74.91**
IASI	20.93	39.91	42.29	43.28	44.16	**44.91**
Landsat	26.55	27.10	27.57	27.94	28.28	**28.48**
M3	39.64	40.46	40.84	41.03	41.14	**41.20**
MODIS	29.68	29.70	29.86	30.02	30.12	**30.14**
MSG	35.16	35.20	35.29	35.38	35.58	**35.75**
PLEIADES	22.37	23.11	23.47	23.67	23.79	**23.85**
SFSI	55.06	55.32	55.44	55.50	55.54	**55.56**
SPOT5	21.83	22.03	22.23	22.31	22.37	**22.38**
VEGETATION	29.24	29.35	29.40	29.41	29.41	**29.42**

Table C-31: Combined SNR Results in dB at 1 Bit/Sample for Different Values of *F*

	***F* = 1**	***F* = 2**	***F* = 4**	***F* = 8**	***F* = 16**	***F* = 32**
AIRS	52.26	54.68	55.96	56.76	57.38	**58.01**
AVIRIS	51.24	51.67	51.97	52.20	52.34	**52.44**
CASI	39.57	39.93	40.13	40.28	40.37	**40.44**
CRISM	35.74	36.72	37.26	37.64	37.89	**38.05**
Hyperion	75.65	76.78	77.48	78.00	78.32	**78.53**
IASI	42.06	44.65	46.04	46.71	47.31	**47.85**
Landsat	30.58	31.16	31.61	31.96	32.23	**32.40**
M3	47.41	47.86	48.08	48.21	48.32	**48.39**
MODIS	35.08	35.12	35.23	35.40	35.58	**35.60**
MSG	40.64	40.62	40.67	40.75	40.91	**41.04**
PLEIADES	27.16	27.63	27.92	28.12	28.25	**28.29**
SFSI	57.71	57.87	57.94	57.98	58.00	**58.01**
SPOT5	25.69	25.89	26.06	26.14	26.20	**26.21**
VEGETATION	34.36	34.42	34.44	34.44	34.44	**34.46**

Table C-32: Combined SNR Results in dB at 2 Bits/Sample for Different Values of *F*

	***F* = 1**	***F* = 2**	***F* = 4**	***F* = 8**	***F* = 16**	***F* = 32**
AIRS	58.31	59.83	60.66	61.30	61.79	**62.23**
AVIRIS	56.44	56.78	57.02	57.23	57.35	**57.43**
CASI	45.18	45.46	45.63	45.75	45.83	**45.89**
CRISM	44.98	45.98	46.73	47.43	47.94	**48.28**
Hyperion	81.30	82.18	82.76	83.19	83.46	**83.61**
IASI	48.34	50.16	51.30	51.82	52.29	**52.72**
Landsat	35.93	36.39	36.78	37.09	37.34	**37.48**
M3	56.71	57.19	57.47	57.69	57.88	**58.02**
MODIS	43.33	43.50	43.57	43.78	44.04	**44.09**
MSG	47.45	47.43	47.49	47.59	47.76	**47.89**
PLEIADES	33.97	34.28	34.53	34.69	34.81	**34.83**
SFSI	62.15	62.30	62.38	62.41	62.43	**62.44**
SPOT5	31.63	31.84	32.01	32.10	32.18	**32.20**
VEGETATION	41.24	41.26	41.27	41.29	41.29	**41.32**

Table C-33: Combined SNR Results in dB at 3 Bits/Sample for Different Values of *F*

	***F* = 1**	***F* = 2**	***F* = 4**	***F* = 8**	***F* = 16**	***F* = 32**
AIRS	62.72	64.19	65.03	65.50	65.87	**66.34**
AVIRIS	61.35	61.65	61.86	62.05	62.16	**62.25**
CASI	50.57	50.84	51.00	51.11	51.19	**51.24**
CRISM	51.61	52.58	53.28	53.97	54.46	**54.75**
Hyperion	86.07	86.78	87.32	87.66	87.85	**87.99**
IASI	53.39	55.16	56.12	56.64	57.19	**57.70**
Landsat	40.17	40.53	40.88	41.13	41.31	**41.43**
M3	61.98	62.32	62.52	62.69	62.82	**62.93**
MODIS	50.22	50.38	50.48	50.68	50.90	**50.96**
MSG	52.36	52.36	52.40	52.46	52.57	**52.66**
PLEIADES	39.72	40.00	40.23	40.38	40.50	**40.51**
SFSI	67.39	67.58	67.67	67.72	67.74	**67.75**
SPOT5	36.95	37.20	37.41	37.51	37.60	**37.63**
VEGETATION	47.14	47.15	47.16	47.19	47.21	**47.24**

Table C-34: Combined SNR Results in dB at 4 Bits/Sample for Different Values of *F*

	***F* = 1**	***F* = 2**	***F* = 4**	***F* = 8**	***F* = 16**	***F* = 32**
AIRS	67.20	68.55	69.27	69.78	70.22	**70.52**
AVIRIS	66.84	67.18	67.37	67.53	67.63	**67.69**
CASI	56.11	56.45	56.64	56.75	56.82	**56.87**
CRISM	57.25	58.04	58.58	59.15	59.58	**59.81**
Hyperion	90.49	91.12	91.54	91.89	92.03	**92.21**
IASI	58.48	60.16	60.97	61.39	61.80	**62.21**
Landsat	44.46	44.81	45.09	45.31	45.49	**45.58**
M3	65.84	66.12	66.28	66.42	66.53	**66.60**
MODIS	56.60	56.71	56.78	56.97	57.19	**57.25**
MSG	56.33	56.32	56.35	56.42	56.51	**56.57**
PLEIADES	45.31	45.61	45.83	45.98	46.09	**46.10**
SFSI	72.24	72.42	72.51	72.55	72.57	**72.57**
SPOT5	42.04	42.24	42.41	42.50	42.59	**42.61**
VEGETATION	52.55	52.54	52.56	52.60	52.62	**52.66**

C6 PERFORMANCE RESULTS

Detailed compression performance results are provided for all test images using the CCSDS 122.1-B-1 encoder without any spectral transform as baseline.

Table C-35 shows lossy compression results in terms of ΔSNR (in dB), defined as $SNR_{algorithm} - SNR_{baseline}$. Positive ΔSNR results indicate better compression performance than the baseline.

Results for the Peak Absolute Error (PAE), defined as the maximum absolute distortion between the original and reconstructed images, are also provided. These are expressed in terms of ΔPAE, defined as $PAE_{algorithm} - PAE_{baseline}$. Negative ΔPAE results indicate better compression performance than the baseline.

The Recommended Standard is optimized for SNR and is not designed to provide competitive PAE values. Users requiring controlled maximum errors are encouraged to consider the CCSDS 123.0-B-2 Recommended Standard for near-lossless compression (reference [4]).

Lossless compressed data rates are provided in table C-36. Average results for each instrument type are also provided for ease of reference.

Table C-35: Lossy Compression Results for All Images[1]

Image	Rate	CCSDS + IWT		CCSDS + POT		CCSDS + AAT		JPEG 2000		JPEG-LS		
		ΔSNR	ΔPAE	ΔSNR	ΔPAE	ΔSNR	ΔPAE	ΔSNR	ΔPAE	Rate	ΔSNR	ΔPAE
airs_gran9	0.25	13.52	-345	17.99	-422	-	-	15.14	-371	0.29	-8.84	-254
	0.5	14.55	-171	17.53	-211	-	-	15.29	-193	0.53	-6.96	-126
	1	13.44	-98	15.51	-108	-	-	14.12	-108	0.85	-6.01	-71
	2	10.74	-46	12.35	-49	7.54	-31	11.77	-49	1.27	-7.70	-30
	3	9.50	-16	11.04	-18	6.92	-14	11.26	-18	2.51	-2.00	-17
	4	9.11	-12	10.15	-9	5.58	-9	11.16	-13	4.79	8.36	-15
f060925t01p00r12_sc00.c.img	0.25	13.15	-712	17.12	-841	25.18	-841	15.28	-825	0.39	-5.22	-817
	0.5	15.47	-390	19.60	-446	26.07	-378	17.14	-420	0.70	-2.85	-409
	1	16.84	-195	19.80	-221	24.48	-205	18.04	-200	1.13	-1.44	-208
	2	15.46	-68	17.43	-74	20.74	-74	16.60	-76	1.67	-3.17	-58
	3	13.60	-28	15.21	-31	18.29	-31	15.10	-31	3.07	1.66	-29
	4	11.59	-20	12.63	-21	16.83	-22	13.44	-21	5.36	10.45	-24
f060925t01p00r12_sc00.uncal	0.25	11.17	-3908	18.19	-5064	24.66	-5120	13.77	-4922	1.31	8.60	-5937
	0.5	13.97	-2433	21.00	-2543	27.11	-2285	15.56	-2440	1.88	10.96	-2963
	1	16.25	-1152	21.22	-1395	24.98	-1268	17.20	-1294	2.56	12.04	-1522
	2	15.37	-401	18.22	-447	20.86	-414	16.00	-444	3.34	9.97	-490
	3	13.58	-149	16.37	-200	18.90	-209	14.47	-196	5.12	14.65	-222
	4	12.53	-82	15.50	-91	17.99	-95	13.51	-88	7.58	23.73	-104
f011020t01p03r05_sc01.uncal	0.25	9.53	-35	15.91	-85	19.38	-86	11.64	-87	0.00	-21.66	117
	0.5	11.37	-26	15.37	-58	17.60	-45	12.69	-51	0.03	-17.96	53
	1	11.04	-27	12.96	-27	14.80	-29	12.03	-29	0.17	-15.86	25
	2	7.69	-12	8.91	-13	11.04	-14	9.15	-13	0.46	-18.05	15
	3	5.28	-5	6.30	-5	10.26	-6	8.44	-6	1.29	-13.11	0
	4	5.21	-2	5.74	-1	16.52	-3	18.92	-3	3.06	-2.73	-3
t0477f06-nuc	0.25	10.65	-1028	13.12	-2641	18.81	-2483	13.54	-2811	0.48	-1.34	-3069
	0.5	12.49	-957	14.78	-1082	20.05	-1211	14.58	-1188	0.80	-0.41	-1232
	1	13.94	-379	15.03	-444	18.30	-479	15.30	-454	1.22	-0.17	-453
	2	12.14	-146	12.67	-155	14.73	-153	13.31	-154	1.76	-2.58	-147
	3	10.18	-43	10.76	-44	12.68	-48	11.59	-47	3.19	2.15	-50
	4	9.24	-25	9.58	-25	11.91	-27	10.95	-26	5.51	11.63	-31
t0477f06-rad-rss	0.25	11.34	-20331	13.59	-20175	18.64	-22553	13.96	-22295	1.82	14.69	-24245
	0.5	12.35	-5247	13.74	-5483	17.96	-6517	14.16	-6362	2.54	15.97	-7326
	1	12.37	-1635	12.71	-1574	16.00	-1671	13.62	-1812	3.41	16.63	-2354
	2	11.03	-487	10.93	-396	13.90	-335	12.31	-584	4.35	15.32	-750
	3	10.85	-309	10.72	-349	13.30	-363	12.17	-353	6.27	21.41	-450
	4	10.74	-105	10.87	-117	13.43	-127	12.34	-116	8.77	31.47	-167
t0477f06-raw	0.25	11.27	303	13.23	-1325	18.65	-931	14.25	-1594	0.49	-0.87	-1734
	0.5	12.73	-421	14.37	-593	19.43	-599	14.89	-659	0.80	-0.03	-695
	1	13.98	-135	14.49	-182	17.96	-226	15.37	-216	1.21	0.10	-213
	2	12.10	-104	12.15	-113	14.38	-108	13.36	-112	1.75	-2.40	-104
	3	10.23	-32	10.30	-31	12.42	-35	11.70	-34	3.18	2.31	-38
	4	9.32	-18	9.29	-19	11.73	-20	11.08	-19	5.50	11.81	-24
frt00009326_07_sc167-nuc	0.25	3.19	-369	3.64	-173	1.26	1922	6.86	-968	0.15	-12.73	-1076
	0.5	3.97	86	4.31	84	0.29	824	6.83	-29	0.33	-10.72	-55
	1	4.49	146	4.95	8	0.06	309	7.22	-25	0.65	-9.62	-17
	2	4.70	-12	5.07	-11	-0.22	204	7.07	-13	1.06	-10.49	-5
	3	3.95	-8	4.39	-8	-0.09	46	6.59	-8	2.07	-5.05	-11

[1] Rate is provided in bits/sample and SNR results in dB.

Image	Rate	CCSDS + IWT		CCSDS + POT		CCSDS + AAT		JPEG 2000		JPEG-LS		
		ΔSNR	ΔPAE	ΔSNR	ΔPAE	ΔSNR	ΔPAE	ΔSNR	ΔPAE	Rate	ΔSNR	ΔPAE
	4	3.37	-4	3.51	-2	-0.12	14	6.24	-5	4.15	4.83	-9
msp00004125_05_sc214	0.25	0.92	-125	1.96	-104	4.11	-42	12.90	-292	0.31	-10.96	-247
	0.5	-0.44	12	1.33	-29	2.84	110	11.87	-109	0.46	-10.68	-93
	1	1.48	10	3.09	1	5.24	161	12.28	-34	0.68	-11.05	-3
	2	4.75	-9	6.12	-15	7.51	-4	12.26	-21	0.99	-12.82	0
	3	5.55	-8	6.26	-10	7.10	-5	10.86	-13	1.57	-7.70	-8
	4	4.14	-3	4.40	-3	5.75	-3	12.30	-5	3.20	2.15	-6
frt00009326_07_sc167vnir	0.25	4.41	28	5.28	21	6.00	56	9.36	-79	0.09	-14.38	40
	0.5	4.25	-55	5.32	-40	6.16	19	12.61	-77	0.20	-10.17	-8
	1	5.72	-33	7.19	-32	8.10	-12	14.57	-56	0.54	-9.10	-10
	2	7.86	-22	8.74	-23	9.11	-14	12.68	-26	1.05	-11.09	-2
	3	5.83	-11	6.46	-11	6.86	-10	10.79	-14	1.88	-6.64	-10
	4	4.26	-3	4.78	-3	5.62	-4	11.34	-5	3.70	2.99	-6
hrl0000648f_07_sc182	0.25	-0.06	-50	0.40	-152	0.96	99	12.44	-841	0.52	-1.70	-1547
	0.5	-0.88	359	-0.03	301	0.55	1004	12.15	-367	0.78	-1.85	-537
	1	1.38	16	2.39	-34	4.36	944	12.70	-174	1.06	-2.17	-215
	2	3.89	-10	4.32	-9	3.55	233	11.52	-45	1.38	-5.33	-40
	3	4.10	-15	4.32	-15	2.75	117	10.11	-20	2.24	-0.75	-27
	4	3.40	-7	3.67	-8	2.58	25	9.39	-9	4.28	9.01	-16
msp00004125_05_sc214-nuc	0.25	5.96	-59	8.62	-29	10.92	-39	10.71	-244	0.19	-19.66	-100
	0.5	5.56	31	8.03	-32	9.35	29	9.51	-61	0.24	-17.34	19
	1	6.43	-21	8.79	-35	9.73	-9	10.20	-29	0.45	-16.09	12
	2	7.28	-8	8.84	-12	9.80	-5	10.48	-11	0.74	-16.29	14
	3	6.03	-5	6.95	-5	8.18	-6	9.51	-6	1.45	-10.64	-1
	4	4.67	-3	5.41	-2	8.97	-3	14.48	-4	3.10	-0.71	-4
hrl0000648f_07_sc182-nuc	0.25	5.44	-684	6.34	-740	5.65	16	9.91	-1457	0.21	-10.67	-1499
	0.5	5.96	-194	6.68	-212	4.15	806	9.62	-250	0.41	-8.88	-279
	1	5.98	-49	6.66	-73	3.01	254	9.44	-89	0.72	-8.50	-74
	2	5.90	-13	6.34	-13	2.33	125	8.66	-14	1.15	-10.04	-5
	3	5.12	-6	5.56	-10	2.61	45	8.06	-9	2.12	-4.63	-11
	4	4.29	-4	4.68	-4	2.29	6	7.46	-5	4.16	5.21	-8
frt00009326_07_sc167vnir-nuc	0.25	9.90	-301	12.36	-440	13.68	-335	13.25	-458	0.08	-18.39	-277
	0.5	10.87	-69	12.20	-85	12.94	-54	13.18	-81	0.17	-14.88	19
	1	9.85	-60	10.64	-66	11.03	-52	11.87	-63	0.39	-13.53	-14
	2	7.39	-10	8.24	-11	8.40	-9	9.91	-13	0.73	-14.81	14
	3	6.04	-4	6.56	-4	7.03	-5	8.98	-6	1.67	-9.68	-1
	4	4.95	-2	5.10	-2	7.33	-3	11.92	-3	3.56	0.40	-4
frt00009326_07_sc167ir-nuc	0.25	2.46	-57	2.92	-100	0.97	1611	6.36	-1070	0.17	-11.65	-1243
	0.5	3.25	-92	3.59	-26	-0.27	2124	6.20	-120	0.37	-9.80	-193
	1	3.83	-50	4.08	-57	-1.20	575	6.44	-58	0.72	-8.76	-73
	2	3.54	-5	3.94	-8	-1.73	136	6.19	-11	1.14	-9.63	-5
	3	2.95	-4	3.34	-7	-2.18	62	5.64	-10	2.17	-4.06	-12
	4	2.54	-1	2.94	-2	-2.14	13	5.48	-4	4.29	5.89	-9
frt00009326_07_sc167	0.25	-2.13	573	-1.65	231	-2.95	732	10.91	-1596	0.48	-1.60	-2564
	0.5	-2.65	763	-1.53	491	-5.17	2732	10.14	-249	0.74	-2.21	-486
	1	-0.28	64	0.77	137	-2.91	1647	10.71	-154	1.05	-2.79	-205
	2	2.62	7	2.98	-1	-0.91	327	10.33	-42	1.35	-5.24	-39
	3	2.88	-7	3.02	-7	-1.91	131	8.95	-19	2.21	-0.55	-26
	4	2.35	-4	2.55	-5	-1.86	56	8.35	-10	4.28	9.24	-16
frt00009326_07_sc167ir	0.25	-2.46	897	-2.15	766	-2.75	431	10.67	-1553	0.58	0.13	-2825
	0.5	-3.25	669	-2.33	603	-5.45	2242	10.09	-569	0.87	-0.47	-1009
	1	-1.71	277	-0.62	263	-4.88	2332	10.82	-123	1.18	-1.02	-213

Image	Rate	CCSDS + IWT		CCSDS + POT		CCSDS + AAT		JPEG 2000		JPEG-LS		
		ΔSNR	ΔPAE	ΔSNR	ΔPAE	ΔSNR	ΔPAE	ΔSNR	ΔPAE	Rate	ΔSNR	ΔPAE
	2	1.21	21	1.69	29	-2.93	588	9.60	-37	1.42	-4.01	-43
	3	2.24	-3	2.18	-4	-3.74	165	8.26	-22	2.29	0.86	-31
	4	1.66	0	1.71	-1	-4.20	87	7.73	-9	4.42	10.70	-16
Geo_Sample_flatfielded	0.25	6.70	-108	7.60	-8	7.91	-84	8.45	-8	0.02	-19.99	-39
	0.5	5.24	18	6.34	0	6.65	-28	7.00	-124	0.07	-16.78	-22
	1	3.88	4	5.33	0	5.66	12	5.80	-54	0.18	-15.08	-11
	2	2.93	3	4.32	0	4.74	-3	5.21	12	0.40	-14.57	6
	3	2.41	-1	3.29	0	3.98	-1	5.01	5	1.21	-9.57	-3
	4	1.84	2	2.50	0	4.02	1	5.71	1	3.26	-0.43	-6
Geo_Sample-nuc	0.25	7.21	-162	8.05	-173	9.49	-113	9.00	-209	0.03	-19.52	-52
	0.5	5.89	20	6.64	-36	7.58	82	7.39	-44	0.10	-16.03	42
	1	4.61	2	5.63	-6	6.19	31	6.19	-11	0.27	-14.99	28
	2	3.69	-5	4.80	6	5.28	10	5.63	-4	0.58	-15.01	14
	3	3.14	0	3.87	-2	4.74	-6	5.43	-5	1.45	-9.79	-3
	4	2.34	-2	2.72	-3	4.08	-4	5.63	-4	3.40	0.05	-6
Geo_Sample	0.25	5.83	-67	5.95	-61	8.40	-71	9.05	-177	0.03	-15.70	-49
	0.5	5.07	92	5.74	87	7.94	60	9.22	-71	0.11	-12.67	-1
	1	4.61	25	5.36	-25	6.50	91	8.12	-41	0.33	-12.24	-2
	2	3.24	-6	4.36	-11	4.79	15	6.41	-13	0.70	-13.70	6
	3	2.80	-1	3.73	-6	4.46	-6	6.27	-7	1.53	-8.36	-5
	4	2.44	-1	2.91	-2	4.09	-3	6.29	-3	3.45	1.63	-5
Desert	0.25	10.35	-369	11.36	-392	-32.02	-609	11.04	-427	0.42	-12.12	-354
	0.5	9.95	-285	10.27	-293	-35.64	-409	10.46	-327	0.63	-10.11	-280
	1	8.81	-122	9.02	-105	-39.83	-184	9.60	-137	0.93	-7.93	-119
	2	8.15	-64	7.86	-54	-45.86	-90	9.38	-71	1.45	-8.06	-57
	3	8.14	-36	7.75	-27	-50.95	-50	9.86	-39	3.12	-1.87	-41
	4	7.49	-8	6.92	-4	-56.29	-16	9.79	-10	5.61	8.25	-14
mountain	0.25	0.40	-8	0.80	-14	2.22	-9	1.28	-6	0.00	-20.31	158
	0.5	0.29	7	1.17	-4	2.72	21	1.07	9	0.02	-21.69	90
	1	0.17	3	1.47	6	3.28	11	0.82	3	0.03	-20.97	47
	2	0.05	0	1.43	-1	2.92	-1	0.82	0	0.19	-19.57	23
	3	-0.28	1	0.66	0	2.00	0	0.74	1	0.93	-14.01	4
	4	-0.59	0	0.05	0	2.55	-1	1.59	0	2.67	-4.05	-2
coast	0.25	3.03	-6	2.54	-8	3.72	-11	4.14	-9	0.00	-31.37	226
	0.5	2.37	-4	2.21	-5	3.13	-6	3.21	0	0.06	-32.49	113
	1	1.21	-2	1.31	-3	2.06	-4	2.31	-3	0.03	-28.36	55
	2	0.01	0	0.27	-1	0.92	-1	2.05	-1	0.11	-27.05	28
	3	-1.17	0	-0.98	0	0.87	0	4.68	0	0.44	-20.24	6
	4	-7.16	0	-6.96	1	1.15	0	11.88	0	1.53	-17.51	0
agriculture	0.25	1.55	-26	2.28	1	3.24	-4	2.68	-24	0.00	-22.09	146
	0.5	1.63	10	2.75	-2	3.84	13	2.58	9	0.04	-24.29	94
	1	1.76	2	2.92	-2	4.37	2	2.71	7	0.03	-20.77	47
	2	1.57	0	2.71	-1	4.05	-2	2.80	-2	0.27	-20.30	23
	3	0.99	0	1.65	0	2.96	0	2.43	0	1.02	-14.52	4
	4	0.40	-1	0.83	-1	4.43	-1	4.50	-1	2.53	-4.32	-2
M3targetB-nuc	0.25	5.14	-115	5.85	-124	5.58	114	9.71	-193	0.06	-23.00	-5
	0.5	4.89	-23	5.33	-32	3.93	100	8.46	-33	0.12	-21.10	62
	1	4.14	-9	4.56	-11	2.82	50	7.44	-13	0.28	-20.32	38
	2	2.70	-2	2.96	-3	1.75	8	5.81	-4	0.58	-21.39	23
	3	1.61	-1	1.66	-1	1.16	1	5.97	-3	1.27	-15.46	3
	4	0.90	0	0.55	2	1.32	1	15.95	-1	2.44	-5.17	-1
M3target2	0.25	2.06	210	2.51	231	-1.07	763	15.74	-580	0.18	-9.77	-577

Image	Rate	CCSDS + IWT		CCSDS + POT		CCSDS + AAT		JPEG 2000		JPEG-LS		
		ΔSNR	ΔPAE	ΔSNR	ΔPAE	ΔSNR	ΔPAE	ΔSNR	ΔPAE	Rate	ΔSNR	ΔPAE
	0.5	2.59	11	2.90	59	-2.91	476	14.69	-80	0.29	-10.43	-37
	1	5.19	-10	5.60	-19	0.38	153	16.26	-49	0.55	-9.51	-10
	2	4.33	-5	4.44	-7	1.25	40	10.31	-14	1.05	-15.74	11
	3	2.43	-2	2.52	-2	0.39	11	7.79	-5	1.33	-11.58	0
	4	1.56	0	1.46	0	0.04	2	12.28	-2	2.50	-0.83	-3
M3globalA-nuc	0.25	9.27	3	10.59	-54	12.38	48	14.04	-112	0.11	-20.32	94
	0.5	10.15	-43	10.94	-46	11.91	-26	14.08	-47	0.18	-17.32	52
	1	10.94	-23	11.42	-23	12.13	-13	14.81	-25	0.41	-16.06	27
	2	11.36	-14	11.32	-14	11.88	-11	14.65	-15	0.76	-16.23	14
	3	9.52	-7	9.30	-6	10.91	-6	17.48	-7	1.49	-10.62	-1
	4	9.04	-2	8.03	-2	14.43	-3	33.96	-3	3.30	-0.63	-3
M3globalA	0.25	11.40	-1817	12.03	-1827	11.21	-1809	24.45	-2181	0.25	-0.68	-2077
	0.5	4.10	-74	5.29	-28	1.07	102	16.64	-143	0.36	-8.96	-82
	1	4.32	-26	5.40	-22	1.16	71	16.92	-45	0.58	-10.62	-7
	2	8.95	-22	9.17	-21	6.79	-2	16.73	-29	0.93	-12.18	-2
	3	9.55	-11	9.10	-12	8.02	-7	16.36	-14	1.59	-7.30	-8
	4	7.44	-6	7.01	-6	6.71	-5	27.04	-7	3.38	2.62	-7
MOD01.A2001222 .1200day-nuc	0.25	1.41	-264	2.86	-540	4.25	-866	3.46	-831	0.57	1.36	-1911
	0.5	0.83	-254	2.43	-349	4.05	-508	2.52	-166	0.84	2.88	-1152
	1	0.25	75	1.98	-150	3.32	-94	1.96	-20	1.20	2.69	-522
	2	-0.51	25	1.26	19	2.18	138	1.43	18	1.65	-0.51	-143
	3	-0.59	4	1.23	-24	1.58	-10	1.46	11	2.73	3.19	-63
	4	-1.06	4	0.78	-7	0.80	-6	0.97	-3	4.35	11.31	-30
MOD01.A2001222 .1200day	0.25	1.45	-248	2.92	-514	4.36	-830	3.96	-793	0.57	1.56	-1878
	0.5	0.79	-214	2.42	-377	4.05	-544	2.98	-137	0.84	2.20	-1152
	1	0.03	29	1.77	-133	3.05	-72	2.24	-45	1.18	2.65	-521
	2	-0.90	35	0.81	29	1.46	10	1.36	4	1.61	-1.54	-127
	3	-1.51	1	0.14	-22	0.21	-29	0.89	7	2.60	1.23	-63
	4	-1.66	12	-0.06	0	-0.46	0	0.58	8	4.07	9.94	-20
MOD01.A2001222 .1200night	0.25	-1.23	466	-0.67	104	0.76	443	4.69	-71	0.09	-10.23	-249
	0.5	-1.70	24	-0.82	117	1.29	186	5.51	-156	0.20	-7.93	-241
	1	-2.12	-2	-0.42	-54	2.17	40	5.66	-114	0.41	-8.02	-129
	2	-2.68	21	-0.18	12	1.92	-13	4.05	-12	0.76	-11.72	-4
	3	-2.12	5	-0.07	-1	1.30	-4	3.42	-5	1.66	-7.01	-8
	4	-2.17	2	-0.49	0	0.79	-2	3.01	-3	3.37	2.82	-7
MOD01.A2001222 .1200_250m-nuc	0.25	0.84	-411	2.03	-495	1.78	-453	2.89	-729	0.42	-5.16	-1243
	0.5	0.96	22	2.56	-183	2.45	-111	2.51	-304	0.69	-3.22	-608
	1	1.21	6	3.75	-215	3.87	-183	2.74	-240	1.04	-2.15	-389
	2	1.39	-9	5.50	-57	5.86	-77	3.36	-73	1.49	-3.62	-112
	3	1.43	-26	5.73	-37	6.25	-38	3.67	-32	2.81	1.32	-61
	4	1.32	1	5.23	-8	5.78	-6	3.75	-6	5.08	10.41	-21
MOD01.A2001222 .1200_250m	0.25	0.86	-257	2.04	-460	1.81	-272	2.95	-733	0.42	-5.27	-1247
	0.5	0.96	55	2.53	-163	2.43	-89	2.51	-231	0.69	-3.39	-578
	1	1.21	2	3.52	-226	3.66	-190	2.68	-241	1.04	-2.35	-399
	2	1.23	-45	4.44	-61	4.64	-81	3.08	-70	1.49	-3.77	-117
	3	1.15	-20	4.63	-38	4.88	-34	3.48	-39	2.79	1.26	-61
	4	1.04	2	4.34	-5	4.68	-6	3.63	-5	5.05	10.28	-20
MOD01.A2001222.1200_500m	0.25	1.34	-381	1.12	165	1.46	-177	6.09	-635	0.79	-0.65	-1469
	0.5	2.51	-777	1.42	-716	1.27	-706	7.11	-1209	1.11	2.94	-1591
	1	2.15	-326	2.38	-386	1.72	-236	8.38	-588	1.49	4.46	-727
	2	1.49	-56	3.07	-107	2.38	-110	8.87	-206	2.03	2.65	-239
	3	1.00	10	3.32	-13	2.43	1	9.46	-47	3.49	7.83	-73

Image	Rate	CCSDS + IWT		CCSDS + POT		CCSDS + AAT		JPEG 2000		JPEG-LS		
		ΔSNR	ΔPAE	ΔSNR	ΔPAE	ΔSNR	ΔPAE	ΔSNR	ΔPAE	Rate	ΔSNR	ΔPAE
	4	0.58	4	3.25	-12	2.31	-8	9.36	-28	5.84	17.07	-40
MOD01.A2001222.1200_500m-nuc	0.25	1.52	-145	1.28	84	1.58	-44	6.15	-597	0.79	-0.60	-1499
	0.5	2.78	-128	1.51	63	1.28	-87	7.35	-511	1.11	2.97	-867
	1	2.58	-340	2.52	-393	1.88	-239	8.84	-595	1.50	4.56	-727
	2	2.18	-27	3.32	-110	3.03	-111	9.20	-200	2.04	2.74	-232
	3	1.70	24	3.77	-31	3.15	-24	9.71	-73	3.53	7.99	-92
	4	1.43	3	3.86	-10	3.13	-8	9.69	-28	5.87	17.38	-40
MOD01.A2001222.1200night-nuc	0.25	-1.08	95	0.09	133	2.04	373	1.39	-18	0.08	-15.93	-109
	0.5	-1.91	5	-0.34	-1	1.74	81	0.84	-77	0.17	-14.12	-115
	1	-2.52	2	0.00	40	1.95	-19	1.05	-34	0.33	-13.09	-38
	2	-2.57	6	0.04	-4	1.69	-6	1.14	-3	0.57	-14.96	7
	3	-2.45	3	-0.16	0	1.20	-3	1.10	-2	1.45	-9.72	-3
	4	-2.26	3	-0.57	1	1.04	-1	1.07	0	3.30	0.25	-4
RC15_Level15	0.25	-0.18	33	1.50	-5	3.26	31	2.28	-59	0.04	-21.09	43
	0.5	-0.17	18	1.42	0	3.12	10	1.71	-12	0.11	-20.35	51
	1	-0.40	2	1.20	-10	2.55	-3	1.59	-10	0.25	-21.01	26
	2	-1.57	1	0.14	-1	0.64	-2	1.19	-1	0.44	-22.24	19
	3	-2.03	0	-0.69	-1	-0.04	0	0.89	-2	0.98	-16.91	1
	4	-3.14	1	-2.05	0	0.02	0	2.24	-1	2.38	-6.86	-2
montpellier_crop_Deca-020	0.25	1.08	-256	0.99	-288	1.98	-337	2.81	-388	0.46	-4.09	-825
	0.5	1.00	-53	0.83	-109	1.75	-100	2.14	-150	0.79	-2.14	-443
	1	0.95	-12	0.79	19	2.04	7	1.95	-47	1.18	-0.91	-217
	2	0.92	4	0.73	-18	2.75	-24	2.18	-46	1.76	-1.85	-107
	3	0.88	-1	0.74	-9	2.96	-22	2.40	-25	3.24	3.83	-56
	4	0.94	2	0.70	-6	2.80	-11	2.30	-14	5.66	13.75	-30
montpellier_crop_Deca-050	0.25	0.85	-204	0.77	-389	1.65	-277	2.39	-479	0.46	-4.08	-849
	0.5	0.44	-13	0.29	-6	0.95	-9	1.49	-112	0.79	-2.05	-418
	1	0.10	43	-0.06	35	0.75	17	1.12	-47	1.18	-0.90	-211
	2	-0.32	5	-0.47	-5	0.79	-16	0.95	-41	1.76	-1.86	-107
	3	-0.21	6	-0.52	-6	0.87	-5	1.12	-20	3.24	3.83	-59
	4	-0.23	8	-0.44	2	0.78	-5	1.23	-10	5.66	13.75	-31
montpellier_crop_Deca-035	0.25	0.97	-366	0.89	-330	1.84	-392	2.60	-469	0.46	-4.13	-859
	0.5	0.71	-159	0.56	-182	1.35	-150	1.80	-261	0.79	-2.11	-565
	1	0.46	15	0.29	36	1.29	9	1.45	-47	1.18	-0.90	-230
	2	0.21	15	-0.01	1	1.63	-9	1.42	-29	1.76	-1.86	-94
	3	0.23	2	-0.06	-7	1.74	-3	1.61	-22	3.24	3.83	-56
	4	0.34	10	0.01	-6	1.62	-7	1.63	-13	5.66	13.75	-32
montpellier-t	0.25	1.34	-730	1.12	-675	2.35	-697	3.30	-891	0.71	0.14	-1631
	0.5	1.59	-169	1.16	-201	2.42	-237	2.96	-367	1.06	1.98	-891
	1	1.95	50	1.61	-59	3.28	-72	3.22	-186	1.58	3.21	-433
	2	2.53	-97	2.17	-74	4.75	-108	3.99	-142	2.27	2.48	-210
	3	2.59	-49	2.38	-42	4.85	-69	4.24	-80	3.93	8.42	-119
	4	2.44	-23	2.14	-21	4.45	-32	4.11	-37	6.40	18.30	-62
SFSI_mantar-nuc	0.25	0.55	-2	0.76	-3	0.79	27	1.17	-11	0.05	-30.83	217
	0.5	0.21	7	0.56	4	0.61	4	1.14	-3	0.04	-20.36	101
	1	-0.10	7	0.47	11	0.29	3	1.44	0	0.09	-18.39	45
	2	-0.30	-1	0.62	-3	0.66	-4	2.24	-5	0.17	-15.52	16
	3	-0.38	1	0.60	-2	0.69	-2	2.62	-3	1.01	-11.39	-1
	4	-0.86	0	-0.25	-1	0.40	-1	2.35	-1	3.04	-2.17	-3
SFSI_mantar_Radrmnoise_	0.25	3.99	242	7.81	23	8.41	83	7.39	-72	0.13	-15.68	7
	0.5	5.49	7	9.44	20	9.52	-35	8.99	-56	0.26	-14.62	3
	1	6.78	10	9.91	26	10.10	-6	10.25	-24	0.49	-14.27	11

Image	Rate	CCSDS + IWT		CCSDS + POT		CCSDS + AAT		JPEG 2000		JPEG-LS		
		ΔSNR	ΔPAE	ΔSNR	ΔPAE	ΔSNR	ΔPAE	ΔSNR	ΔPAE	Rate	ΔSNR	ΔPAE
	2	7.53	-8	10.01	24	10.15	-10	10.61	-23	0.77	-16.35	2
	3	6.20	-17	8.75	-11	8.88	-17	9.58	-23	1.46	-11.37	-18
	4	4.69	-8	8.93	-9	9.65	-10	10.46	-10	3.02	-1.84	-11
SFSI_mantar_Raw	0.25	4.11	-75	4.36	-81	4.49	-80	5.07	-84	0.05	-25.02	137
	0.5	3.28	-9	3.69	-9	3.80	-10	4.49	-17	0.06	-19.09	86
	1	2.42	-7	3.01	-10	3.15	-9	4.11	-12	0.11	-14.01	32
	2	1.50	-2	2.54	-5	2.30	-4	4.29	-7	0.26	-12.89	13
	3	1.44	0	2.42	-2	2.22	-1	4.60	-3	1.43	-9.26	-1
	4	1.15	0	1.66	-1	2.36	-1	4.28	-2	3.29	0.50	-4
toulouse_spot5_xs_extract1	0.25	-0.07	-13	0.93	-12	0.93	-17	0.70	-22	0.02	-20.71	112
	0.5	-0.32	0	1.32	-3	1.24	-11	0.40	2	0.26	-17.77	35
	1	-0.56	-11	1.91	-17	1.86	-27	0.13	-8	0.18	-12.13	0
	2	-1.06	8	2.24	0	2.14	-2	0.04	10	0.63	-12.96	13
	3	-1.02	5	2.22	-1	2.16	-1	0.15	4	1.73	-7.82	-2
	4	-0.97	1	1.74	0	1.94	-1	0.23	1	3.87	2.12	-4
vgt1_1b	0.25	0.19	-174	-0.04	-144	-0.61	-101	4.60	-349	0.16	-12.36	-351
	0.5	0.68	-62	1.10	0	0.36	-20	3.78	-99	0.30	-10.80	-154
	1	0.85	-18	1.51	-45	0.85	-15	3.55	-68	0.52	-9.99	-74
	2	0.43	3	1.33	0	0.48	3	3.19	-9	0.86	-11.62	-7
	3	0.01	1	1.20	-1	0.21	3	3.00	-5	1.85	-6.70	-11
	4	-0.32	1	0.93	-2	0.07	0	2.73	-3	3.84	2.81	-9
vgt1_1c	0.25	1.84	-96	2.54	-131	2.34	-143	4.34	-232	0.11	-13.76	-236
	0.5	1.81	-42	2.55	-99	2.41	-107	3.85	-122	0.25	-11.78	-146
	1	1.51	-32	2.44	-42	2.25	-53	3.59	-56	0.49	-10.88	-68
	2	0.72	9	2.18	-5	1.83	-3	3.21	-10	0.82	-12.36	-6
	3	0.32	1	1.94	-6	1.60	-2	2.96	-5	1.82	-7.44	-10
	4	0.05	0	1.54	-4	1.41	-1	2.65	-4	3.81	2.11	-9

Table C-36: Lossless Compression Results for All Images

	CCSDS 122.1 (No transform)	CCSDS 122.1 + IWT	CCSDS 122.1 + POT	JPEG-LS	JPEG2000
AIRS	6.923	5.025	**4.711**	6.365	4.754
airs_gran9	6.923	5.025	**4.711**	6.365	4.754
AVIRIS	7.079	5.183	**4.875**	6.888	5.039
f011020t01p03r05_sc01.uncal	4.789	3.525	3.316	4.579	**3.257**
f060925t01p00r12_sc00.c.img	7.122	4.874	**4.623**	6.922	4.729
f060925t01p00r12_sc00.uncal	9.325	7.150	**6.686**	9.163	7.132
CASI	8.465	6.658	6.688	8.168	**6.523**
t0477f06-nuc	7.378	5.606	5.579	7.086	**5.430**
t0477f06-rad-rss	10.613	8.744	8.812	10.348	**8.698**
t0477f06-raw	7.404	5.625	5.674	7.071	**5.440**
CRISM	6.167	5.408	5.335	5.450	**4.640**
frt00009326_07_sc167-nuc	6.122	5.428	5.373	5.715	**4.976**
frt00009326_07_sc167	6.919	6.501	6.488	5.837	**5.396**
frt00009326_07_sc167ir-nuc	6.306	5.747	5.694	5.870	**5.284**
frt00009326_07_sc167ir	7.170	6.891	6.892	5.988	**5.732**
frt00009326_07_sc167vnir-nuc	5.368	4.120	4.023	5.079	**3.716**
frt00009326_07_sc167vnir	5.889	4.889	4.808	5.219	**4.009**
hrl0000648f_07_sc182-nuc	6.224	5.336	5.263	5.716	**4.811**
hrl0000648f_07_sc182	6.908	6.300	6.262	5.830	**5.192**
msp00004125_05_sc214-nuc	5.101	4.010	3.792	4.566	**3.474**
msp00004125_05_sc214	5.658	4.861	4.757	4.679	**3.805**
Hyperion	5.281	4.661	4.558	4.910	**4.301**
Geo_Sample-nuc	5.274	4.617	4.493	4.943	**4.287**
Geo_Sample	5.559	4.960	4.872	4.996	**4.424**
Geo_Sample_flatfielded	5.009	4.406	4.309	4.792	**4.191**
IASI	7.119	5.532	5.677	7.122	**5.128**
Desert	7.119	5.532	5.677	7.122	**5.128**
Landsat	4.042	3.881	3.906	3.695	**3.668**
agriculture	4.399	4.085	4.093	3.990	**3.919**
coast	3.260	3.173	3.284	2.932	**2.864**
mountain	4.467	4.384	4.340	**4.163**	4.221
M3	5.074	4.039	4.031	4.385	**3.192**
M3globalA-nuc	5.110	3.363	3.343	4.815	**2.797**
M3globalA	5.767	4.324	4.321	4.889	**3.152**
M3targetB-nuc	4.268	3.807	3.780	3.887	**3.259**

	CCSDS 122.1 (No transform)	CCSDS 122.1 + IWT	CCSDS 122.1 + POT	JPEG-LS	JPEG2000
M3targetB	5.152	4.661	4.678	3.949	**3.559**
MODIS	6.668	6.827	6.480	**6.027**	6.258
MOD01.A2001222.1200_250m-nuc	6.924	6.715	**6.106**	6.627	6.545
MOD01.A2001222.1200_250m	6.891	6.726	**6.222**	6.610	6.535
MOD01.A2001222.1200_500m-nuc	8.177	7.941	7.567	7.427	**6.774**
MOD01.A2001222.1200_500m	8.127	8.041	7.619	7.394	**6.781**
MOD01.A2001222.1200day-nuc	6.127	6.752	6.585	**5.416**	6.490
MOD01.A2001222.1200day	5.922	6.626	6.489	**5.031**	6.404
MOD01.A2001222.1200night-nuc	5.323	5.667	5.352	**4.822**	5.196
MOD01.A2001222.1200night	5.856	6.144	5.901	**4.889**	5.339
MSG	3.899	4.178	4.068	**3.789**	3.959
RC15_Level15	3.899	4.178	4.068	**3.789**	3.959
PLEIADES	7.664	7.476	7.612	**7.431**	7.516
montpellier-t	8.215	**7.783**	7.912	7.986	7.841
montpellier_crop_Deca-020	7.480	7.277	7.409	**7.246**	7.309
montpellier_crop_Deca-035	7.480	7.387	7.525	**7.246**	7.421
montpellier_crop_Deca-050	7.479	7.457	7.600	**7.246**	7.494
SFSI	5.020	4.591	4.462	4.558	**4.148**
SFSI_mantar-nuc	4.847	4.846	4.757	4.569	**4.499**
SFSI_mantar_Rad_rmnoise	4.898	3.935	3.735	4.348	**3.341**
SFSI_mantar_Raw	5.314	4.991	4.893	4.756	**4.605**
SPOT5	5.656	5.716	**5.331**	5.443	5.678
toulouse_spot5_xs_extract1	5.656	5.716	**5.331**	5.443	5.678
Vegetation	5.516	5.436	5.453	5.335	**5.319**
vgt1_1b	5.578	5.536	5.557	**5.349**	5.359
vgt1_1c	5.454	5.335	5.297	5.320	**5.279**

ANNEX D

SELECTION CRITERIA

D1 INTRODUCTION

In addition to the mandatory requirements, the following selection criteria were considered in the evaluation of the proposed compressors:

a) implementation complexity;

b) compression effectiveness;

c) flexibility;

d) error containment;

e) robustness;

f) user-friendliness;

g) license considerations;

h) extension of an existing standard.

Each of these criteria is described in further detail below.

D2 IMPLEMENTATION COMPLEXITY

The implementation complexity considerations were

– throughput of hardware implementation;

– algorithmic complexity (number of arithmetic operations per sample or other relevant metric); and

– memory/buffering requirements.

D3 COMPRESSION EFFECTIVENESS

Compression effectiveness refers to the ability of a compressor to provide high-fidelity (or lossless) reconstructed data using a small compressed size, without regard to other concerns such as speed or memory use.

For compressors offering flexibility in adjusting compression parameters, compression effectiveness results that indicate performance in a variety of configurations should be provided, for example, scan-based vs. frame-based, small vs. large error-containment segment size, and default vs. custom decorrelating transform.

Lossless compression effectiveness is measured by the bit rate required to achieve lossless compression on each test data set.

Lossy compression effectiveness is measured by the rate-distortion performance on the test data sets. Lossy compression results were produced at or near compressed bit rates 0.1, 0.25, 0.5, 1, 2, and 4 bits/sample, and at or near compression ratios of 10:1 and 20:1.

At the bit rates examined, SNR (with signal power) was the main metric considered. PAE results were also considered, although with lower priority than SNR.

D4 FLEXIBILITY

Desirable flexibility in the compressor included features such as the following:

- Flexibility in the ability to control the tradeoff between rate and distortion, for example, bit-accurate rate control, fine granularity of quality levels, and the ability to specify rate and/or quality;
- the ability to provide both lossless and lossy compression. For lossy compressors, the ability to provide effective compression over a wide range of bit rates or fidelity was desirable;
- the ability to tune the algorithm post-launch to improve performance, for example, uploading new compression parameters;
- flexibility in the structure of the compressed bit stream, for example, progressive transmission, the ability to provide higher reconstructed fidelity, or downlink priority in certain spectral bands or regions-of-interest.

D5 ERROR CONTAINMENT

The effectiveness of the error-containment mechanism to deal with packet losses on the communications channel was considered. Considerations include the ability to adjust the size of error-containment segments and the impact of using small error-containment segments on compression effectiveness. In certain applications, error concealment following data loss or corruption was also considered.

D6 ROBUSTNESS

It was considered desirable for compressors not to be dramatically affected by instrument defects, including the following:

- detector defects such as dead pixels, frozen pixels, and spikes;
- instrument defects such as spectral and spatial misregistration, keystone, and smile.

D7 USER-FRIENDLINESS

A desirable feature of compressors was that the algorithm achieves good performance without significant expertise on the part of the user implementing the algorithm. Ideally, algorithms should be operated as a black box; that is, no tuning of coding tables would be required from user.

D8 LICENSE CONSIDERATIONS

In addition to the mandatory requirement that CCSDS member agencies can use a recommended compressor without payment of a license fee, other license considerations were relevant. For example, the ability of a third party to use a recommended algorithm for a commercial venture without paying a license fee was considered desirable because it provides further incentive for that third party to develop a hardware implementation.

D9 EXTENSION OF AN EXISTING STANDARD

Leveraging of an existing CCSDS Recommended Standard, that is, incorporating an existing standard or functional part thereof in the new standard, was a desirable feature.

ANNEX E

RATE-ALLOCATION WEIGHT DERIVATIONS

E1 INTRODUCTION

This annex provides the derivations of the spectral band weights to be used in the rate-allocation process when the IWT or the POT are used in the spectral transform stage. Annex subsection E2 provides derivations for the IWT, and E3 provides derivations for the POT.

E2 RATE-ALLOCATION WEIGHT DERIVATIONS FOR THE IWT

This subsection shows how to calculate the band weights used in the rate-allocation methods, described in section 4 of this document and in annex subsection F2 of reference [1], when the IWT is employed.

As described in subsection 4.4.3.1 of reference [1], each IWT decomposition level divides an input sequence into two sequences, L and H, also known as subbands. The input of the first level is the sequence X to be transformed, and the output subbands are denoted L_1 and H_1. The input of each subsequent level i, $i > 1$, is L_{i-1}, and the output is subbands L_i and H_i.

As described in subsection 4.4.4.1 of reference [1], forward calculation of the IWT on an input sequence produces six output sequences: L_5, H_5, H_4, H_3, H_2, and H_1. If $|\cdot|$ denotes the number of elements in a subband, then the first $|L_5|$ spectral bands have the same weight, which is associated with subband L_5. The next $|H_5|$ spectral bands have a weight associated with subband H_5, and so forth. In the case of $|L_i| = 1$, for some i, the weight associated with L_5 is that of L_i, and $|H_j| = 0$ for $j > i$.

If $N_{TZ} = 1$, the IWT has no effect, and the weight of the single transformed band produced, L_0, is 1.

For $N_{TZ} > 1$, the weight $weight_b$ of any subband b (where $b = L_l$ or $b = H_l$, $1 \le l \le 5$) can be calculated as follows:

a) A sequence of length 256 is created, and interpreted as the concatenated output of a 5 level IWT, as depicted in figure 4-2 of reference [1].

b) One of the central elements of subband b is set to 1. All other elements of the sequence are set to 0.

c) An inverse five-level IWT is applied to the sequence obtained in step b). The result is a sequence $\hat{X}$ of length 256.

d) The weight of $\hat{X}$ is calculated as

$$\text{weight}_b = \left\|\hat{X}\right\|^2 = \sum_{i=0}^{255} \hat{X}_i^2.$$

Weights for all subbands of a 5-level IWT are provided in table E-1. Weights for L_0, …, L_4 are also provided for the case $N_{TZ} < 32$, where less than 5 levels can be applied.

Table E-1: Subband Weights for the Five-Level IWT

Subband *b*	$weight_b$
H_1	0.71875
H_2	0.921875
H_3	1.5859375
H_4	3.04296875
H_5	6.021484375
L_0	1
L_1	1.5
L_2	2.75
L_3	5.375
L_4	10.6875
L_5	21.34375

E3 RATE-ALLOCATION WEIGHT DERIVATIONS FOR THE POT

E3.1 GENERAL

The calculation of weights for the POT transform is dependent on the number of levels of the transform applied and the structure resulting from the number of segments to which the POT is applied.

This annex first introduces how to calculate the weight of the output sequence for a single pairwise operation, and then explains how to extend the calculation for the complete multi-level structure of the POT. A more detailed explanation can be found in reference [8].

E3.2 CALCULATION FOR A SINGLE PAIRWISE OPERATION

In a balanced pairwise operation, both input sequences have the same weight 2^q, for some rational number q. The principal output of the pairwise operation has a weight of 2^{q+1}, while the detail output has a weight of 2^{q-1}, as depicted in figure E-1.

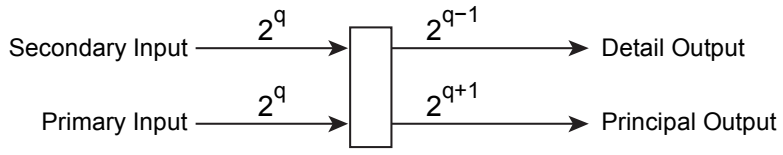

Figure E-1: Weight Calculation for a Balanced Pairwise Operation

For an unbalanced pairwise operation, the primary input was produced by a pairwise operation in the previous level, while the secondary output was not modified by a pairwise operation in that level. Therefore, the primary and secondary inputs have weights 2^{q-1} and 2^{q} respectively, for some rational number q. The principal and detail outputs of the unbalanced operation have weights 2^{q+1} and 2^{q-2}, respectively, as depicted in figure E-2.

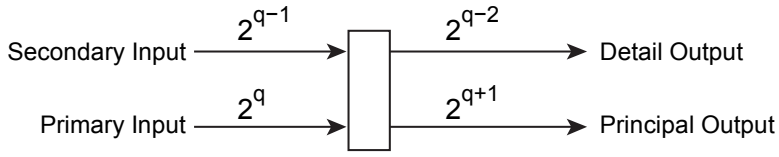

Figure E-2: Weight Calculation for an Unbalanced Pairwise Operation

E3.3 CALCULATION FOR THE MULTI-LEVEL STRUCTURE

The initial weights for the input sequences of the first level of the POT are 2^{0}. Using the expressions provided in E3.2 and E3.3, it is straightforward to calculate the weights of all output components of any multi-level structure.

An example depicting the weight evolution of the POT applied to five input sequences is provided in figure E-3.

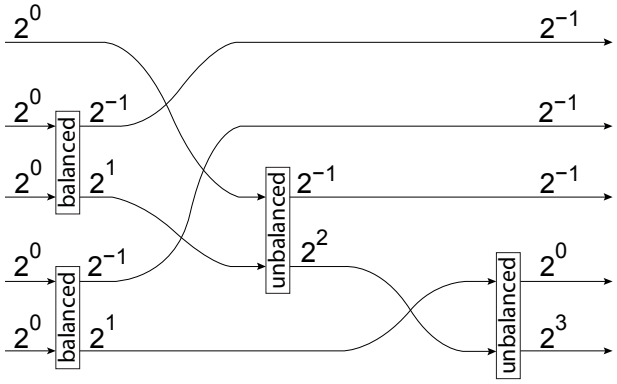

Figure E-3: Weight Calculation for Two POT Levels and Five Input Sequences

ANNEX F

ABBREVIATIONS AND ACRONYMS

AAT	Arbitrary Affine Transform
DWT	Discrete Wavelet Transform
IWT	Integer Wavelet Transform
KLT	Karhunen-Loève Transform
NLR	Near-Lossless Rate
PAE	Peak Absolute Error
POT	Pairwise Orthogonal Transform
RAS	Rate-Allocation Subset

ANNEX G

GLOSSARY

The following terms are used to help explain concepts used in this book. They are not part of the Recommended Standard.

Compressed data rate:	The average number of bits used to encode each sample of an image in its compressed representation.
Discontinuity artifact:	Artifact that may appear at the edge between reconstructed image regions when these regions are transformed using sufficiently different training parameters.
Embedded:	Property of compressed image streams that, when truncated at a given bit rate, allows near-optimal reconstruction of the image at that rate.
Fixed-rate allocation:	Strategies for the distribution of a rate budget so that the length of the compressed bitstream divided by the number of pixels in the original image is approximately constant and equal to the target bitrate.
Fixed-quality allocation:	Strategies to produce compressed bitstreams of variable length so that reconstructed images have approximately the same quality.
Rate-allocation subset:	Each of the subsets into which the image segments are divided for rate-allocation.
Spectral segment:	Set of segments associated with the same spatial area, that is, the *i*th segment of each spectral band.
Dynamic range expansion:	Increment of the range of possible values of an image after applying a transform.
Streaking artifact:	Discontinuities along the y direction due to nonuniformities in the sensor array used for registering an image.

ANNEX H

DEFAULT COMPRESSION PARAMETERS

The following user-defined parameters are employed to generate the experimental results presented in this document when applicable and unless explicitly stated otherwise.

When the IWT is employed for lossy compression, the band weights used for rate allocation are as described in annex subsection E2. When the POT is used, weights for rate allocation are as described in annex subsection E3. When the AAT is employed or no spectral transform is applied, weights for rate allocation are equal to 1 for all bands.

When the AAT is employed, matrix Q and vector V for each image are calculated to approximate the KLT, as described in 6.2.2. Therefore, all regions within a collection have the same Q and V values, and the *repeat* field (described in subsection 5.4.4.4 of reference [1]) is set to '1' to minimize the required side information.

Table H-1: Default Parameter Values

Description	Symbol	Default value (lossless)	Default value (lossy)
Upshift	U	0	2
Downshift	D	0	0
Number of segments from the same band in a collection	R	$\left\lceil \frac{N_X}{8} \right\rceil$	$\left\lceil \frac{N_X}{8} \right\rceil$
Segment size (in blocks per segment)	S'	128	128
Output word size	W	1	1
Type of DWT used in the 2D coder		Integer DWT	Float DWT
Rate-allocation algorithm		N/A	Lagrange
Spectral transform type		IWT	IWT
Region size	F	32	32
POT arithmetic precision	Ω	12	12
POT flip mode		Bypass	Bypass
AAT arithmetic precision	Ψ	15	15
Bands kept in AAT	N_{TZ}	N_Z	N_Z