The New CCSDS Image Compression Recommendation

Pen-Shu Yeh[1], Philippe Armbruster[2], Aaron Kiely[3], Bart Masschelein[4], Gilles Moury[5], Christoph Schaefer[6], Carole Thiebaut[7]

[1]NASA/GSFC, Code 567, Greenbelt, MD 20771, USA, 301-286-4477, pen-shu.yeh-1@nasa.gov;
[2]ESA/ESTEC, Keplerlaan 1 P.O.Box 299, Noordwijk, 2200AG, The Netherlands, Philippe.Armbruster@esa.int;
[3]NASA/JPL, 4800 Oak Grove Drive, M/S 238-420, Pasadena, CA 91109, USA, Aaron.B.Kiely@jpl.nasa.gov;
[4]IMEC, Kapeldreef 75, Leuven, 3001, Belgium, bart.masschelein@imec.be;
[5,7]CNES, 18, Avenue Edouard Belin, Toulouse, F-31401 Toulouse Cedex9, France, gilles.moury@cnes.fr,
Carole.Thiebaut@cnes.fr
[6]EADS/Astrium GmbH, An der B31, 88039 Friedrichshafen, Germany, christoph.schaefer@astrium.eads.net

Abstract—The Consultative Committee for Space Data Systems (CCSDS) data compression working group has recently adopted a recommendation for image data compression, with a final release expected in 2005. The algorithm adopted in the recommendation consists of a two-dimensional discrete wavelet transform of the image, followed by progressive bit-plane coding of the transformed data. The algorithm can provide both lossless and lossy compression, and allows a user to directly control the compressed data volume or the fidelity with which the wavelet-transformed data can be reconstructed. The algorithm is suitable for both frame-based image data and scan-based sensor data, and has applications for near-Earth and deep-space missions. The standard will be accompanied by free software sources on a future web site. An Application-Specific Integrated Circuit (ASIC) implementation of the compressor is currently under development. This paper describes the compression algorithm along with the requirements that drove the selection of the algorithm. Performance results and comparisons with other compressors are given for a test set of space images.

TABLE OF CONTENTS

1. INTRODUCTION 1
2. ALGORITHM DESCRIPTION 2
3. PERFORMANCE 4
4. STATUS 5
5. CONCLUSION 6
ACKNOWLEDGEMENT 7
REFERENCES 7
BIOGRAPHY 7

1. INTRODUCTION

The benefits of data compression to space missions include increased ability to collect science data, and reductions in onboard storage and telemetry bandwidth requirements. Because of these benefits, the Consultative Committee for Space Data Systems (CCSDS) has been engaged in recommending data compression standards for space applications.

The first CCSDS data compression recommendation, adopted in 1997, standardized a version of the lossless Rice compression algorithm [1]. Space missions benefiting from this recommendation range from deep space probes to near Earth observatories.

In 1998, the CCSDS data compression working group began to assess the feasibility of establishing an image compression recommendation suitable for spaceborne applications. The working group agreed that a suitable compressor must meet the requirements listed in Table 1, which were intended to reflect the envisioned application

Table 1. CCSDS Image Compression Requirements

1	Process both frame and non-frame (push-broom) data
2	Offer adjustable coded data rate or image quality (up to lossless)
3	Accommodate from 4-bit to 16-bit input pixels
4	Provide real-time processing with space qualified electronics (≥20 Msamples/sec, ≤1 watt/Msamples/sec, based on year 2000 space electronics technology)
5	Require minimal ground operation
6	Limit the effects of a packet loss to a small region of the image.

---

[1] 0-7803-8870-4/05/$20.00© 2005 IEEE
[2] IEEEAC paper #1165, Version 8, Updated January 25, 2005

for real-time hardware compression onboard a spacecraft.

Apart from the requirements listed in Table 1, perhaps the biggest consideration in the algorithm selection process was to optimize rate-distortion performance. The ability to perform progressive compression was viewed as a highly desirable feature, but not mandatory. It was the hope of the working group that if any patents were included in the recommendation, a royalty-free license could be offered to all CCSDS member agencies.

The working group also assembled a set of 20 test images including Earth observations, star fields, galaxies and solar images. The dynamic ranges of the test images include 8-bit, 10-bit, 12-bit and a 16-bit radar image.

Candidate algorithms were proposed, and performance evaluations were conducted based on both quantitative rate-distortion evaluations and subjective assessments of image quality. In addition, implementation architecture studies were performed to assess the real-time processing capabilities of the proposed algorithms. Implementation complexity played a significant role in the final algorithm selection. In particular, an early analysis of ASIC implementation complexity suggested that the JPEG2000 coder [2] was at least a factor of two more complex than other coding options being considered. For spacecraft applications, this could have a significant impact on the achievable processing rate.

A consensus was reached in 2003, when a wavelet-based compression algorithm was selected. The selected algorithm combined elements from different algorithms that were initially proposed, along with modifications to reduce complexity.

In Section 2, we describe the compression algorithm. Compression performance results of the algorithm on the test images are given in Section 3. Section 4 describes the current status of the recommendation.

2. ALGORITHM DESCRIPTION

The recommended algorithm consists of two functional modules as depicted in Figure 1: a Discrete Wavelet Transform (DWT) module that performs decorrelation, and a Bit-Plane-Encoder (BPE) that encodes the decorrelated data. This general image compression approach is widely used, see, for example, references [2, 3].

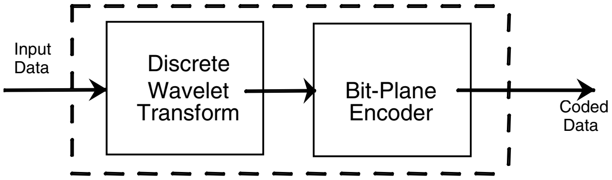

Figure 1 - The Two Functional Modules of the Algorithm

2.1 Discrete Wavelet Transform

The recommendation specifies two DWTs that may be used. When applied to one-dimensional data, both transforms effectively use 9 filter taps to compute low-pass output, and 7 filter taps to compute high-pass output. Each filter is thus referred to as a "9/7" DWT under the usual naming convention. The two filters differ in the need for floating-point arithmetic. The floating-point filter [2] requires

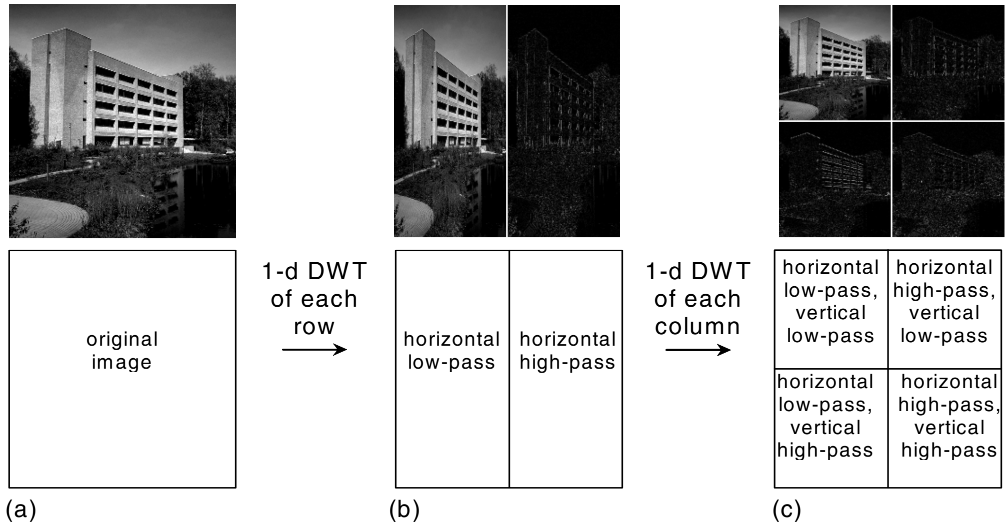

Figure 2 - Single Level Two-Dimensional DWT Decomposition of an Image

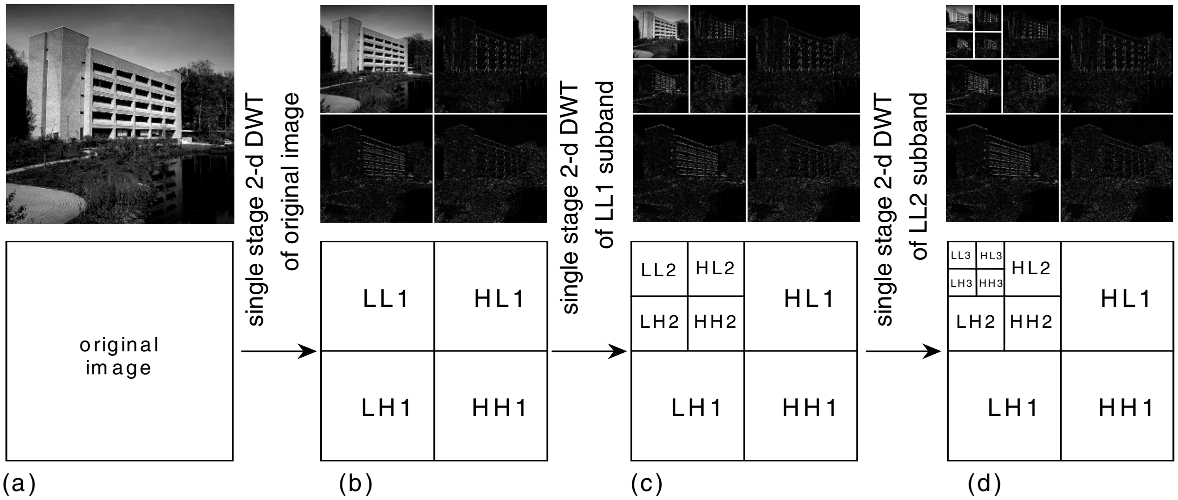

Figure 3 - Example of 3-Level Two-Dimensional DWT Decomposition of an Image

floating-point calculations and gives improved performance at low bit rates, while the integer filter [4] permits lossless compression and requires no floating-point operations. There are many variations in methods for computing integer and floating point 9/7 DWTs, and the reader is encouraged to refer to [5] for exact specifications of the forward and inverse transforms that are to be used with this recommendation.

A single-stage two-dimensional DWT is computed by first applying the one-dimensional DWT to the rows of the image, and then to the columns of the transformed image, as illustrated in Figure 2. Subsequent stages of decomposition are applied to the low-pass horizontal / low-pass vertical subband output from the previous stage, producing the pyramidal decomposition described in [6]. The standard calls for 3 stages of DWT decomposition, decomposing an image into 10 subbands, as illustrated in Figure 3. Increasing the number of levels of wavelet decomposition offers the potential for increased compression effectiveness, but was not selected for the recommendation because of the resulting increase in implementation complexity.

The BPE described in Section 2.2 is used to encode the subbands produced by the two-dimensional DWT decomposition. For effective operation, the BPE relies on the same bit plane in each of the subbands having approximately the same relative priority in terms of contribution to overall image distortion. For the integer transform, this requires the subbands to be scaled. The scaling factors are chosen to be powers of two so that scaling can be performed using bit-shift operations. Under the floating-point DWT, no scaling is performed, but DWT coefficients are rounded to the nearest integer.

2.2 Bit Plane Encoder

The BPE processes wavelet coefficients in groups of 64 coefficients referred to as a *block*. A block loosely corresponds to a localized region in the original image. A block consists of a single coefficient from the lowest spatial frequency subband, referred to as the *DC coefficient*, and 63 *AC coefficients*, as illustrated in Figure 4. Blocks are processed in raster scan order, i.e., rows of blocks are processed from top to bottom, and proceeding from left to right horizontally within a row.

This structure is used to jointly encode information pertaining to groups of coefficients within the block because these coefficients exhibit significant statistical correlation.

The set of blocks in an image is partitioned into groups called *segments*. The blocks in a segment are consecutive in raster-scan order. Coding of DWT coefficients proceeds segment-by-segment and each segment is coded independently of the others. The number of blocks in a segment can be assigned by the user to any value between 16 and 2^{20} inclusive; the value might be chosen based on the memory available to store the segment. The use of small segments provides the potential benefit of reducing implementation memory requirements and of confining the effects of transmission errors to smaller regions of the image. However, when the compressed data volume is controlled by a constraint on the compressed size of each segment, the use of smaller segments generally results in reconstructed segments of an image having varying quality levels, and causes reduced overall rate-distortion performance.

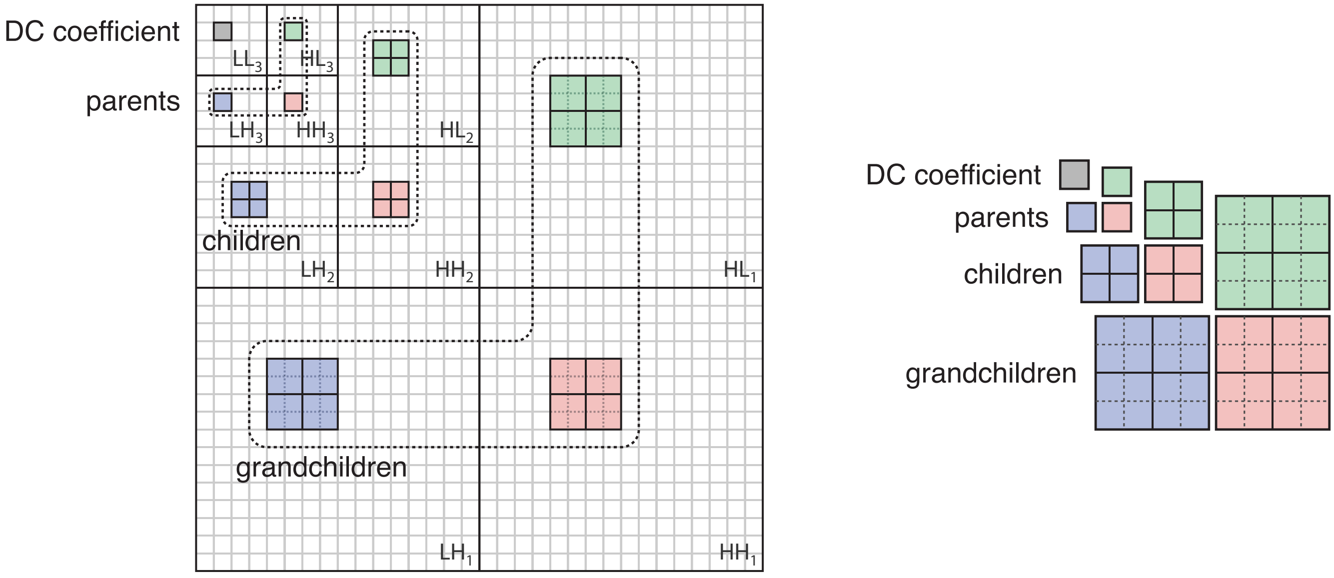

wavelet-transformed image

a single block

Figure 4 - In this schematic of a wavelet-transformed image, the 64 shaded pixels comprise a single block

Within a segment, the BPE first encodes a quantized version of the DC coefficients for the segment by applying the Rice coding algorithm to differences between successive quantized coefficients. Bits providing further DC coefficient resolution are included (uncoded) as part of the subsequent bit-plane coding process.

Next, the BPE successively encodes bit planes of coefficient magnitudes in a segment, proceeding from most-significant to least-significant bit plane, inserting AC coefficient sign bit values at appropriate points in the encoded data stream. The resulting encoded bitstream constitutes an embedded data format that provides progressive transmission within a segment; DWT coefficient resolution effectively improves by a factor of 2 as encoding proceeds from one bit plane to the next.

Coefficients within a block are arranged in groups, each with at most 4 coefficients. Conceptually, at a given bit plane, a binary word can be used to describe an update to each coefficient in the group for which all more significant magnitude bits are zero. These words are entropy coded using one of a handful of variable-length binary codes; the specific code is selected adaptively. The entropy coded data are arranged so that all parent coefficients in the segment are updated first, followed by children, and then grandchildren coefficients. Finally, the segment includes (uncompressed) update bits for the coefficients in the segment for which more significant magnitude bits are not all zero.

The tradeoff between reconstructed image quality and compressed data volume for each segment can be controlled by specifying the maximum number of bytes in each compressed segment, and a "quality" limit that constrains the amount of DWT coefficient information to be encoded. Compressed output for a segment is produced until the byte limit or quality limit is reached, whichever comes first. The encoded bitstream for a segment can be further truncated (or, equivalently, coding can be terminated early) at any point to further reduce the data rate, at the price of reduced image quality for the corresponding segment.

3. PERFORMANCE

The quantitative performance of the new recommendation has been evaluated on the test images listed in Table 2. These images are available at CCSDS public website[3]. Two examples of test images are given in Figure 5.

As an indication of compression performance, we compare the two DWT options of the recommendation with the JPEG2000 standard when used with the 9/7 floating-point DWT. For this comparison, we simulate performance for "push-broom" spacecraft compression applications. In the case of the CCSDS recommendation, this means defining a segment of blocks to correspond to the image width, and imposing a fixed rate constraint on each compressed segment. Similar constraints are imposed on the JPEG2000 coder by using the scan-based mode introduced by SAIC and CNES with 8-line precincts [7]. We evaluated the performance at bit rates ranging from 1/4 bit/pixel up to 2 bits/pixel.

---

[3] http://public.ccsds.org/sites/cwe/sls-dc/Public/Forms/AllItems.aspx

Table 2. List of Test Images

Image	**Source**	**Size**	**Bits/ pixel**
marstest	Mars Pathfinder	512×512	8
spot_la_b1/b2/b3	SPOT Imaging	500×500	8
spot_panchromatic	SPOT Imaging	1000×1000	8
forest_2kb1/b4	NOAA Polar Orbiter	2048×2048	10
ice_2kb1/b4	NOAA Polar Orbiter	2048×2048	10
india_2kb1/b4	NOAA Polar Orbiter	2048×2048	10
north_atlantic_1kb1/b4	NOAA Polar Orbiter	1024×1024	10
ocean_2kb1/b4	NOAA Polar Orbiter	2048×2048	10
solar	Big Bear Solar Observatory	1024×1024	12
sun_spot	Big Bear Solar Observatory	512×512	12
wfpc	Hubble Space Telescope	800×800	12
foc	Hubble Space Telescope	1024×512	12
sar16bit	ERS-1	512×512	16

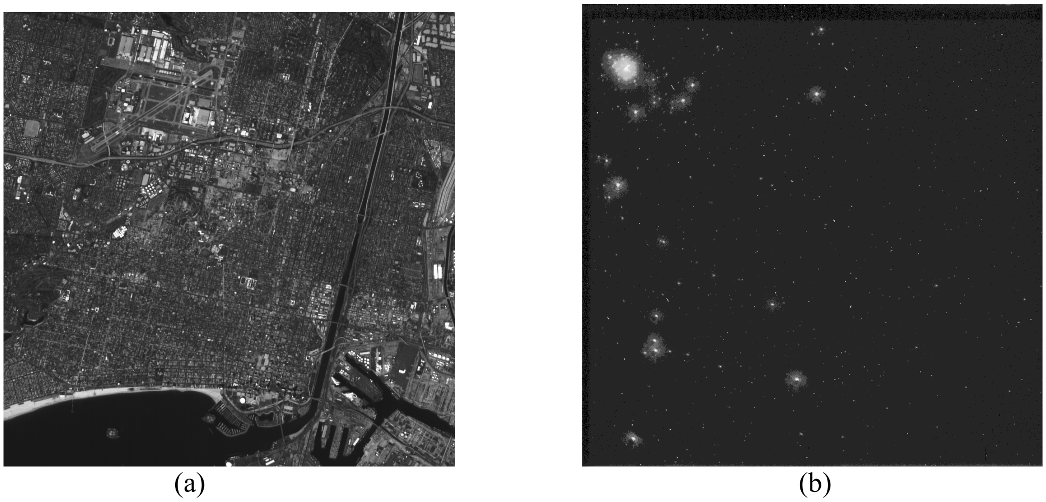
(a) (b)

Figure 5 - Examples of Test Images: (a) SPOT Panchromatic Image and (b) Wide Field Planetary Camera Image

Table 3 shows the Peak-Signal-To-Noise Ratio (PSNR) in dB and the maximum absolute error averaged over images with the same dynamic range. It is seen from these results that the new CCSDS recommendation has performance similar to that of the JPEG2000 standard when both methods use the floating point 9/7 DWT under the "push-broom" constraints described above. As one might expect, use of the integer 9/7 DWT results in more than 1 dB loss in performance at higher bit rates.

4. STATUS

A first version draft of the new image compression recommendation (known as a *red book* in CCSDS parlance[4]) has been approved for agency review. The data compression working group will issue a second draft after

---

[4] The CCSDS web site, www.ccsds.org, describes the meaning of the different books and includes downloadable versions of CCSDS recommendations.

Table 3. Performance Comparison for Push-Broom Mode

Rate (bits/pixel)	**PSNR** (dB)			**Maximum Absolute Error**		
	CCSDS		JPEG2000	CCSDS		JPEG2000
	Floating-point DWT	Integer DWT	Floating-point DWT	Floating-point DWT	Integer DWT	Floating-point DWT
Average for 8-bit test images						
2.00	41.37	40.22	41.47	18.00	15.20	13.20
1.00	35.76	35.18	35.52	34.40	29.00	32.20
0.50	32.37	31.95	32.06	60.60	52.20	52.80
0.25	29.89	29.50	29.48	93.40	81.80	89.20
Average for 10-bit test images						
2.00	54.70	53.26	54.92	26.00	23.50	18.30
1.00	47.76	47.10	47.80	63.20	53.40	53.90
0.50	42.97	42.60	42.90	115.50	95.10	113.20
0.25	39.36	39.12	39.32	204.80	188.30	195.30
Average for 12-bit test images						
2.00	65.93	64.30	66.49	33.33	28.00	22.70
1.00	61.18	60.17	61.20	59.00	50.67	46.30
0.50	58.57	57.87	58.48	88.67	78.67	83.70
0.25	56.62	56.12	56.29	142.00	141.33	139.30

taking into account review comments from different agencies. Following agency review of the second red book, it is expected that a formal recommendation, a *blue book*, will be released in 2005.

The compression working group is also producing a *green book* which is not part of the recommendation, but will serve as a user's guide for implementers. The green book will cover subjects such as system issues relating to error propagation and rate control, implementation schemes for wavelet transform using localized transform, and detailed study results.

Several implementations are being pursued concurrently both for the purpose of validating the recommendation, and also to provide a technology demonstration for space implementation. Software implementations have been produced at JPL by Mikhail Garvey and GSFC. Additional software implementations are under development at the University of Idaho, in an effort led by Prof. Gary Maki, and the University of Nebraska in an effort led by Prof. Khalid Sayood. These codecs are written in C and are in the process of cross-verification. A JAVA implementation based on earlier documentation was developed at the University of Barcelona under the direction of Prof. Joan Serra. A hardware ASIC implementation is being developed at the University of Idaho's Center for Advanced Microelectronics and Biomolecular Research (CAMBR) facility[5] where the Radiation-Hardness-By-Design (RHBD) technique [8] has been developed and is being applied to the algorithm to produce high-speed space-qualified circuits. The projected throughput is over 20 Msamples/sec. This implementation separates the DWT and BPE into two ASICs.

The software development and verification is expected to be completed before the publication of the green book, which will then include an open-source website for users to download and execute the codes. The ASIC flight hardware will be available commercially.

5. CONCLUSION

The CCSDS data compression working group has finalized an algorithm for image data compression, intended for on-board spacecraft use. The algorithm yields nearly the demanding rate-distortion performance of the commercial JPEG2000 standard but reduces on-board implementation complexity.

The recommendation makes use of 9/7 DWTs. An integer DWT can be used for applications requiring lossless compression, or to avoid floating point operations in the

---

[5] www.cambr.uidaho.edu

DWT calculation. A floating-point DWT can be used for improved rate-distortion performance.

The DWT is followed by a bit-plane encoder that produces an encoded bitstream providing progressive transmission within a coded segment.

The algorithm is applicable to a variety of imaging instruments, and is suitable for push-broom sensors requiring immediate processing of data.

The final recommendation is expected to be released in 2005. An open source C software implementation is expected to be available soon, and an ASIC hardware implementation is currently under development.

ACKNOWLEDGEMENT

Portions of the research described in this paper were conducted by Pen-Shu Yeh for the U. S. Government. Portions of the research described in this paper were carried out by Aaron Kiely at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

[1] CCSDS 121.0-B-1, Lossless Data Compression, CCSDS, 1997.

[2] *JPEG2000: Image Compression Fundamentals, Standards and Practice*, D. Taubman, M. Marcellin, Kluwer Academic Publishers, 2002.

[3] A. Said and W. Pearlman, "A New, Fast, and Efficient Image Codec Based on Set Partitioning in Hierarchical Trees," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 6, no. 3, pp. 243–250, June 1993.

[4] A. R. Calderbank, I Daubechies, W. Sweldens, B.-L. Yeo, "Wavelet Transforms that Map Integers to Integers," *Appl. Comput. Harmon. Anal.*, vol. 5, pp. 332–369, July 1998.

[5] CCSDS 122.0-R-1, *Image Data Compression*, CCSDS, 2004

[6] S. G. Mallat, "A Theory for Multiresolution Signal Decomposition: The Wavelet Representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 7, pp. 674--693, July 1989.

[7] P.-S. Yeh, G. Moury, and P. Armbruster, "The CCSDS Data Compression Recommendations: Development and Status," *Proc. SPIE Application of Digital Image Processing*, July 7-10, 2002, Seattle, WA.

[8] J. Gambles and G. Maki, "Radiation Effects and Hardening Techniques for Spacecraft System Microelectronics," *IAF World Space Congress*, Reference IAC-02-I.05.08, October, 2002.

BIOGRAPHY

Pen-Shu Yeh *works for NASA's Goddard Space Flight Center. She has been leading the development of data compression and onboard processing technology for over sixteen years at Goddard. She has supported various space missions in implementing data compression and currently chairs the Data Compression Working Group within CCSDS. Pen-Shu Yeh received a Ph.D in Electrical Engineering in 1981 from Stanford University after completing a BSEE at the National Taiwan University and a MSEE at the University of Washington in Seattle. Her research interests include signal processing, pattern recognition, computer vision and implementation using radiation-hard space electronics.*

Philippe Armbruster *received an Engineering degree of Physics and Electronics in 1983 and then a PhD in Signal and Image Processing from the Louis Pasteur University of Strasburg, France. After having completed the development of a high resolution image digitizing and processing system, he started working in 1989 at the technical center of the European Space Agency, ESTEC, located in The Netherlands. He was first in charge of developing signal processing devices and electronic units for on-board satellite payload data processing applications. Being nominated head of the Data Systems division in 2003, he is involved in several technological developments focusing on ground/space communication, instruments inter-operability via standardized interfaces, generic building blocks and openly specified techniques such as data compression algorithms for space applications.*

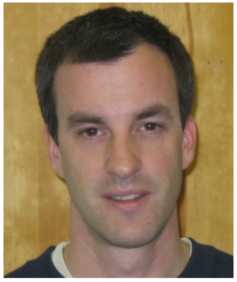

***Aaron Kiely** received the B.S. degree in Electrical Engineering from Virginia Tech in 1989, the M.S.E. and Ph.D. degrees in Electrical Engineering: Systems from the University of Michigan in 1990 and 1993, respectively, and the M.S. degree in Aerospace Engineering from the University of Southern California in 1999. Since 1993 he has worked in the Communications Architectures and Research Section at JPL, where he conducts research in data compression and error-correcting codes. Aaron has also served on the faculty of Caltech, where he has periodically taught a course on error-correcting codes. He led the development of the ICER wavelet-based image compressor being used by the Mars Exploration Rovers. Aaron's current work includes research on wavelet-based compression methods for hyperspectral imagery.*

***Bart Masschelein** was born in Izegem, Belgium. He received an Industrial Engineering degree of Electronics in 1998 from the Katholieke Hogeschool Brugge Oostende, Oostende (Belgium). He received his Master of Science in Electronic Design at the KHBO (Ostend, Belgium) in cooperation with Leeds Metropolitan University in 1999. On August 1998 he started working at the Multimedia Image Compression System group of IMEC/DESICS, currently called MultiMedia group. His first research topics were wavelet-based still-image compression systems for space application, which led to the Local Wavelet Transform, an instruction-based memory-efficient wavelet implementation. Currently, he is involved in the implementation of a memory-efficient Scalable Video Codec, as developed within the MPEG-21 committee. He has also been supporting the European Space Agency in developing the new CCSDS compression recommendation.*

***Gilles Moury** graduated from Ecole Polytechnique in 1983 and from Sup'Aero in 1985. He joined CNES – the French space agency – at the Toulouse space center in 1985. He was responsible for the development of a number of on-board data processing and storage equipments for satellites, like the Solid State Recorder of the SIGMA mission, the Image Compression Module of the Clementine & Cassini missions, the image compression algorithm and equipments of SPOT5 & HELIOS2 missions. He is now head of the on-board data processing section at CNES in Toulouse within the technical directorate. He represents CNES at the international standardization committee CCSDS, in the source coding, TM/TC and channel coding panels.*

***Christoph Schaefer** obtained his Ph.D. from the Institute of Applied Physics in Tübingen in 1982 after a formal education in mathemetics and physics at the Universities of Maryland (USA) and Tübingen (German). He has worked for IBM T.J.Watson Research in Yorktown Heights, New York, from 1983-1986 in the field of numerical electron optics. From 1986-1996, he has been affiliated with Dornier GmbH in Friedrichshafen, Germany and has been working in the field of computation of Radar Cross Section and the fields of robotics and image processing, with particular emphasis on vision-based object recognition and obstacle detection. He has been a Senior Advisor for Synthetic Aperture Radar & Image Processing Systems for EADS/Astrium GmbH in Friedrichshafen since 1996, with focus on space-borne data handling systems. His interests include signal processing, software and hardware implementation, and architectures.*

***Carole Thiebaut** was born in North of France in 1977. She received an Engineering degree of Electronics and Signal Processing in 2000 from the Institut National Polytechnique de Toulouse (ENSEEIHT). She received her Ph.D. in Sept. 2003 in Signal and Image Processing with a Ph.D fellowship from the CNRS (National Centre of Scientific Research). During these three years in the CESR (Centre for the Study of Radiation in Space), she has led studies on astronomical image analysis, object detection, automatic classification using neural networks, analysis of astrophysical signals with time-frequency methods. With the support of a post-doctoral fellowship from CNES, the French Space Agency, she started working on optical observations of space debris in Oct. 2003 for CNES. Since Sept. 2004, she is working on on-board data processing at CNES.*