dscal
DIGITAL SYSTEMS & COMPUTER ARCHITECTURE LABORATORY

High Performance CCSDS Image Compression Implementations on Space-Grade SRAM FPGAs

N. Kranitis, G. Theodorou, A. Tsigkanos, A. Paschalis

Dept. of Informatics & Telecommunications

National & Kapodistrian University of Athens (NKUA)

SpacE FPGA Users Workshop, 3rd Edition

15-17 March 2016, ESTEC

Outline

- Motivation
- CCSDS 122.0-B-1 IP Core
  - Image Data Compression
- CCSDS 123.0-B-1 IP Core
  - Multispectral & Hyperspectral Image Compression
- Conclusion

Motivation

- Huge volume of remote sensing data from high-resolution, high-speed imager payloads
- Limited spacecraft data storage resources
- Limited downlink bandwidth
- Dynamic adaptability
  - An adaptable instrument can extend mission lifetime
  - Time-Space Partitioning (savings in mass, volume, power)
- Radiation hardness (TID, SEU, SEFI, SET, etc.)
- High-performance for Gigabit data-rate applications

High-speed on-board image compression implemented on a reconfigurable rad-hard platform!

Virtex-5QV: Space-grade SRAM FPGA

- High-density & high-performance reconfigurable SRAM FPGA
- Radiation-hardened by design (RHBD)
  - SEU hardened flip-flops
  - BRAM EDAC (SECDED)
- Radiation immunity
  - TID 1000krad(Si)
  - Exceptional hardness to SEUs
    - 4-5 upsets/year in GEO (MTTF>2 years)*
    - ≈ 1000x improvement on SEU rate (per bit vs. Virtex-4QV)
  - Exceptional hardness to SEFIs
    - Mean Time to SEFI is 9,930 years/device in GEO
    - ≈ 100x improvement on SEFI rate (per device vs. Virtex-4QV)
  - Total SEL immunity (>100MeV.cm²/mg)
  - Data path protection from SETs

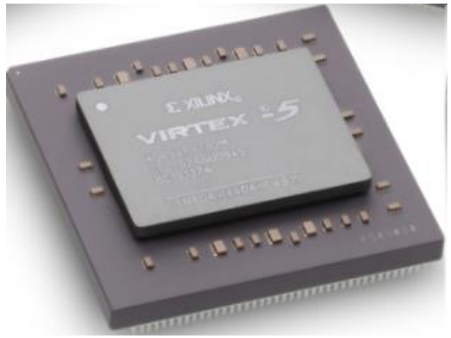

An excellent adaptable and rad-hard platform!

* G. Swift, C. Carmichael, G. Allen, G. Madias, E. Miller, and R. Monreal, "Compendium of XRTC Radiation Results on All Single-Event Effects Observed in the Virtex-5QV", ReSpace/MAPLD 2011

CCSDS Image Compression algorithms

- CCSDS developed image compression algorithms specifically for on-board use
  - Addressing memory and computational resources challenges
  - Excellent trade-off between compression effectiveness & HW complexity

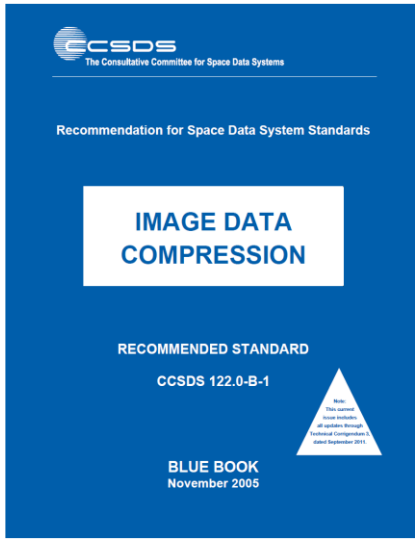

CCSDS
Consultative Committee for Space Data Systems

- Lossy Multispectral & Hyperspectral compression under development
  - CCSDS-122.1-B: Extends CCSDS 122.0-B-1 by defining a one-dimensional pre-processing spectral decorrelating transform and associated output data structures
  - CCSDS-123.1-B: Extends CCSDS 123.0-B-1 by defining a quantization feedback loop and associated output data structures to provide low-complexity near-lossless compression

CCSDS 122.0-B-1: Algorithm Overview

- Recommended standard for monoband image data compression
  - Compression of grayscale images up to 16-bit
- Suitable for use on-board spacecraft
  - Trade-off between compression effectiveness and complexity
  - High-speed & low-power HW implementation feasible
- Widespread applicability to many imaging instruments
  - Frame-based images (i.e. CCD arrays)
  - Strip-based images (i.e. push-broom type sensors)
- Can provide both lossless and lossy compression
  - Rate and quality limited lossy compression
- Segment-based compression
  - Trade-off between data protection and small memory requirements vs. compression effectiveness and rate-distortion performance
- Suitable for Selective Image Compression with ROI support
  - Our CCSDS-IDC implementation was the first that provided selective image compression (PROBA3/ASPIICS SWT, July-2014)

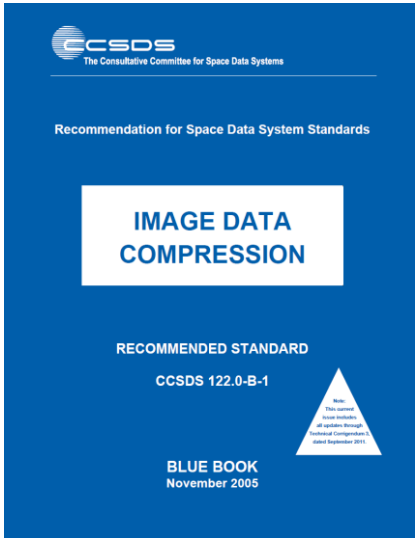

CCSDS 122.0-B-1: Algorithm Overview

- CCSDS 122.0-B-1 (IDC) Blocks
  - Discrete Wavelet Transform (DWT)
    - Performs decorrelation
  - Bit Plane Encoder (BPE)
    - Encodes the decorrelated data

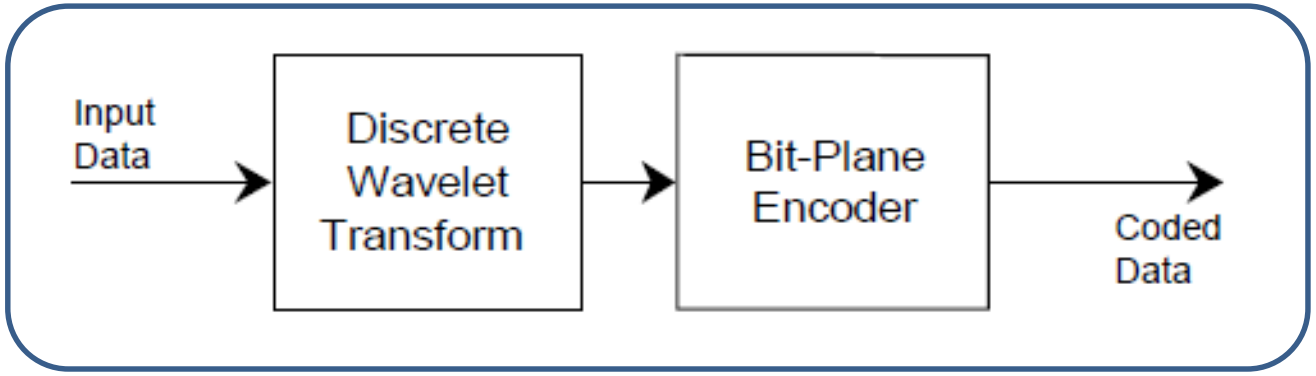

CCSDS 122.0-B-1: Algorithm Overview

- Discrete Wavelet Transform
  - 9/7 Integer DWT (lossless)
  - 9/7 Float DWT (lossy)
  - 3-level 2-d decomposition of an image

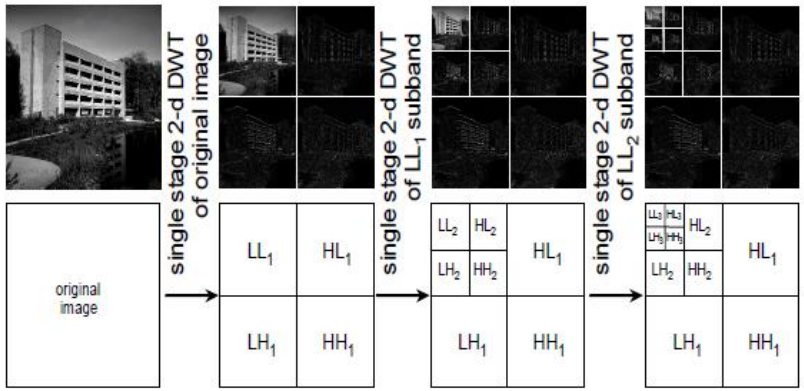

DWT Dataflow

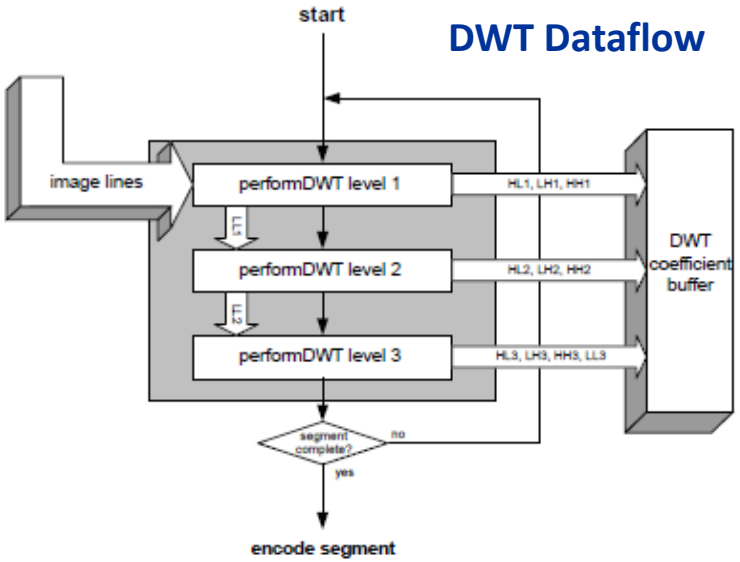

- *Segment*: a group of *S* consecutive blocks
- *Block*: localized region in original image
  - Consists of 64 coefficients (1 DC + 63 AC)
  - AC coefficients arranged into 3 families: F1, F2, F3
  - Blocks processed by BPE in raster scan order
- *Gaggle: 16 blocks*

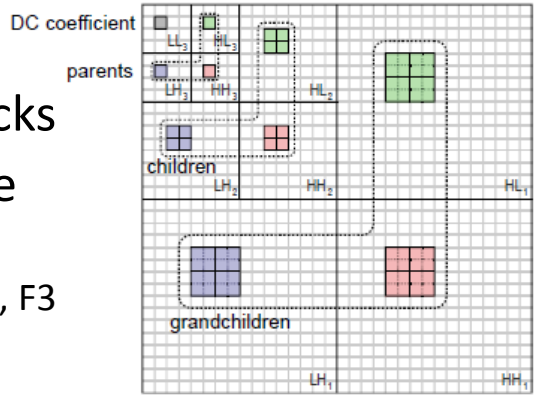

wavelet-transformed image

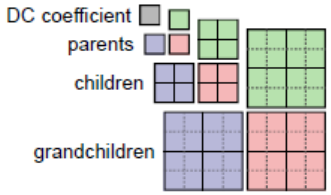

a single block

CCSDS 122.0-B-1: Algorithm Overview

- Bit Plane Encoder
  - Coding of a segment begins when all DWT coefficients are available in segment buffer
  - Coding proceeds segment-by-segment, each segment coded independently
  - Entropy coding is performed bit plane after bit plane, using variable-length codewords

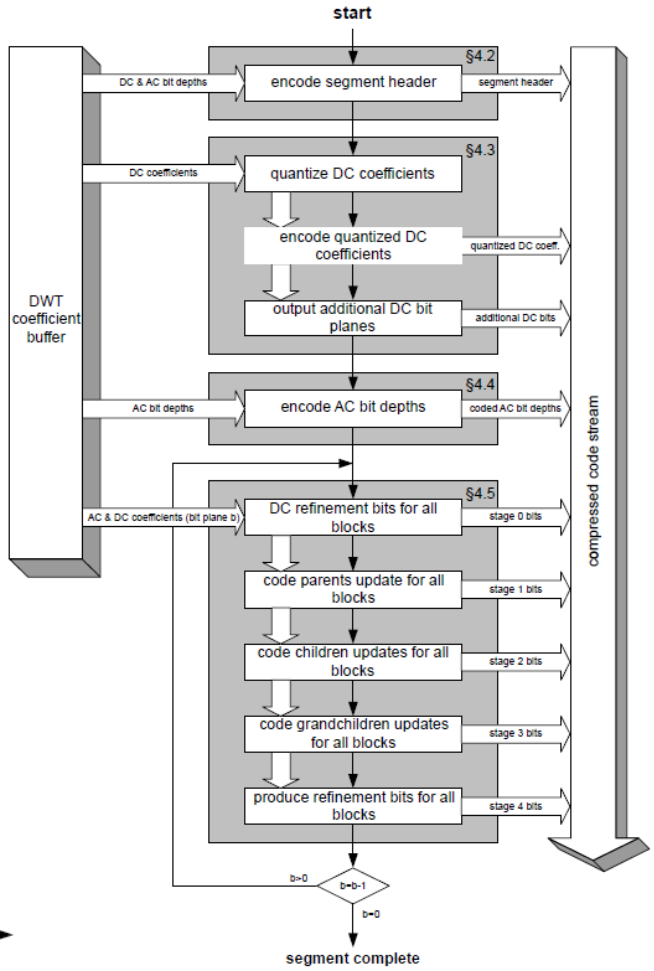

Segment Header (§4.2)
Initial coding of DC coefficients (§4.3)
Coded AC coefficient bit depths (§4.4)
Coded bit plane b = BitDepthAC-1 (§4.5)
Coded bit plane b = BitDepthAC-2 (§4.5)
. . .
Coded bit plane b = 0 (§4.5)

Structure of an encoded segment

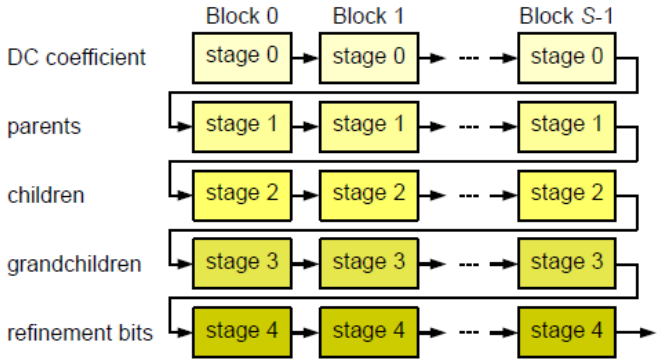

Structure of an encoded bit-plane

BPE Dataflow

CCSDS 122.0-B-1: Existing Implementations

- Software implementations on space platforms
  - UT699 (LEON3-FT), SpW-RTC (LEON2)
  - Less than 1Mbps data-rate performance
- Existing hardware implementations
  - ASIC chip set (CAMBR - U. Idaho, NASA)
    - 2 ASICS (DWT + BPE) + External SRAM
    - 320Mbps @ 16-bit (20Mpixels/s)
  - CCSDS Wavelet Image COMpression ASIC (Airbus DS, ESA)
    - Single-chip solution
    - 60Mpixels/s
  - One-time programmable FPGA (IDA/TU Braunschweig)
    - Microsemi RTAX2000S + External SDRAM
    - 90Mbps @ 12bit (PROBA-V) , 130Mbps@14bit (EnMAP)
  - Reconfigurable FPGA (DSCAL/NKUA)
    - Virtex-5QV (Single chip solution)
    - 78MSamples/s
    - First implementation of CCSDS-IDC for an ESA mission, in a reconfigurable FPGA (PROBA-3/ASPIICS PDR March 2013)

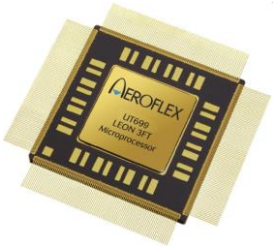

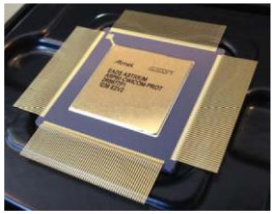

CCSDS 122.0-B-1 IP Core: Features summary

- State-of-the-art data-rate performance
  - 128 MSamples/s (2Gbps @ 16-bit) on V5QV
- Fully parallel, pipelined architecture
- Single FPGA solution. No external memory is required
- High rate-distortion performance for lossy mode
  - Using large values of segment size (S)
- Selective image data compression with ROI support
- Simple, FIFO-based streaming I/O data I/F
- Slave AMBA APB configuration I/F
- Pixel dynamic range: Up to 16 bits
- Parallel, configurable 8/16/32-bit output

Selective Image Data Compression with ROI support

ROI: Lossless
Outside ROI: Only DC coefficients

CCSDS 122.0-B-1 IP Core: Block Diagram

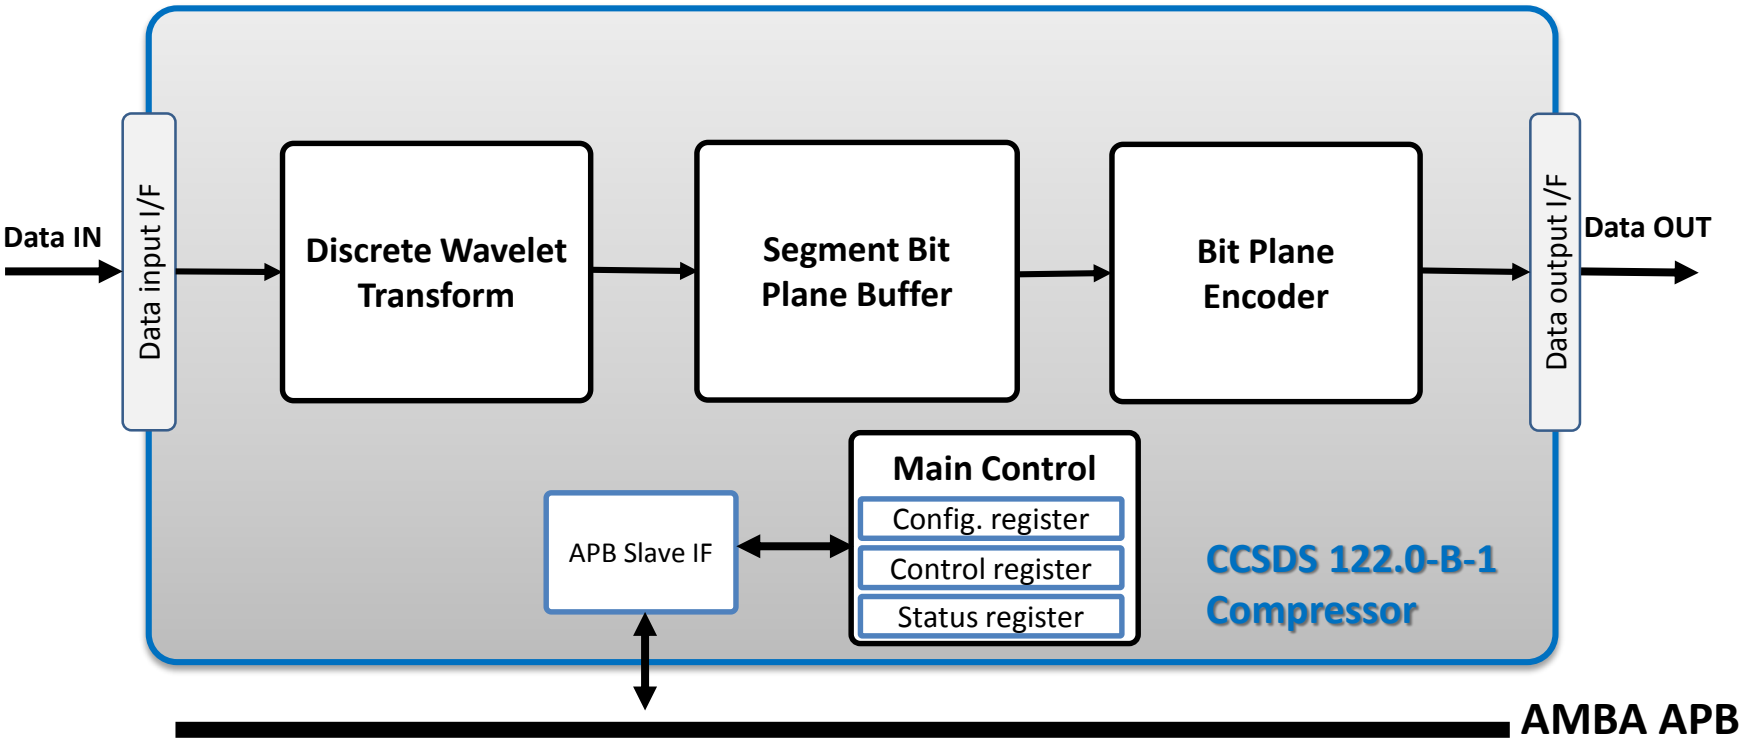

- Data I/O interface
  - Simple, parallel FIFO-based streaming I/F
- Configuration interface
  - Memory mapped configuration registers accessible through Slave AMBA APB I/F

CCSDS 122.0-B-1 IP Core: DWT

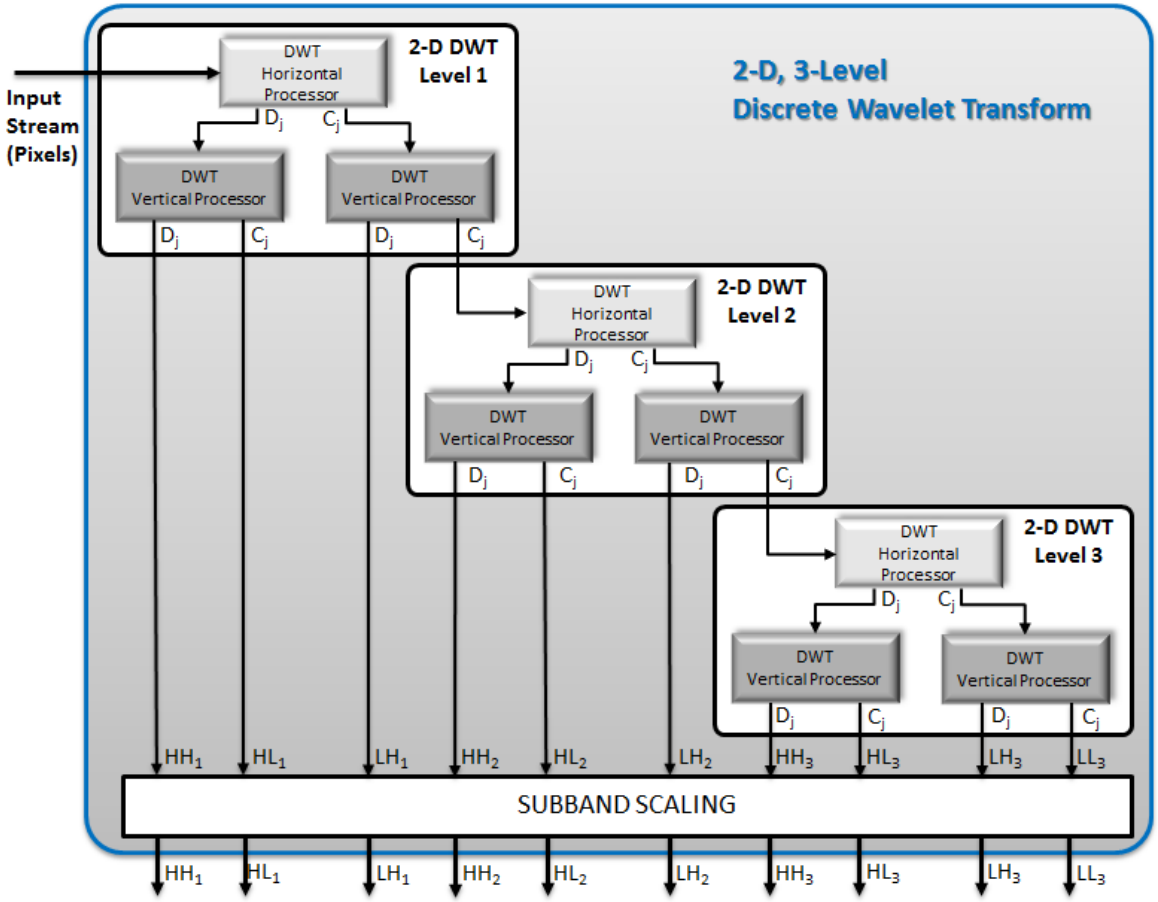

- Non-linear, integer approximation to a 9/7 DWT that
- Provides both lossless & lossy compression
- Parallel, pipelined architecture
- Subband scaling using standard weights (powers of 2)

CCSDS 122.0-B-1 IP Core: SBPB

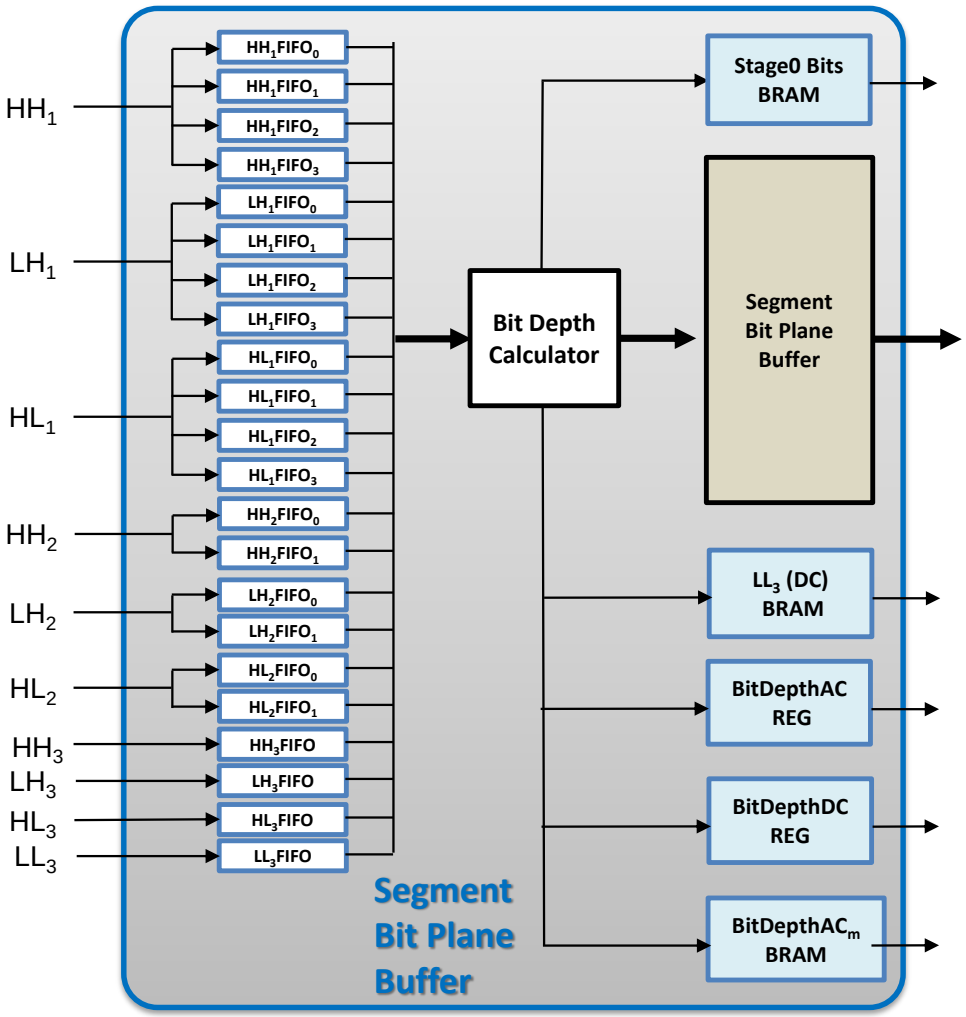

- Hosts the segment DWT coefficients buffer
- Features an efficient DWT coefficient data organization so that BPE can fetch a block with minimum latency
- Performs all bit depth calculations

CCSDS 122.0-B-1 IP Core: BPE

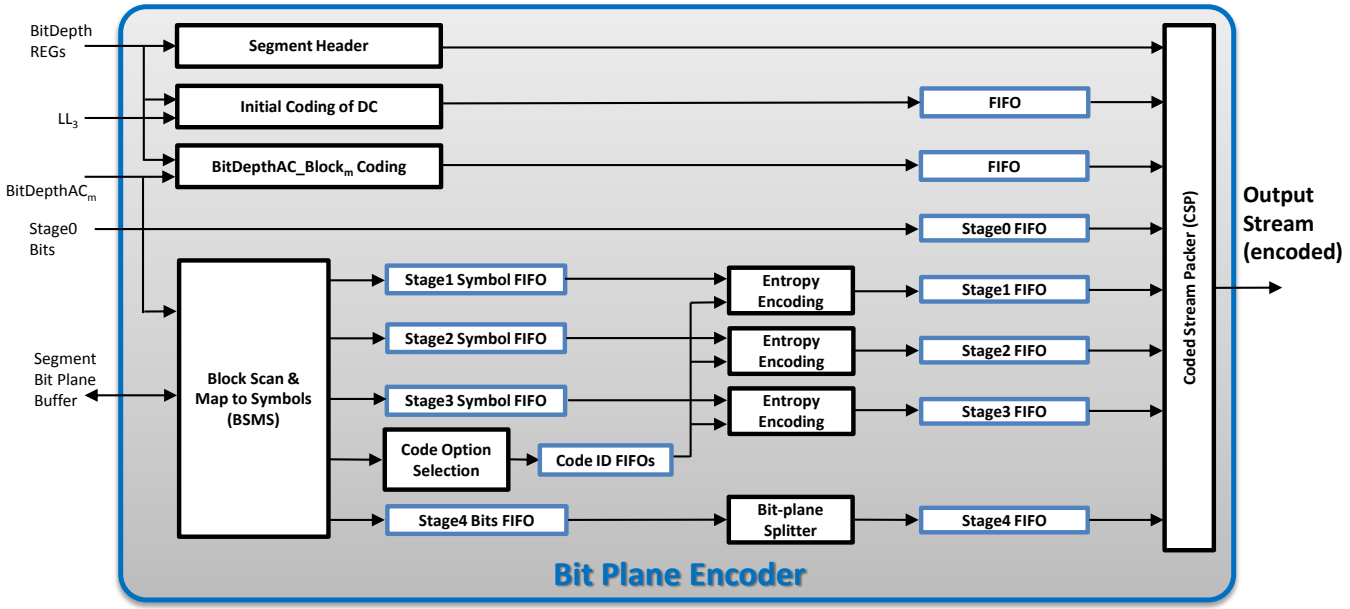

- Fully parallel, fully pipelined architecture
  - Exploits inherent parallelism of BPE processing tasks and high density SRAM FPGAs to boost throughput performance
  - Parallel execution of Segment Header coding, initial coding of DC coefficients, coding of BitDepthAC_Block$_m$ values and coding of bit planes
  - Parallel coding of parent (stage1), children (stage2) and grandchildren (stage3) symbols
  - Pipelining in AC coefficient processing between: binary word generation, symbol mapping, entropy encoding of these symbols and coded stream packetization

CCSDS 122.0-B-1 IP Core: Experimental Results

- FPGA implementation statistics using:
  - Image size: 2Kx2K
  - Segment size: S=128 blocks

Target device	X5VFX130T-1 (Commercial Equivalent of Virtex-5QV)
Device utilization	60% slices, 66% BRAMs
Max frequency	133 MHz
Max throughput	**> 128 MSamples/s (2Gbps @ 16-bit)**

CCSDS 122.0-B-1 IP Core: Experimental Results

- Data-rate (throughput) performance @ 133MHz
  - Compression mode: Lossless
  - Segment size: S=128 blocks

Image	Image description	Throughput performance	
		MSamples/s	Gbps
sar_16bit	ERS-1 ESA, 512x512, 16-bit	128	2.0
P_160_B_F	Picard Imager (IAS) - CNRS, 2048x2048,16-bit	132	2.1
india_2kb1	NOAA Polar Orbiter (AVHRR) – NOAA, 2048x2048,10-bit	132	1.3
ice_2kb1	NOAA Polar Orbiter (AVHRR) – NOAA, 2048x2048,16-bit	132	2.1
marstest	Mars Pathfinder (Sojourner) – NASA, 512x512, 8-bit	128	1.0
random	Random image, 2048x2048, 16-bit	132	2.1
black	Black image, 2048x2048, 16-bit	132	2.1

CCSDS 123.0-B-1: Algorithm Overview

- Recommended standard for Lossless Multispectral and Hyperspectral image compression
- Based on Fast Lossless (FL) algorithm (NASA)
  - Addition of block-adaptive encoder (CCSDS 121.0-B-2) as an alternative option to FL sample-adaptive encoder
- Suitable for use on-board spacecraft
  - Trade-off between compression effectiveness and complexity
  - High-speed HW implementation feasible
- Supports different scan orders for prediction and encoding
  - Band-Interleaved-by-Pixel (BIP)
  - Band-Interleaved-by-Line (BIL)
  - Band-SeQuential (BSQ)

CCSDS 123.0-B-1: Algorithm Overview

- CCSDS 123.0-B-1 Blocks
  - Predictor
    - Estimates predicted sample value based on nearby (3D neighborhood) samples
    - Maps prediction residual into a nonnegative integer
    - Adaptively adjusts prediction weights for each spectral band using adaptive linear prediction
  - Encoder
    - Losslessly encodes the mapped prediction residuals
    - Sample-adaptive (FL) or Block-adaptive (CCSDS 121.0-B-2)

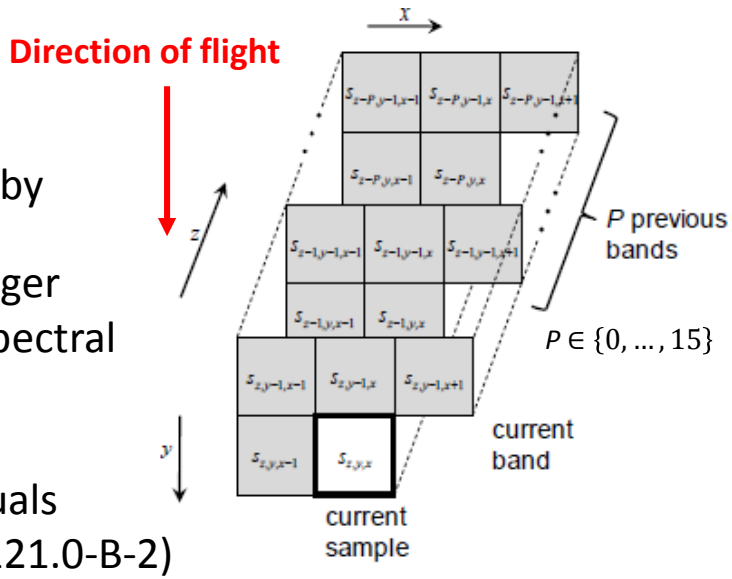

Prediction 3D Neighborhood

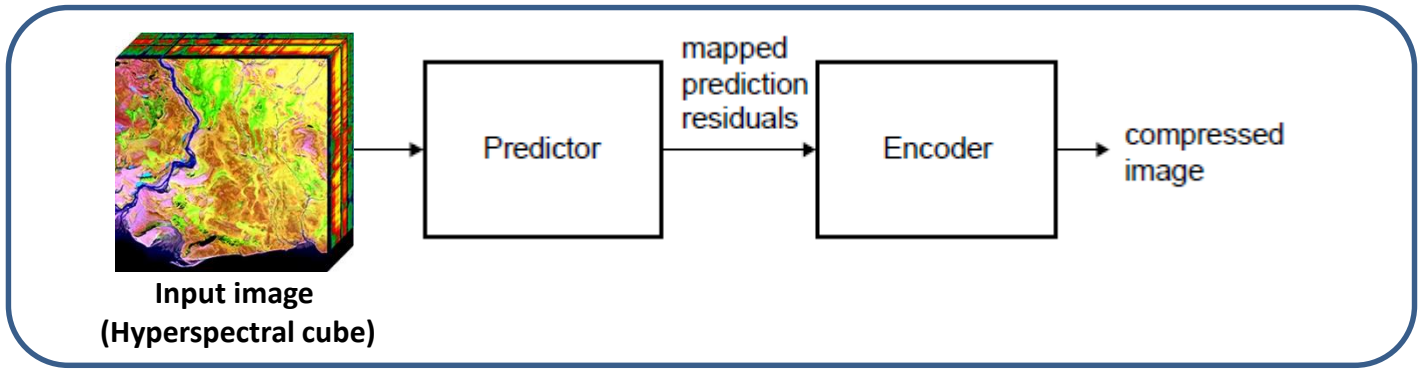

CCSDS 123.0-B-1: Algorithm Overview

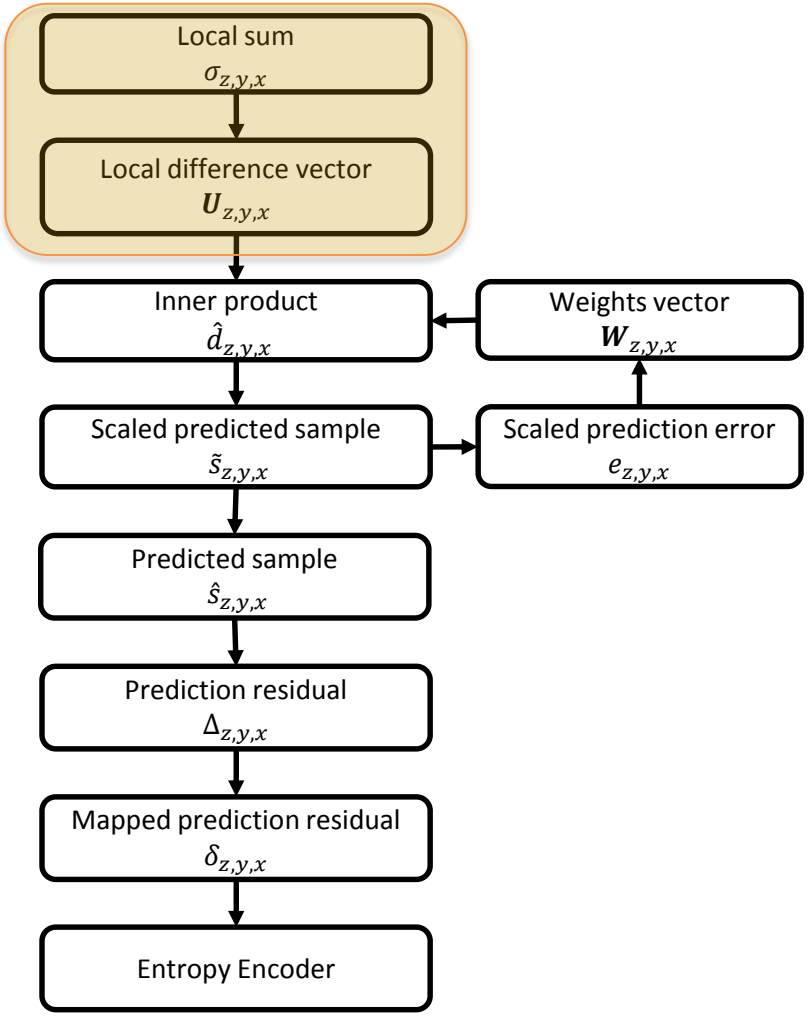

Algorithm flowchart

Local sum: Preliminary estimate of predicted sample

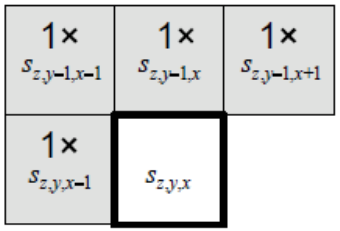

Neighbor oriented

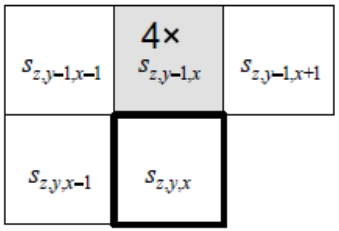

Column oriented

Local differences: Differences between computed local sum and scaled original pixel

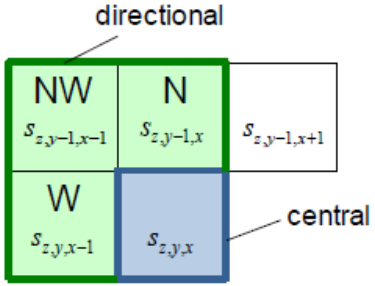

$$\mathbf{U}_{z,y,x} = \begin{bmatrix} d^{N}_{z,y,x} \\ d^{W}_{z,y,x} \\ d^{NW}_{z,y,x} \\ d_{z-1,y,x} \\ d_{z-2,y,x} \\ \vdots \\ d_{z-P,y,x} \end{bmatrix} = \begin{bmatrix} 4s_{z,y-1,x} - \sigma_{z,y,x} \\ 4s_{z,y,x-1} - \sigma_{z,y,x} \\ 4s_{z,y-1,x-1} - \sigma_{z,y,x} \\ 4s_{z-1,y,x} - \sigma_{z-1,y,x} \\ 4s_{z-2,y,x} - \sigma_{z-2,y,x} \\ \vdots \\ 4s_{z-P,y,x} - \sigma_{z-P,y,x} \end{bmatrix}$$

- Reduced prediction mode: Central Local Differences
- Full prediction mode: Central + Directional Local Differences

- Whiskbroom & frame imagers: Full mode + neighbor oriented
- Push-broom imagers: Reduced mode + column oriented

CCSDS 123.0-B-1: Algorithm Overview

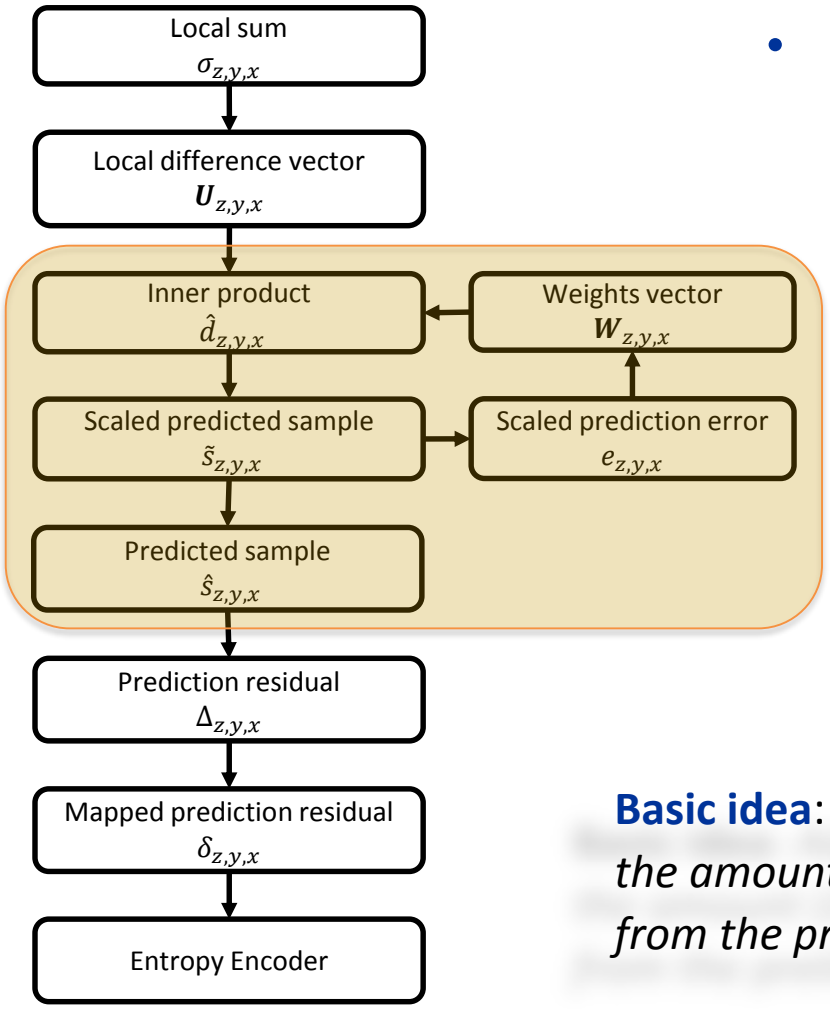

Algorithm flowchart

- **Prediction**
  - An estimation of local difference $\hat{d}_{z,y,x}$ (scalar) is calculated using a weight vector $\boldsymbol{W}_{z,y,x}$

$$\hat{d}_{z,y,x} = \boldsymbol{W}_{z,y,x}^{T} \cdot \boldsymbol{U}_{z,y,x}$$

  - The scaled predicted sample $\tilde{s}_{z,y,x}$ is calculated by adjusting preliminary estimate $\sigma_{z,y,x}$ by $\hat{d}_{z,y,x}$
  - The prediction error $e_{z,y,x}$ along with a weight update scaling exponent $\rho_{z,y,x}$ is used to adapt the weight vector $\boldsymbol{W}_{z,y,x}$ using the sign algorithm
  - The predicted sample $\hat{s}_{z,y,x}$ is calculated

Basic idea: *Adaptively adjust prediction weights to predict the amount ($\hat{d}_{z,y,x}$) by which the sample value $s_{z,y,z}$ differs from the preliminary estimate $\sigma_{z,y,x}$*

CCSDS 123.0-B-1: Algorithm Overview

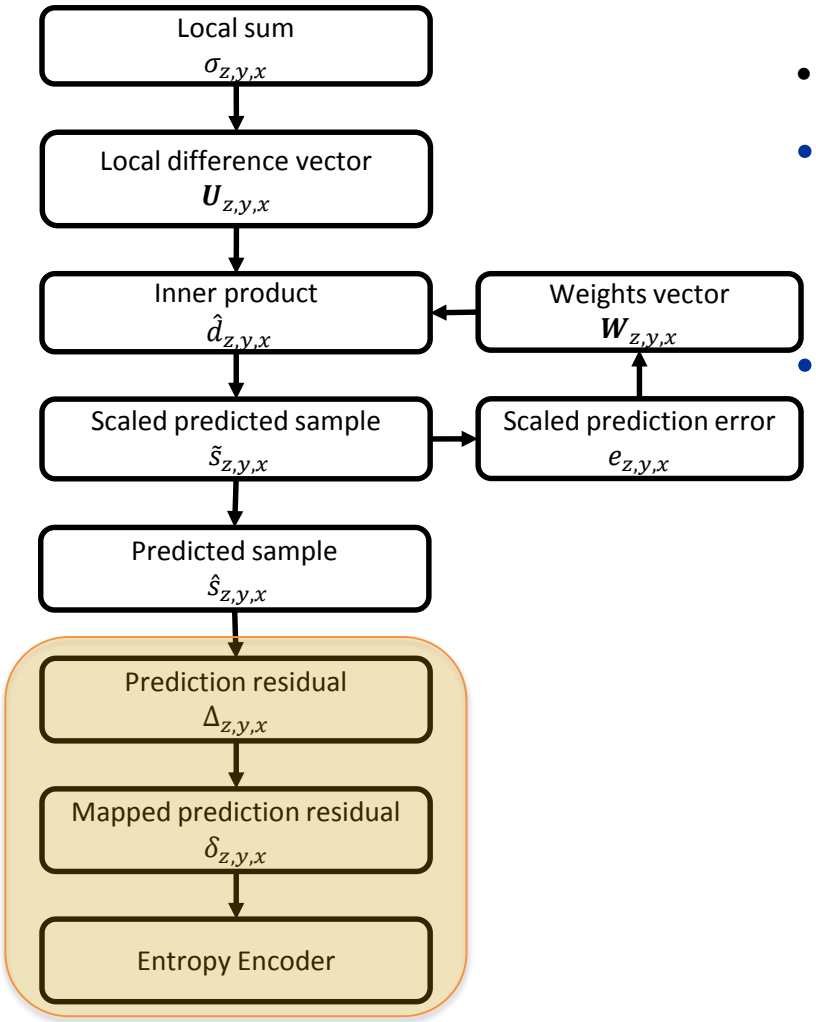

Algorithm flowchart

- The signed prediction residual $\Delta_{z,y,x}$ is calculated
- **Mapping**
  - Prediction residuals $\Delta_{z,y,x}$ are mapped to nonnegative mapped prediction residuals $\delta_{z,y,x}$
- **Entropy Encoder**
  - Sample-adaptive (FL)
    - Mapped prediction residuals are encoded using variable, length-limited Golomb-power-of-2 codes, adjusted based on statistics for each spectral band
    - Compressed size independent of encoding order
  - Block-adaptive (CCSDS 121.0-B-2)
    - Mapped prediction residuals are partitioned into blocks, independently Rice encoded
    - Compressed size depends in encoding order since $\delta_{z,y,x}$ may be from the same or different spectral bands

Sample-adaptive encoding usually more effective than block-adaptive

CCSDS 123.0-B-1 IP Core: Features summary

- State-of-the-art data-rate performance
  >110 MSamples/s (1.7 Gbps @ 16-bit) on V5QV
- Fully parallel, pipelined architecture
- Fully configurable
- Simple, FIFO-based streaming I/O data I/F for pixel data (current, N, NE)
- Slave AMBA APB configuration I/F
- Parallel, 64-bit encoded bitstream output
- Supports optional block-adaptive encoding
  - Available as external CCSDS 121.0-B-2 IP core **

Algorithm Features	CCSDS 123.0-B-1 IP Core
Pixel dynamic range	$2 \leq D \leq 16$
Image Dimensions	$0 < N_X \leq 4096$ $0 < N_Y \leq 4096$ $3 < N_Z \leq 4096$
Encoding order	Band-Interleaved-by-Pixel (BIP)
Number of Prediction bands	P=3*
Prediction mode	Full Reduced
Local sum calculation	Neighbor oriented Column oriented
Entropy encoder	Sample-adaptive Block-adaptive **

* E. Augé, J.E. Sánchez, A. Kiely, I. Blanes, J. Serra-Sagrista, "Performance impact of parameter tuning on the CCSDS-123 lossless multi- and hyperspectral image compression standard" Journal of Applied Remote Sensing (JARS), Volume 7, Issue 1, August 2013

** N. Kranitis, I. Sideris, A. Tsigkanos, G. Theodorou, A. Paschalis, R. Vitulli, "An Efficient FPGA Implementation of CCSDS 121.0-B-2 Lossless Data Compression algorithm for Image Compression", Journal of Applied Remote Sensing (JARS), Volume 9, Issue 1, Special Issue on Onboard Compression and Processing for Space Data Systems, May 2015

CCSDS 123.0-B-1 IP Core: Block Diagram

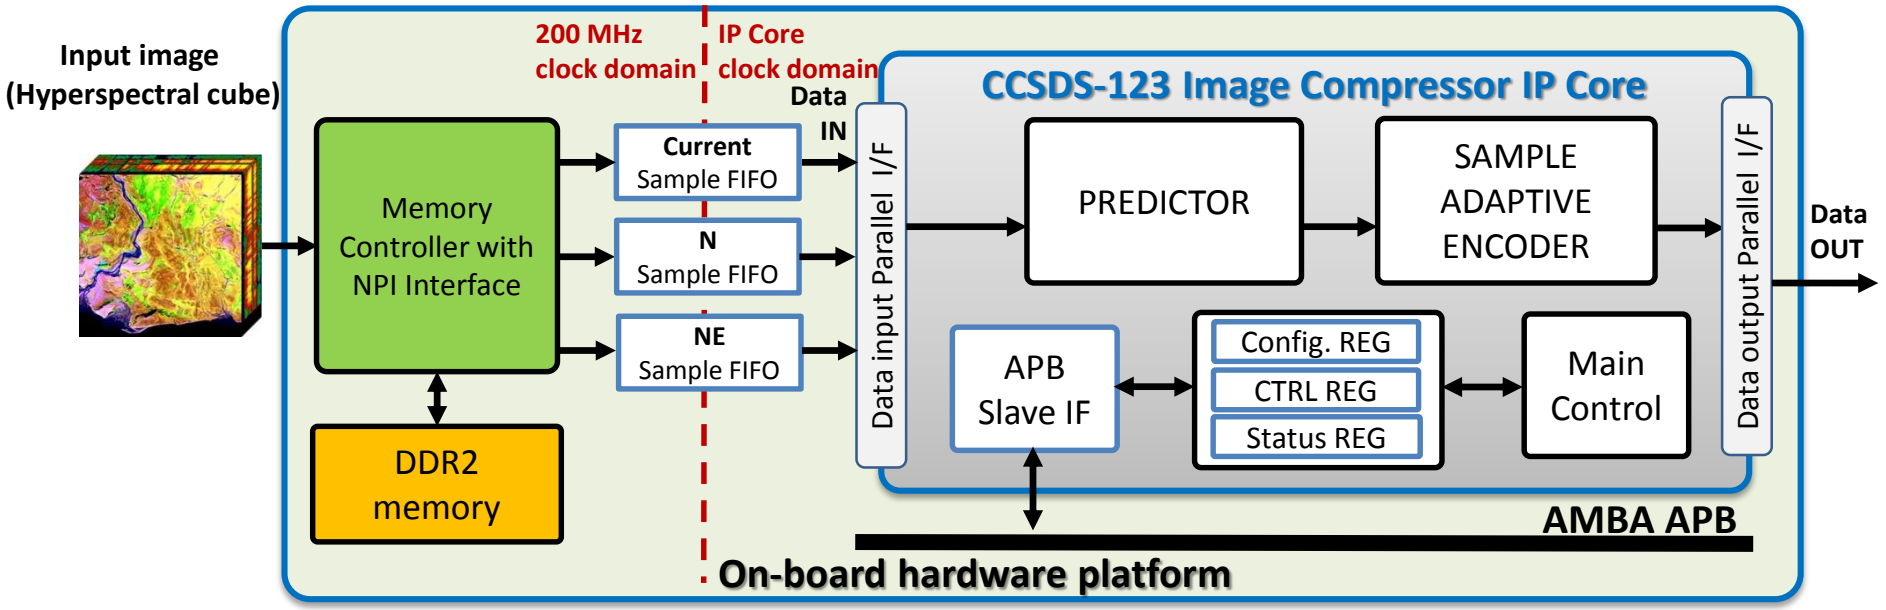

- Data I/O interface
  - Simple, parallel FIFO-based streaming interface for pixel data (current, N, NE)
- Configuration interface
  - Memory mapped configuration registers accessible through a Slave AMBA APB I/F
- External memory controller (Xilinx MPMC) provides access to DDR2 memory
  - Native Port Interface (NPI) streaming interface to sensor and CCSDS 123.0-B-1 IP Core (current, N, NE)
- Modest memory requirements
  - External DDR2 with size enough to hold one spacial-spectral slice of data (N_x*N_z)

CCSDS 123.0-B-1 IP Core: Predictor

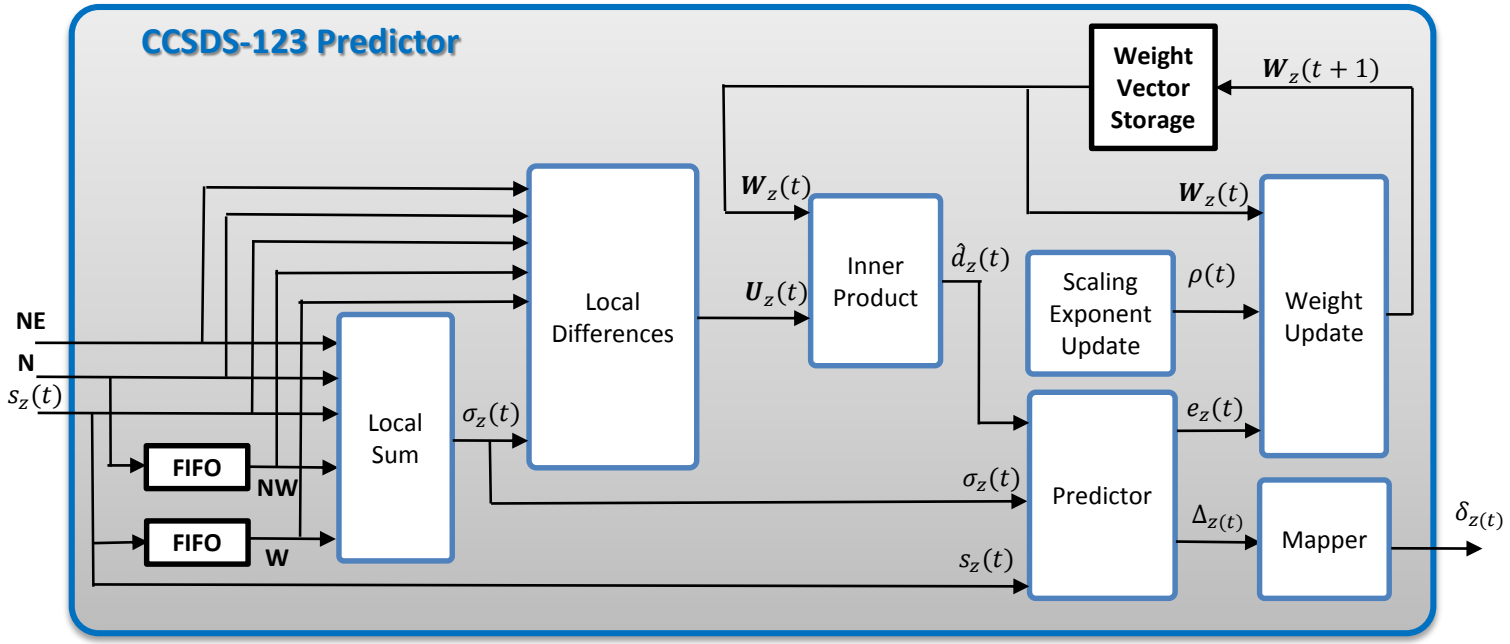

- Sample ordering affects:
  - Local difference and weight vector storage requirements
  - Data dependencies and throughput
  - Difficult pipelining or parallelization under BIL or BSQ
- Band-Interleaved-by-Pixel (BIP)
  - Does not require storage for local differences vector
  - Each spectral band has its own weight vector thus easier to pipeline or parallelize

CCSDS 123.0-B-1 IP Core: Entropy Encoder

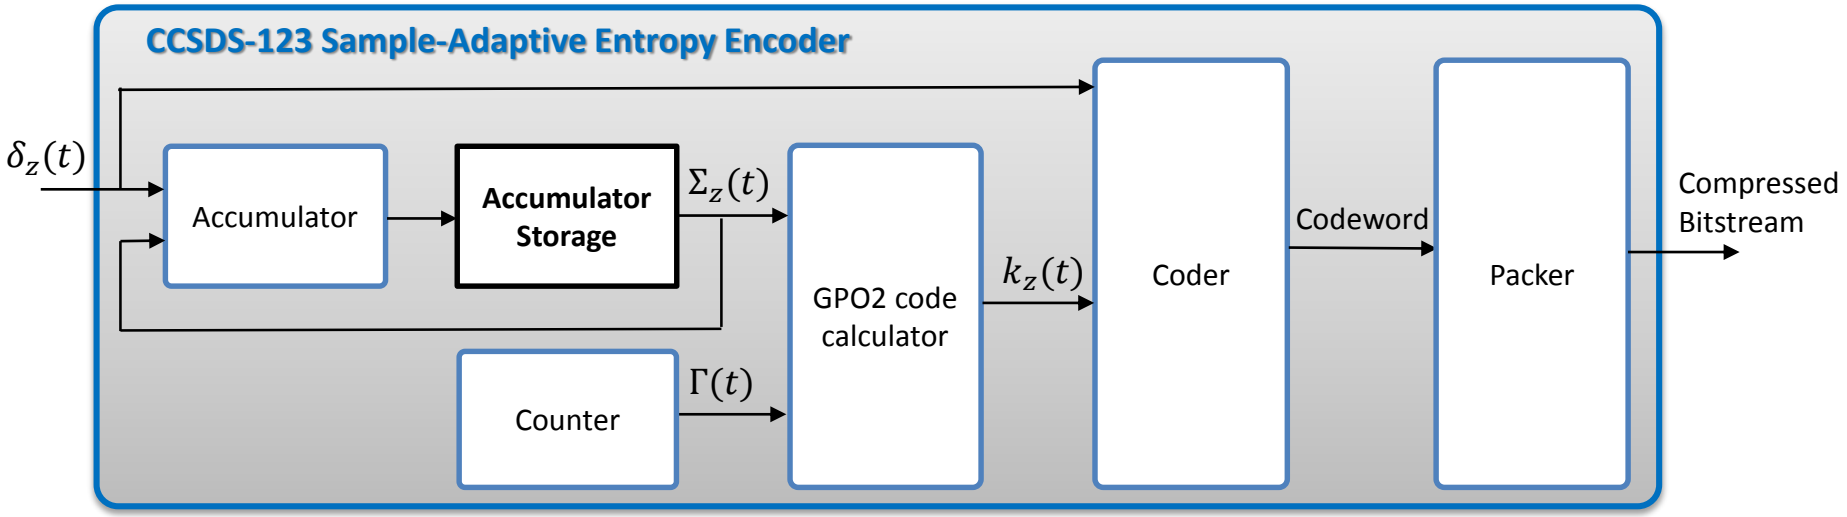

- Sample-adaptive encoder maintains separate coding statistics for each band
- Sample ordering affects accumulator and counter storage requirements
- Band-Interleaved-by-Pixel (BIP)
  - No need to store counter values
  - Each spectral band has its own accumulator vector thus easier to pipeline or parallelize

CCSDS 123.0-B-1 IP Core: Experimental Results

- FPGA implementation statistics using:
  - Image size: N_x, N_y, N_z up to 4096
  - Number of Prediction bands: P=3

Target device	X5VFX130T-1 (Commercial Equivalent of Virtex-5QV)
Device utilization	10% slices, 7% BRAMs
Max frequency	110 MHz
Max throughput	**110 MSamples/s (1.7Gbps @ 16-bit)**

CCSDS 123.0-B-1 IP Core: Experimental Results

- Data-rate (throughput) performance measured @ 110MHz
  - Used the CCSDS corpus of Hyperspectral & Multispectral images
  - Compression parameters according to Green Book (CCSDS 120.2-G-1)

	Image	Image size (N_x, N_y, N_z)	Throughput performance	
			MSamples/s	Gbps
Hyperspectral Images	aviris_hawaii_sc01_raw	614x512x224, 12-bits	110	1.3
	aviris_maine_sc10_raw	680x512x224, 12-bits	110	1.3
	aviris_yellowstone_sc0_raw	680x512x224, 16-bits	110	1.7
	aviris_yellowstone_sc3_raw	680x512x224, 16-bits	110	1.7
	aviris_yellowstone_sc11_raw	680x512x224, 16-bits	110	1.7
	Hyperion_Cuprite	256x1024x242, 12-bits	110	1.3
	crism_frt00009326_07_sc167	640x510x545, 12-bits	110	1.3
Multispectral Image	modis_A2001222.0835night	2030x1354x17, 12-bits	110	1.3

CCSDS 123.0-B-1: Comparisons

- Existing FPGA implementations
  - All use only sample-adaptive entropy encoder

	Existing FPGA implementations		This work (DSCAL/NKUA)
	HyLoC (IUMA/ULPGC, TAS Spain) OBPDC14	**NASA FL** (JPL, B&A) AHS14	
Sample ordering	BSQ	BIP	BIP
Number of bands used for prediction (P)	3	3	3
Max bit depth (D)	16	13	16
Image dimensions (N_x, N_y, N_z)	680, 512, 224	640, 32, 427	4096, 4096, 4096
Target device	Xilinx Virtex-5 (FX130T)	Xilinx Virtex-5 (SX50T)	Xilinx Virtex-5 (FX130T)
Device utilization Slices Slice registers Slice LUTs BRAMs/FIFOs DSP48E	 842 1535 2342 0 1	 NA 1586 12697 8 3	 2170 2667 5378 22 6
Max frequency (MHz)	134	40	110
Max throughput (MSamples/s)	**11**	**40**	**110**

CCSDS IP Cores: Verification

- Both IP cores verified extensively using:
  - RTL simulation-based verification (Mentor Graphics Questa)
  - FPGA-in-the-loop based verification
  - Code coverage measurements (100%)
  - A significant amount of test images (CCSDS images, random, corner cases, etc.)
- Both IP cores validated using a development board as hardware demonstrator
  - Alpha Data ADM-XRC-5T2 using XC5VFX130T -1 commercial equivalent of V5QV
- Both IP cores validated using standard bit-accurate golden reference models
  - CCSDS-IDC: University of Nebraska-Lincoln (UNL) software
  - CCSDS 123.0-B-1: ESA software implementation in C

Alpha Data ADM-XRC-5T2

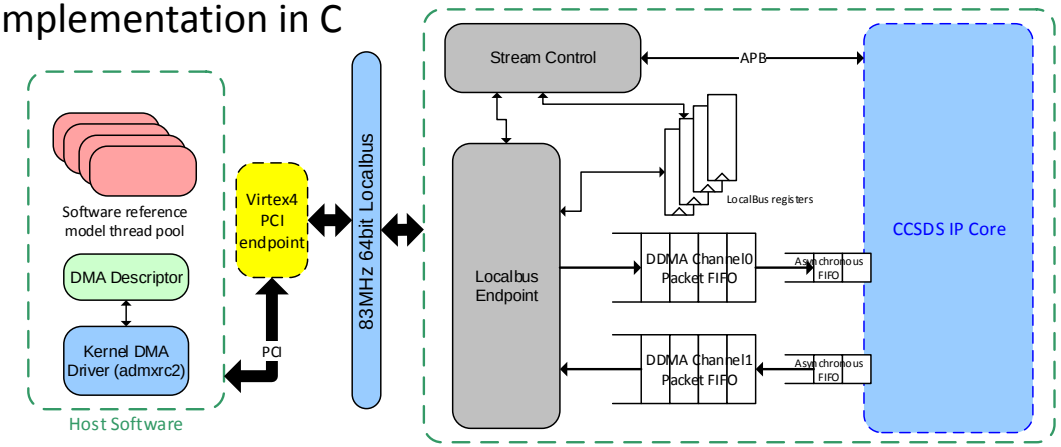

Demonstrator platform

Conclusions

- High-speed on-board image compression is necessary to handle remote sensing data challenges
- CCSDS developed image compression algorithms specifically for on-board use
- High-performance space-grade SRAM FPGA implementations were presented

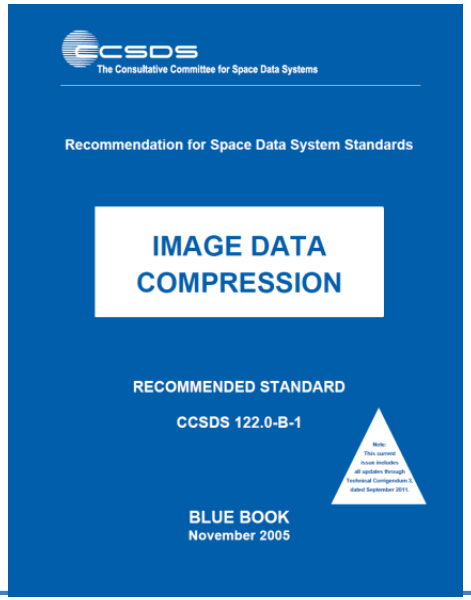

THANK YOU FOR YOUR ATTENTION!

nkran@di.uoa.gr

Questions?