BS ISO/IEC 18477-6:2016

BSI Standards Publication

Information technology — Scalable compression and coding of continuous-tone still images

Part 6: IDR Integer Coding

...making excellence a habit.™

National foreword

This British Standard is the UK implementation of ISO/IEC 18477-6:2016.

The UK participation in its preparation was entrusted to Technical Committee IST/37, Coding of picture, audio, multimedia and hypermedia information.

A list of organizations represented on this committee can be obtained on request to its secretary.

This publication does not purport to include all the necessary provisions of a contract. Users are responsible for its correct application.

Published by BSI Standards Limited 2016

ISBN 978 0 580 87699 8

ICS 35.040

Compliance with a British Standard cannot confer immunity from legal obligations.

This British Standard was published under the authority of the Standards Policy and Strategy Committee on 29 February 2016.

Amendments/corrigenda issued since publication

Date	Text affected

BS ISO/IEC 18477-6:2016

INTERNATIONAL STANDARD

ISO/IEC 18477-6

First edition
2016-02-01

Information technology — Scalable compression and coding of continuous-tone still images —

Part 6:
IDR Integer Coding

Technologies de l'information — Compression échelonnable et codage d'images plates en ton continu —

Partie 6: Codage de nombre entier par IDR

Reference number
ISO/IEC 18477-6:2016(E)

Contents

Page

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work. In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular the different approval criteria needed for the different types of document should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see www.iso.org/directives).

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights. Details of any patent rights identified during the development of the document will be in the Introduction and/or on the ISO list of patent declarations received (see www.iso.org/patents).

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation on the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the WTO principles in the Technical Barriers to Trade (TBT) see the following URL: Foreword - Supplementary information

The committee responsible for this document is ISO/IEC JTC 1, *Information technology*, SC 29, *Coding of audio, picture, multimedia and hypermedia information*.

ISO/IEC 18477 contains the following parts under the general title *Information technology — Scalable compression and coding of continuous-tone still images*:

— *Part 1: Scalable compression and coding of continuous-tone still images*

— *Part 2: Extensions for high dynamic range images*

— *Part 3: Box file format*

— *Part 6: IDR Integer Coding*

— *Part 7: HDR Floating-Point Coding*

— *Part 8: Lossless and Near-lossless Coding*

— *Part 9: Alpha Channel Coding*

The following parts are under preparation:

— *Part 4: Conformance testing*

— *Part 5: Reference software*

Introduction

This part of ISO/IEC 18477 specifies a coded codestream format for storage of continuous-tone high and low dynamic range photographic content. JPEG XT part 6 is a scalable image coding system supporting multiple component images consisting of integer samples of a bit precision between 9 and 16 bits. The format itself is based on the Box Based format specified in ISO/IEC 18477-3, which ensures that legacy applications conforming to Rec. ITU-T T.81 | ISO/IEC 10918-1 are able to reconstruct a lower quality, low dynamic range, eight bits per sample version of the image.

Today, the most widely used digital photography format, a minimal implementation of JPEG (specified in ITU Recommendation T.81 | ISO/IEC 10918-1), uses a bit depth of 8; each of the three channels that together compose an image pixel is represented by 8 bits, providing 256 representable values per channel. For more demanding applications, it is not uncommon to use a bit depth of 16, providing 65 536 representable values to describe each channel within a pixel, resulting on over $2,8 \times 10^{14}$ representable colour values.

Most common photo and image formats use an 8-bit or 16-bit unsigned integer value to represent some function of the intensity of each colour channel. While it might be theoretically possible to agree on one method for assigning specific numerical values to real world colours, doing so is not practical. Since any specific device has its own limited range for colour reproduction, the device's range may be a small portion of the agreed-upon universal colour range. As a result, such an approach is an extremely inefficient use of the available numerical values, especially when using only 8 bits (or 256 unique values) per channel. To represent pixel values as efficiently as possible, devices use a numeric encoding optimized for their own range of possible colours or gamut.

JPEG XT is primarily designed to provide coded data containing intermediate dynamic range and wide colour gamut content while simultaneously providing 8 bits per pixel low dynamic range images using tools defined in ISO/IEC 18477-1, which is itself a subset of Rec. ITU-T T.81 | ISO/IEC 10918-1. The goal is to provide a backwards compatible coding specification that allows legacy applications and existing toolchains to continue to operate on codestreams conforming to this part of ISO/IEC 18477.

JPEG XT has been designed to be backwards compatible to legacy applications while at the same time having a small coding complexity; JPEG XT uses, whenever possible, functional blocks of Rec. ITU-T T.81 | ISO/IEC 10918-1 to extend the functionality of the legacy JPEG Coding System. It is optimized for storage and transmission of intermediate dynamic range and wide colour gamut images while also enabling low-complexity encoder and decoder implementations.

This part of ISO/IEC 18477 is an extension of ISO/IEC 18477-1, a compression system for continuous tone digital still images which is backwards compatible with Rec. ITU-T T.81 | ISO/IEC 10918-1. That is, legacy applications conforming to Rec. ITU-T T.81 | ISO/IEC 10918-1 will be able to reconstruct streams generated by an encoder conforming to this part of ISO/IEC 18477, though will possibly not be able to reconstruct such streams in full dynamic range, full quality or other features defined in this Recommendation| International Standard.

This part of ISO/IEC 18477 is itself based on ISO/IEC 18477-3 which defines a box-based file format similar to other JPEG standards. The aim of this part of ISO/IEC 18477 is to provide a migration path for legacy applications to support, potentially in a limited way, coding of intermediate dynamic range images, that is images represented by sample values requiring 9 to 16 bits precision. While the legacy Rec. ITU-T T.81 | ISO/IEC 10918-1 already defines a coding mode for 12 bit sample precision, images encoded in this mode cannot be decoded by applications implementing only the 8 bit mode. Unlike the legacy standard, this part of ISO/IEC 18477 defines a scalable coding engine supporting all bit depths between 9 and 16 bits per sample while also staying compatible with legacy applications. Such applications will continue to work, but will only able to reconstruct an 8 bit standard low dynamic range (LDR) version of the full image contained in the codestream. This part of ISO/IEC 18477 specifies a coded file format, referred to as JPEG XT, which is designed primarily for storage and interchange of continuous-tone photographic content.

Information technology — Scalable compression and coding of continuous-tone still images —

Part 6:
IDR Integer Coding

1 Scope

This part of ISO/IEC 18477 specifies a coding format, referred to as JPEG XT, which is designed primarily for continuous-tone photographic content.

2 Normative references

The following documents, in whole or in part, are normatively referenced in this document and are indispensable for its application. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

ISO/IEC 18477-1:2015, *Information technology — Scalable compression and coding of continuous-tone still images — Part 1: Scalable compression and coding of continuous-tone still images*

ISO/IEC 18477-3:2015, *Information technology — Scalable compression and coding of continuous-tone still images —Part 3: Box-based file format*

Rec. ITU-T T.81 | ISO/IEC 10918-1, *Information technology — Digital compression and coding of continuous tone still images — Requirements and guidelines*

Rec. ITU-T BT.601, *Studio encoding parameters of digital television for standard 4:3 and wide screen 16:9 aspect ratios*

3 Terms and definitions, abbreviated terms, and symbols

3.1 Terms and definitions

For the purposes of this document, the following definitions apply.

3.1.1
AC coefficient
any DCT coefficient for which the frequency is not zero in at least one dimension

3.1.2
ASCII encoding
encoding of text characters and text strings according to ISO/IEC 10646-1

3.1.3
base decoding path
process of decoding legacy codestream and refinement data to the base image, jointly with all further steps until residual data is added to the values obtained from the residual codestream

3.1.4
base image
collection of sample values obtained by entropy decoding the DCT coefficients of the legacy codestream and the refinement codestream, and inversely DCT transforming them jointly

3.1.5
binary decision
choice between two alternatives

3.1.6
bitstream
partially encoded or decoded sequence of bits comprising an entropy-coded segment

3.1.7
block
8 × 8 array of samples or an 8 × 8 array of DCT coefficient values of one component

3.1.8
box
structured collection of data describing the image or the image decoding process embedded into one or multiple APP_{11} marker segments

Note 1 to entry: See ISO/IEC 18477-3:2015, Annex B for the definition of boxes.

3.1.9
byte
group of 8 bits

3.1.10
coder
embodiment of a coding process

3.1.11
coding
encoding or decoding

3.1.12
coding model
procedure used to convert input data into symbols to be coded

3.1.13
(coding) process
general term for referring to an encoding process, a decoding process, or both

3.1.14
compression
reduction in the number of bits used to represent source image data

3.1.15
component
two-dimensional array of samples having the same designation in the output or display device

Note 1 to entry: An image typically consists of several components, e.g. red, green, and blue.

3.1.16
continuous-tone image
image whose components have more than one bit per sample

3.1.17
DC coefficient
DCT coefficient for which the frequency is zero in both dimensions

3.1.18
decoder
embodiment of a decoding process

3.1.19
decoding process
process which takes as its input compressed image data and outputs a continuous-tone image

3.1.20
dequantization
inverse procedure to quantization by which the decoder recovers a representation of the DCT coefficients

3.1.21
discrete cosine transform
DCT
either the forward discrete cosine transform or the inverse discrete cosine transform

3.1.22
downsampling
procedure by which the spatial resolution of a component is reduced

3.1.23
encoder
embodiment of an encoding process

3.1.24
encoding process
process which takes as its input a continuous-tone image and outputs compressed image data

3.1.25
entropy-coded (data) segment
independently decodable sequence of entropy encoded bytes of compressed image data

3.1.26
entropy decoder
embodiment of an entropy decoding procedure

3.1.27
entropy decoding
lossless procedure which recovers the sequence of symbols from the sequence of bits produced by the entropy encoder

3.1.28
entropy encoder
embodiment of an entropy encoding procedure

3.1.29
entropy encoding
lossless procedure which converts a sequence of input symbols into a sequence of bits such that the average number of bits per symbol approaches the entropy of the input symbols

3.1.30
grayscale image
continuous-tone image that has only one component

3.1.31
high dynamic range
image or image data comprised of more than eight bits per sample

3.1.32
Huffman decoder
embodiment of a Huffman decoding procedure

3.1.33
Huffman decoding
entropy decoding procedure which recovers the symbol from each variable length code produced by the Huffman encoder

3.1.34
Huffman encoder
embodiment of a Huffman encoding procedure

3.1.35
Huffman encoding
entropy encoding procedure which assigns a variable length code to each input symbol

3.1.36
intermediate dynamic range
image or image data comprised of more than eight bits per sample

3.1.37
joint photographic experts group
JPEG
informal name of the committee which created this part of ISO/IEC 18477

Note 1 to entry: The “joint” comes from the ITU-T and ISO/IEC collaboration.

3.1.38
legacy codestream
collection of markers and syntax elements defined by Rec. ITU-T T.81 | ISO/IEC 10918-1 bare any additional syntax elements defined by the family ISO/IEC 18477 standards, i.e. the legacy codestream consists of the collection of all markers except those APP_{11} markers that describe JPEG XT boxes by the syntax defined in ISO/IEC 18477-3:2015, Annex A

3.1.39
legacy decoding path
collection of operations to be performed on the entropy coded data as described by Rec. ITU-T T.81 | ISO/IEC 10918-1 jointly with the Legacy Refinement scans before this data is merged with the residual data to form the final output image

3.1.40
legacy decoder
embodiment of a decoding process conforming to Rec. ITU-T T.81 | ISO/IEC 10918-1, confined to the lossy DCT process and the baseline, sequential, or progressive modes, decoding at most four components to eight bits per component

3.1.41
legacy image
arrangement of sample values as described by applying the decoding process described by Rec. ITU-T T.81 | ISO/IEC 10918-1 on the entropy coded data as defined by said standard

3.1.42
lossless
descriptive term for encoding and decoding processes and procedures in which the output of the decoding procedure(s) is identical to the input to the encoding procedure(s)

3.1.43
lossless coding
mode of operation which refers to any one of the coding processes defined in this part of ISO/IEC 18477 in which all of the procedures are lossless

Note 1 to entry: See ISO/IEC 18477-8.

3.1.44
lossy
descriptive term for encoding and decoding processes which are not lossless

3.1.45
low dynamic range
image or image data comprised of data with no more than eight bits per sample

3.1.46
marker
two-byte code in which the first byte is hexadecimal FF and the second byte is a value between 1 and hexadecimal FE

3.1.47
marker segment
marker together with its associated set of parameters

3.1.48
pixel
collection of sample values in the spatial image domain having all the same sample coordinates, e.g. a pixel may consist of three samples describing its red, green, and blue value

3.1.49
precision
number of bits allocated to a particular sample or DCT coefficient

3.1.50
procedure
set of steps which accomplishes one of the tasks which comprise an encoding or decoding process

3.1.51
quantization value
integer value used in the quantization procedure

3.1.52
quantize
act of performing the quantization procedure for a DCT coefficient

3.1.53
residual decoding path
collection of operations applied to the entropy coded data contained in the residual data box and residual refinement scan boxes up to the point where this data is merged with the base image to form the final output image

3.1.54
residual image
extension image
sample values as reconstructed by inverse quantization and inverse DCT transformation applied to the entropy-decoded coefficients described by the residual scan and residual refinement scans

3.1.55
residual scan
additional pass over the image data invisible to legacy decoders which provides additive and/or multiplicative correction data of the legacy scans to allow reproduction of high dynamic range or wide colour gamut data

3.1.56
refinement scan
additional pass over the image data invisible to legacy decoders which provides additional least significant bits to extend the precision of the DCT transformed coefficients

3.1.57
sample
one element in the two-dimensional image array which comprises a component

3.1.58
sample grid
common coordinate system for all samples of an image

Note 1 to entry: The samples at the top left edge of the image have the coordinates (0,0), the first coordinate increases towards the right, the second towards the bottom.

3.1.59
scan
single pass through the data for one or more of the components in an image

3.1.60
scan header
marker segment that contains a start-of-scan marker and associated scan parameters that are coded at the beginning of a scan

3.1.61
superbox
box that carries other boxes as payload data

3.1.62
table specification data
coded representation from which the tables used in the encoder and decoder are generated and their destinations specified

3.1.63
(uniform) quantization
procedure by which DCT coefficients are linearly scaled in order to achieve compression

3.1.64
upsampling
procedure by which the spatial resolution of a component is increased

3.1.65
vertical sampling factor
relative number of vertical data units of a particular component with respect to the number of vertical data units in the other components in the frame

3.1.66
zero byte
0x00 byte

3.1.67
zig-zag sequence
specific sequential ordering of the DCT coefficients from (approximately) lowest spatial frequency to highest

3.2 Symbols

X Width of the sample grid in positions

Y Height of the sample grid in positions

Nf Number of components in an image

$s_{i,x}$ Subsampling factor of component in horizontal direction

$s_{i,y}$ Subsampling factor of component in vertical direction

H_i Subsampling indicator of component in the frame header

V_i Subsampling indicator of component in the frame header

$v_{x,y}$ Sample value at the sample grid position x,y

R_h Additional number of DCT coefficient bits represented by refinement scans in the legacy decoding path, $8+R_h$ is the number of non-fractional bits (i.e. bits in front of the "binary dot") of the output of the inverse DCT process in the legacy decoding path.

R_r Additional number of DCT coefficient bits represented by refinement scans in the residual decoding path. $P+R_r$ is the number of non-fractional bits of the output of the invers DCT process in the residual decoding path, where P is the frame-precision of the residual image as recorded in the frame header of the residual codestream.

R_b Additional bits in the HDR image. 8+Rb is the sample precision of the reconstructed HDR image.

3.3 Abbreviated terms

For the purposes of this part of ISO/IEC 18477, the following abbreviated terms apply.

ASCII American Standard Code for Information Interchange

LSB Least Significant Bit

MSB Most Significant Bit

HDR High Dynamic Range

IDR Intermediate Dynamic Range

LDR Low Dynamic Range

TMO Tone Mapping Operator

DCT Discrete Cosine Transformation

4 Conventions

4.1 Conformance language

This part of ISO/IEC 18477 consists of normative and informative text.

Normative text is that text which expresses mandatory requirements. The word "shall" is used to express mandatory requirements strictly to be followed in order to conform to this part of ISO/IEC 18477 and from which no deviation is permitted. A conforming implementation is one that fulfils all mandatory requirements.

Informative text is text that is potentially helpful to the user, but not indispensable and can be removed, changed, or added editorially without affecting interoperability. All text in this part of ISO/IEC 18477 is normative, with the following exceptions: The Introduction, any parts of the text that are explicitly labelled as "informative", and statements appearing with the preamble "NOTE" and behaviour described using the word "should". The word "should" is used to describe behaviour that is encouraged but is not required for conformance to this part of ISO/IEC 18477.

The keywords "may" and "need not" indicate a course of action that is permissible in a conforming implementation.

The keyword “reserved” indicates a provision that is not specified at this time, shall not be used, and may be specified in the future. The keyword “forbidden” indicates “reserved” and in addition indicates that the provision will never be specified in the future.

4.2 Operators

NOTE Many of the operators used in this part of ISO/IEC 18477 are similar to those used in the C programming language.

4.2.1 Arithmetic operators

+ Addition

− Subtraction (as a binary operator) or negation (as a unary prefix operator)

* Multiplication

/ Division without truncation or rounding

4.2.2 Logical operators

|| Logical OR

&& Logical AND

! Logical NOT

∈ x ∈ {A, B} is defined as (x == A || x = = B)

∉ x ∉ {A, B} is defined as (x != A && x != B)

4.2.3 Relational operators

> Greater than

>= Greater than or equal to

< Less than

<= Less than or equal to

== Equal to

!= Not equal to

4.2.4 Precedence order of operators

Operators are listed below in descending order of precedence. If several operators appear in the same line, they have equal precedence. When several operators of equal precedence appear at the same level in an expression, evaluation proceeds according to the associativity of the operator either from right to left or from left to right.

Operators	**Type of operation**	**Associativity**
(), [ ], .	Expression	Left to Right
−	Unary negation	
*, /	Multiplication	Left to Right
+, −	Addition and Subtraction	Left to Right
<, >, <=, >=	Relational	Left to Right

4.2.5 Mathematical functions

x	Ceil of x. Returns the smallest integer that is greater than or equal to x.
x	Floor of x. Returns the largest integer that is lesser than or equal to x.
|x|	Absolute value, is −x for x < 0, otherwise x.
sign(x)	Sign of x, 0 if x is zero, +1 if x is positive, −1 if x is negative.
clamp(x,min,max)	Clamps x to the range [min,max]: Returns min if x < min, max if x > max or otherwise x.
x^a	Raises the value of x to the power of a. x is a non-negative real number, a is a real number. x^a is equal to exp[a × log(x)] where exp is the exponential function and log() the natural logarithm. If x is 0 and a is positive, x^a is defined to be 0.

5 General

The purpose of this Clause is to give an informative overview of the elements specified in this part of ISO/IEC 18477. Another purpose is to introduce many of the terms which are defined in Clause 3. These terms are printed in *italics* upon first usage in this Clause.

There are three elements specified in this part of ISO/IEC 18477.

a) An *encoder* is an embodiment of an *encoding process*. An encoder takes as input *digital source image data* and *encoder specifications*, and by means of a specified set of *procedures* generates as output *codestream*.

b) A *decoder* is an embodiment of a *decoding process*. A decoder takes as input a codestream, and by means of a specified set of procedures generates as output *digital reconstructed image data*.

c) The *codestream* is a compressed image data representation which includes all necessary data to allow a (full or approximate) reconstruction of the sample values of a digital image. Additional data might be required that define the interpretation of the sample data, such as colour space or the spatial dimensions of the samples.

5.1 High level overview on JPEG XT ISO/IEC 18477-6

This part of ISO/IEC 18477 allows lossy coding of intermediate dynamic range of photographic images in a way that is backwards compatible to Rec. ITU-T T.81 | ISO/IEC 10918-1. Decoders compliant to the latter standard will be able to parse codestreams conforming to this part of ISO/IEC 18477 correctly, albeit in less precision, with a limited dynamic range, and loss in sample bit precision.

This part of ISO/IEC 18477 includes multiple tools to reach the above functionality, defined in Annex B and on. It is itself based on the file format specified in ISO/IEC 18477-3 and uses the syntax elements and tools defined there. While ISO/IEC 18477-3 only defines a syntax, this part of ISO/IEC 18477 extends the syntax of ISO/IEC 18477-3 to allow the representation of intermediate dynamic range images. It also defines a decoding process that reconstructs sample values from conforming files. A high level overview on both the syntax and the decoding will be given in this section.

The syntax of an ISO/IEC 18477-6 compliant codestream is specified in ISO/IEC 18477-3, that is, this part of ISO/IEC 18477 uses a syntax element denoted as "box" to annotate its syntactical elements. The definition of the box syntax element is not repeated here, and readers are referred to ISO/IEC 18477-3 for further details. Additional boxes besides those already specified in ISO/IEC 18477-3 are defined here, in specific, the Residual Data box, the Refinement Data box, and the Residual Refinement box at the top level of the file, and various sub-boxes of the Merging Specification box defining the decoding process. The Merging Specification Superbox is already defined in ISO/IEC 18477-3, all additional box types are specified in Annex B.

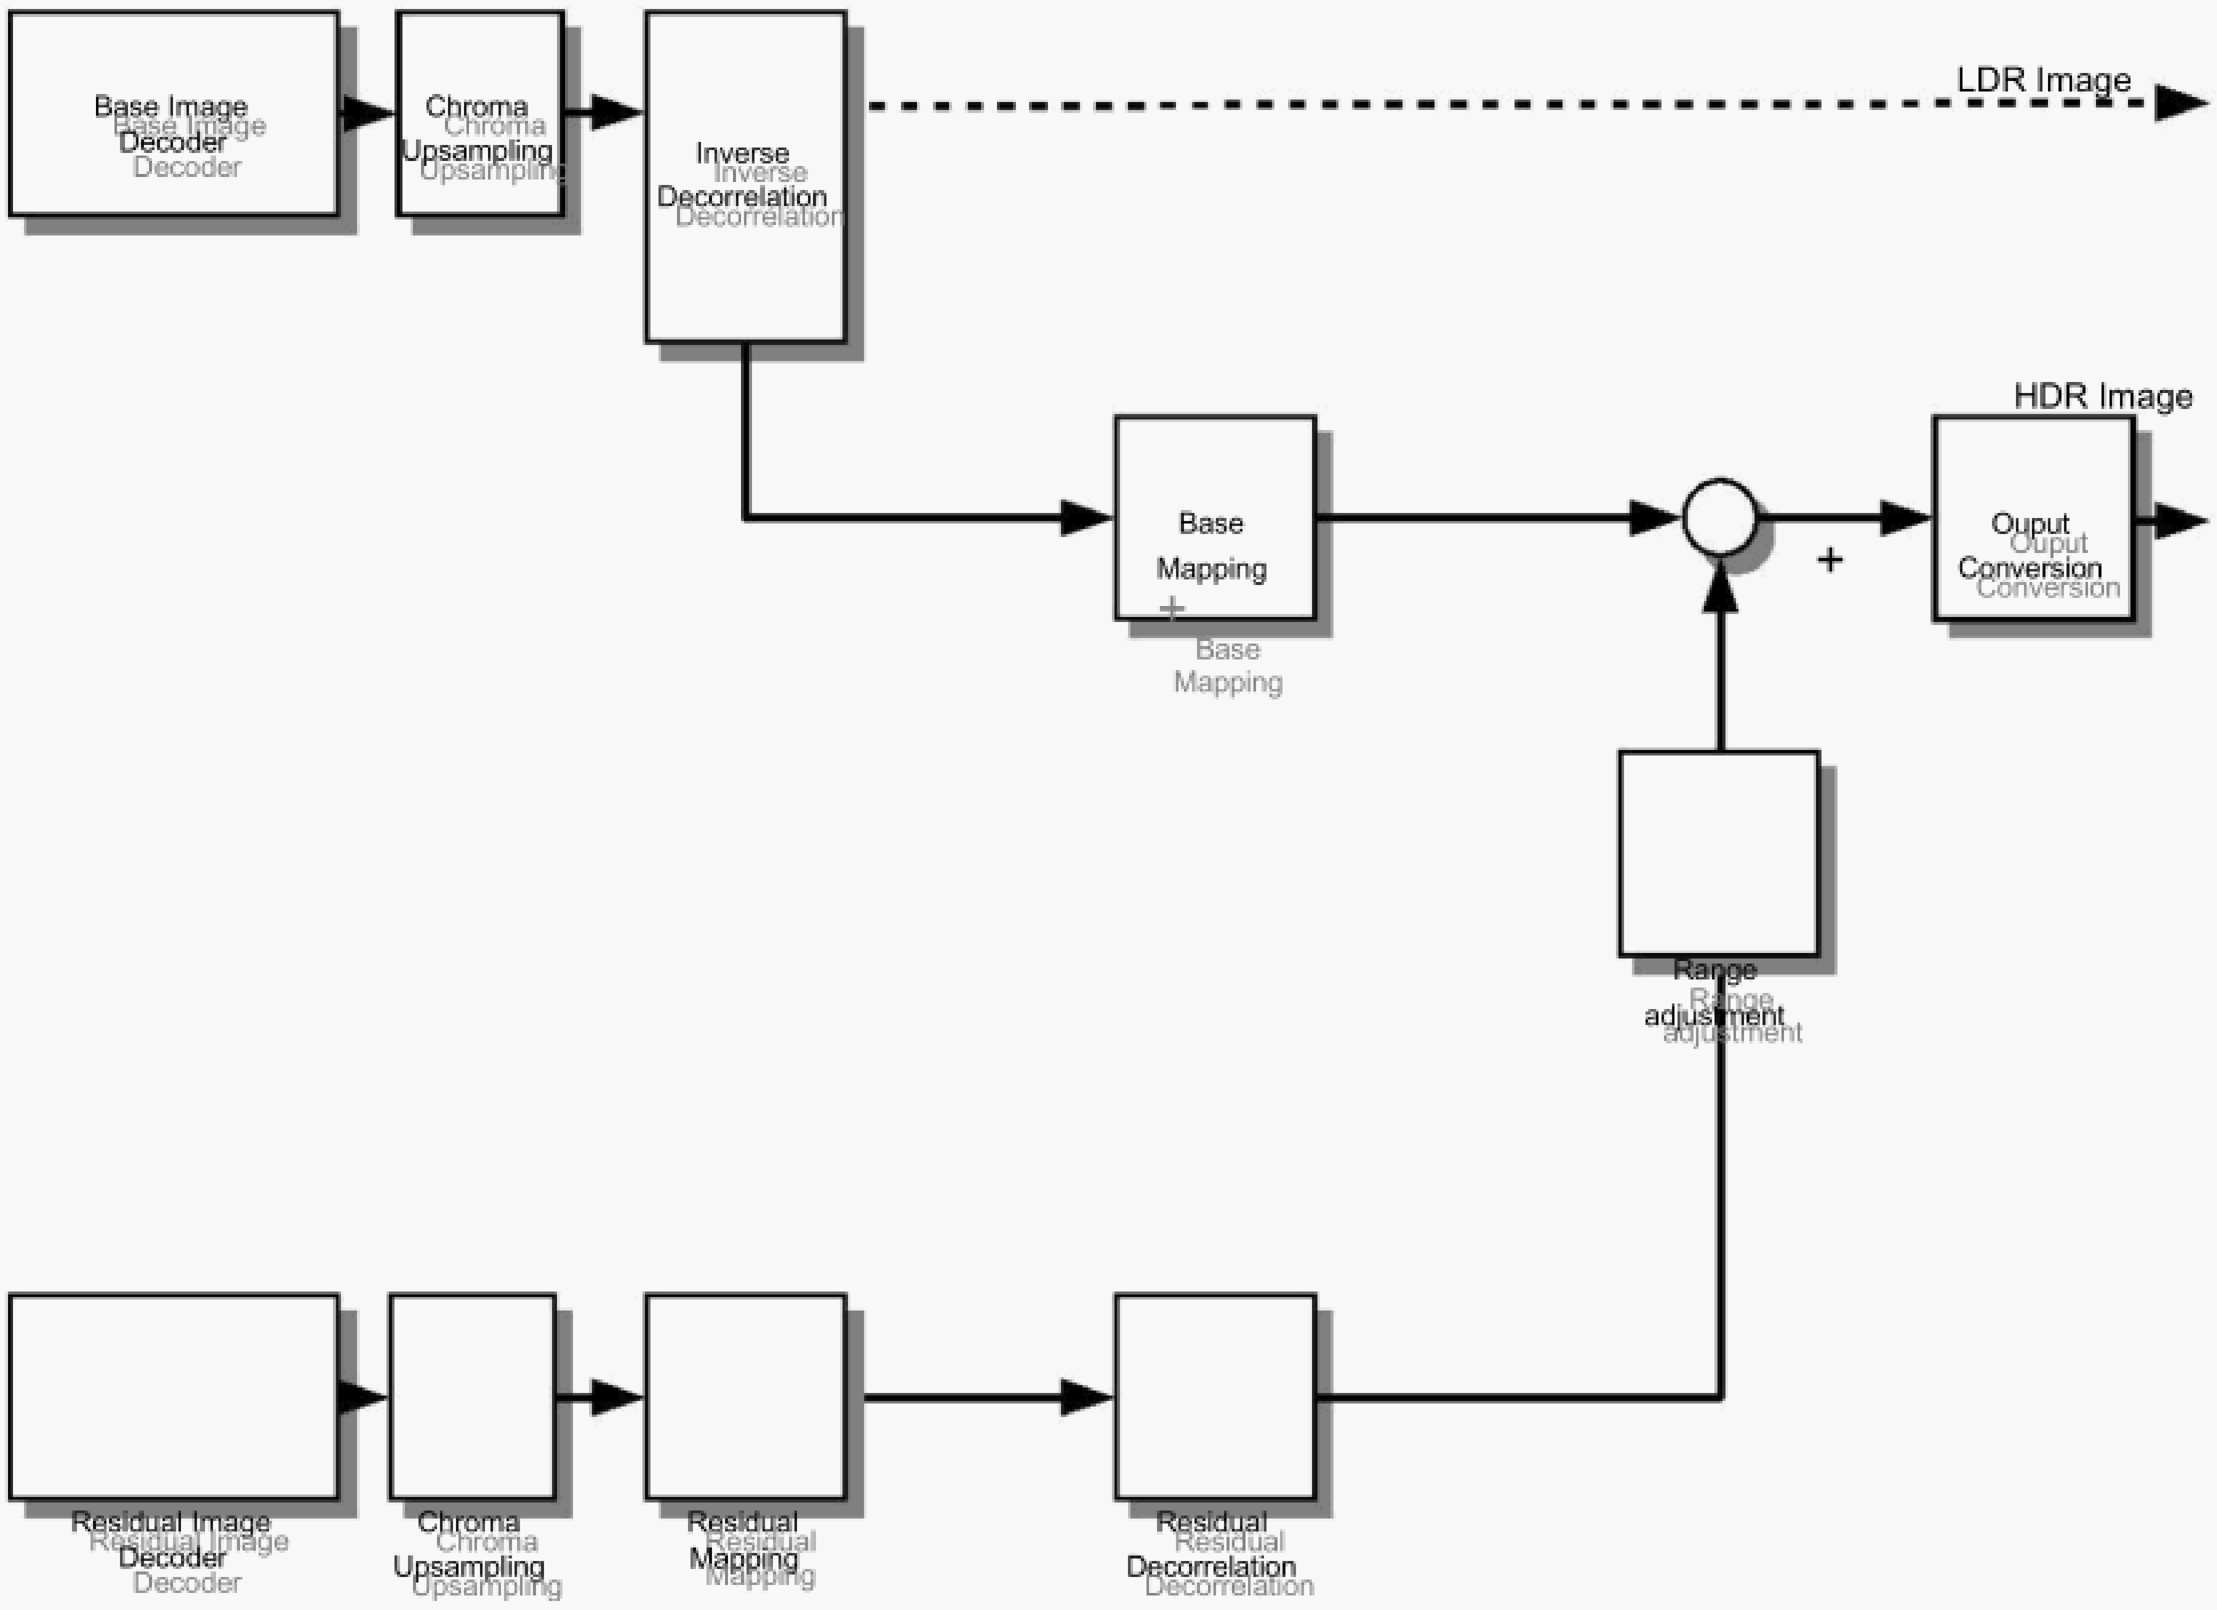

Figure 1 — Overview on the decoding process

This part of ISO/IEC 18477 extends the legacy decoding process by two mechanisms (see Figure 1). The **Refinement Scan**, specified in Annex D, increases the bit precision of the DCT coefficients, i.e. it operates in the DCT domain. The mechanism used here is very similar to that of Subsequent Approximation scans specified in Rec. ITU-T T.81 | ISO/IEC 10918-1:1994, Annex G: A baseline, extended or progressive Huffman scan as defined in Annex G of the legacy standard defines the 12 most significant bits of the DCT coefficients. These initial scans are represented in the legacy codestream and are visible for any ISO/IEC 18477-1 compliant decoder. Refinement scans decode now into up to four additional least significant bits in the same way subsequent approximation within Rec. ITU-T T.81 | ISO/IEC 10918-1 decode least significant bits of a progressive scan pattern. The difference between Refinement scans and Subsequent Approximation scans is only that in the latter case the number of least significant bits is annotated in the scan header of the legacy codestream, whereas Refinement scans are hidden from legacy applications and do not alter the scan header of the legacy codestream. Their number is indicated in the Refinement Specification box within the Merging Specification box and not in the legacy codestream.

While Refinement Scans extend the bit precision within the DCT domain by up to four bits and hence allow backwards compatible coding of images of up to 12 bits sample precision, **Residual Scans** extend the sample precision in the spatial (image) domain. While the entropy coded data of Residual Scans is hidden in the Residual Data box from legacy applications, its decoding process is **identical** to that of the legacy data: A baseline, extended or progressive Huffman scan decodes the data in the Residual Data box to DCT coefficients, inverse quantization and inverse Discrete Cosine Transformation (DCT) compute from these coefficients the **residual image data**. An image merging process, defined in Annex A, computes from the **precursor image** reconstructed from the base image and the **residual image** a final IDR output image. This merging process first performs chroma upsampling to reconstruct a single sample on each point of the sample grid of the base image. Chroma upsampling is specified in ISO/IEC 18477-1:2015, Annex A. It then converts the colour space of the base image first from YCbCr into the Rec. ITU-T BT.601 colourspace, followed by an additional linear transformation transforming the Rec. ITU-T BT.601 primary colours into the primary colours of the target IDR colourspace. For practical reasons, these two transformations are combined into a single linear transformation matrix. This linear transformation is followed by a nonlinear point transformation acting separately on each of the output channels sample by sample. This point transformation can be either specified by a parametric curve or by an explicit lookup table. The output of this decoding path is transformed again by an optional colour transformation forming the **precursor image** which represents a rough imprecise approximation of the final IDR image, already in the correct IDR output colour space.

Processing continues with the decoding of the **Residual Image**: DCT coefficients of the residual image are decoded from the information in the Residual Data box, and their bit precision is extended by additional refinement scans decoded from the data in the **Residual Refinement box**. Processing proceeds with inverse quantization and inverse DCT transformation. The output undergoes chroma upsampling to generate a single sample per sample grid coordinate. The next processing step performs a nonlinear point transformation on each of the reconstructed channels, separately for each sample, resulting in an error image in a YCbCr type of colour space. Samples undergo then an inverse linear decorrelation transformation to map the sample values from the intermediate YCbCr colourspace into the target colourspace. This transformation is typically identical to the transformation matrix in the base decoding path, but does not need to be. The result of this operation is the **residual image.**

To form the final output image, sample values of the **precursor image** and the **residual image** are added together, plus an offset to make the residual image symmetric around zero. Results are then clamped to the range of the intermediate range output image.

The detailed specification of the decoding and merging process is found in Annex A.

5.2 Profiles

The Profiles define the implementation of a particular technology within the functional blocks of Figure 1. The profiles are described in Annex E.

5.3 Encoder requirements

There is **no requirement** in this part of ISO/IEC 18477 that any encoder shall support all profiles. An encoder is only required to meet the compliance tests and to generate the codestream according to the syntax and to limit the coding parameters to those valid within the profile it conforms to. Profiles are defined in Annex E.

5.4 Decoder requirements

A decoding process converts compressed image data to reconstructed image data. It **may** follow the decoding operation specified in the Recommendation | International Standard and ISO/IEC 18477-1 to generate an LDR image from the legacy codestream, and it **shall** follow the operations in **this** part of ISO/IEC 18477 to decode an IDR image from the data in the full file. The decoder shall parse the codestream syntax to extract the parameters, the residual image and the legacy image. The parameters shall be used to merge the residual image with the base image into the reconstructed IDR Image.

In order to comply with this part of ISO/IEC 18477, a decoder

a) **may convert** a codestream conforming to this part of ISO/IEC 18477 **without considering the information in any box** into to a low dynamic range image, and

b) **shall** convert a conforming codestream within the profile it claims to be conforming to into an intermediate dynamic range image.

Annex A
(normative)

Encoding and decoding process

A.1 Decoding process (normative)

The decoding process relies on a layered approach to extend JPEG's capabilities. The encoder decomposes an IDR image into a base layer, which consists of a tone-mapped version of the IDR image and an IDR residual layer. In addition to the residual layer, the codestream includes a description of an approximate inverse tone mapping operation that allows the decoder to reconstruct from the LDR image an approximate IDR image; the errors of this approximation process are corrected by the residual codestream included the residual data box and residual refinement box (see Annex B). Both the description of the tone mapping and the residual image are included in boxes invisible to legacy decoders. Such decoders will thus only see the tone mapped LDR image. While the base image complies to ISO/IEC 18477-1 and thus supports only the 8-bit extended or baseline, extended or progressive Huffman modes, the residual image may optionally be encoded in the 12-bit Huffman or progressive modes.

Figure A.1 illustrates the functionality of a compliant decoder:

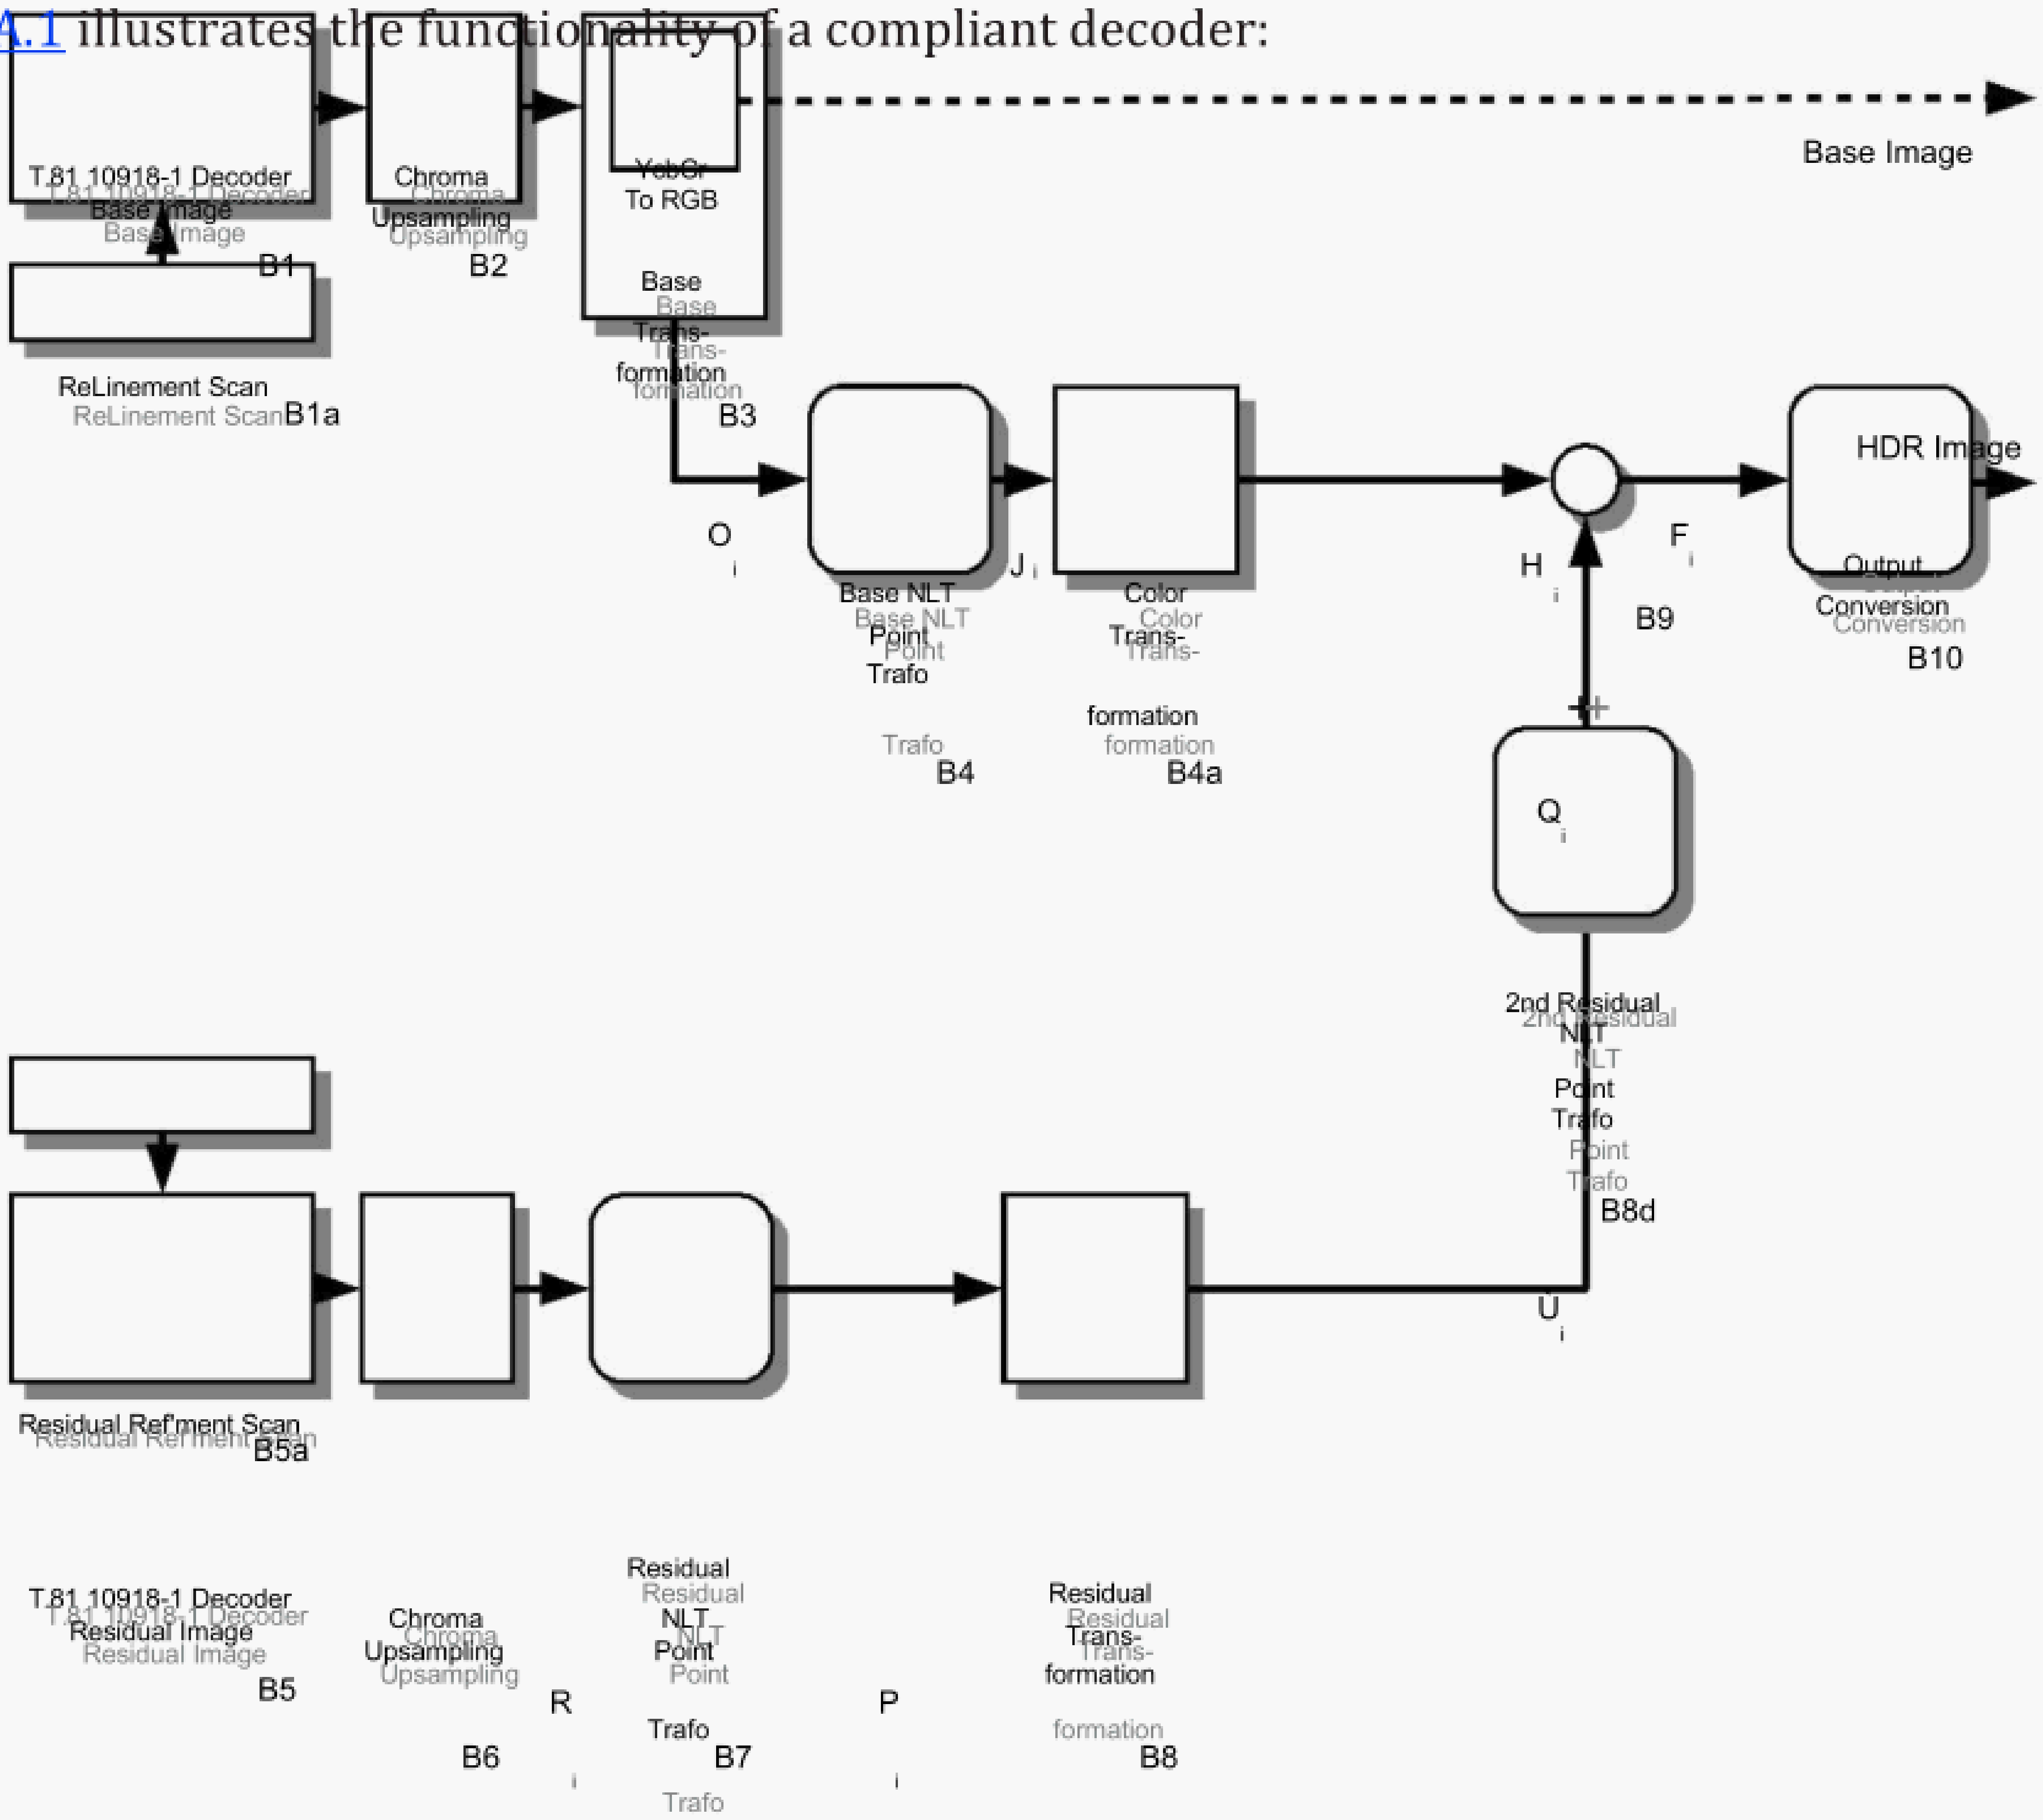

NOTE Bold lines carry three (or one, for greyscale) components. Round boxes implement point-transformations, square boxes (except B1, B1a, B5, B5a) multiplications by 3 × 3 matrices. Letters denote signal names.

Figure A.1 — High level overview of the decoding process of a compliant decoder

This subclause specifies the reconstruction process of an intermediate dynamic range image from a LDR image and a residual image decoders shall follow. This process consists of the following steps, see also Figure A.1.

— In steps B1 and B1a, reconstruct the base image from legacy codestream and the refinement codestream if a Refinement Data box is present. Refinement coding is specified in Annex D.

— In step B1 and B1a, apply the Inverse Quantization and Inverse Discrete Cosine Transformation as in Rec. ITU-T T.81 | ISO/IEC 10918-1.

— In step B2, the upsampling process specified in ISO/IEC 18477-1:2015, Annex A shall be followed to generate samples for all positions on the sample grid.

— In step B3, the linear transformation as selected by the Base Transformation box defined in Annex B shall be applied to inversely decorrelate the image components. Table C.1 defines which transformation to pick. The output of this block consists of either one or three samples per grid point O_i, depending on the number of components in the base image. The output of this transformation is rounded to integers and clipped to $[0,2^{R_h+8}-1]$ where R_h is the number of refinement scans in the base image (see Annex D).

— In step B4, a nonlinear point transformation shall be applied to each of the output components O_i. This process is selected according to the Base nonlinear Point Transformation subbox of the Merging Specification box, implementing the L_i Luts of Figure 1 and following the specifications of ISO/IEC 18477-3:2015, Annex C. The outputs of this process are the predicted high dynamic range samples J_i. As above, i = 1..Nf.

— In step B4a, a colour transformation is applied to the input values J_i resulting in the output pixel values H_i. The transformation is selected by the Colour Transformation subbox of the Merging Specification box, which selects one of the transformations defined in Annex C. If Nf equals 1, no transformation is performed.

— In steps B5 and B5a, the residual image shall be reconstructed from the data contained in the Residual Codestream box and the Residual Refinement box. The codestream contained in this box follows the specifications defined in Rec. ITU-T T.81 | ISO/IEC 10918-1. If a Residual Refinement box is present, the precision of the samples of the residual codestream shall be extended by refinement coding as specified in Annex D. The number of components of the residual image shall be equal to the number of components signalled in the base image.

— In steps B5 and B5a, apply Inverse Quantization and Inverse Discrete Cosine Transformation as in Rec. ITU-T T.81 | ISO/IEC 10918-1.

— In step B6, residual data is upsampled to the common sample grid following the specification of ISO/IEC 18477-1:2015, Annex A. The result of this operation is the residual image R_i.

— In step B7, apply a point transformation to R_i that is defined by the Residual nonlinear Point Transformation subbox of the Merging Specification box. The outputs of this operation are one or three residual sample values per pixel denoted by P_i.

— The outputs of the residual decoding process are one or three integer sample values P_i per sample grid point.

— In step B8, an inverse colour decorrelation shall be applied to the P_i data. The transformation is defined by the Residual Transformation box, which is a subboxes of the Merging Specification Box. The outputs of this process are the inversely decorrelated prediction errors Ω_i.

— In step B8d, a secondary nonlinear point transformation selected by the **Secondary Residual Nonlinear Point Transformation box** is applied. This box selects for each component nonlinearity, either a Floating Point Lookup box or a Parametric Curve box that maps the inversely decorrelated residual errors Ω_i into the final residual errors Q_i. If this box is not present, Q_i is set to Ω_i.

— In step B9, the intermediate dynamic range output F_i, i.e. the output of the decoding process, is reconstructed from H_i, the predicted high dynamic range signal, and the inversely decorrelation prediction errors Q_i. The output bit precision R_b is taken from the Output Conversion subbox of the Merging Specification box defined in Annex B. For this, compute

$$F_i := H_i + Q_i - 2^{R_b+8-1}$$

— The values F_i shall be rounded to the nearest integer, clamped to the output range of $[0, 2^{R_b}-1]$ and then form the final output of the decoding process.

A.2 Encoding process (informative)

This part of ISO/IEC 18477 does not define a normative encoding process. Any encoding process that generates a file format that is compliant to this part of ISO/IEC 18477 is acceptable.

The input to an embodiment of a JPEG XT encoder as described in this subclause consists of an image pair, an IDR image of the full dynamic range and an LDR image a legacy JPEG decoder would reconstruct the codestream to. IDR and LDR image are typically related by a tone mapping operation (TMO), though the selection of a TMO is outside the scope of this part of ISO/IEC 18477. This tone mapping step from IDR to LDR image is logically performed outside of the JPEG XT encoder, though specific encoder implementations may perform this step, i.e. the generation of an LDR image, as part of their implementation.

Two possible encoder mechanisms will be described in this subclause. The **closed loop encoder** depicted in Figure A.2 first estimates a suitable inverse tonemapping from the input image and suitable image parameters or two input images if desired. This inverse tonemapping is encoded in an Integer Table Lookup box or a Parametric Curve box as necessary, and the Base Nonlinear Point Transformation box of the Merging Specification Box points to the Table Lookup or Parametric Curve box. The Colour Transformation box and Base Transformation box select a suitable transformation from the IDR colour space and the colour space of legacy JPEG applications, which is given by Rec. ITU-T BT.601. Depending on the desired colour spaces and compression performance, the Colour Transformation box may be dropped signalling an identity transformation.

Processing of data is now performed as follows: Input IDR data is first transformed into an alternative colourspace by the inverse of the linear decorrelation given by the Colour Transformation box if present. Computation of the inverse matrix is up to the encoder. Data is then clamped to range if necessary and an inverse of the nonlinear point transformation lookup process, given by ISO/IEC 18477-3:2015, Annex C is applied. As above, computing an approximate inverse of the Table Lookup or Parametric Curve is up to the encoder and not specified by this part of ISO/IEC 18477.

The output of this process is decorrelated by the inverse of the base correlation matrix, computing samples in the YCbCr colourspace defined by Rec. ITU-T BT.601. The inverse of the Base Transformation is again computed by the encoder and not signalled; only the transformation matrix for the decoder process is included in the codestream if different from the standard YCbCr to RGB transformation. The outputs of this decorrelation process are first clamped to range, then subsampled if desired and transformed by the DCT, quantized and encoded.

To compute the residual image in the **closed loop encoder design**, the DCT coefficients computed for the base image are firist dequantized, then inversely DCT decoded to the **precursor image** as defined in A.1. That is, the encoding process includes a partial base image decoder implementation that mimics the decoder. The outputs of this partial decoding process are the sample values of the precursor image. The residual image is then computed by subtracting the precursor sample values from the IDR input image sample values and adding an offset of 2^{R_b+8-1}.

The residual image values are then inversely decorrelated by the inverse of the Residual Transformation and mapped into the residual coding domain by the inverse of the Residual Point transformation. Computation of suitable inverses is again up to the encoder; only the decoding matrix and the decoding nonlinear point transformations are signalled in the codestream. The resulting data are then subsampled if desired, DCT transformed, and entropy coded by the baseline, extended or progressive Huffman process. Refinement scans may be added if necessary.

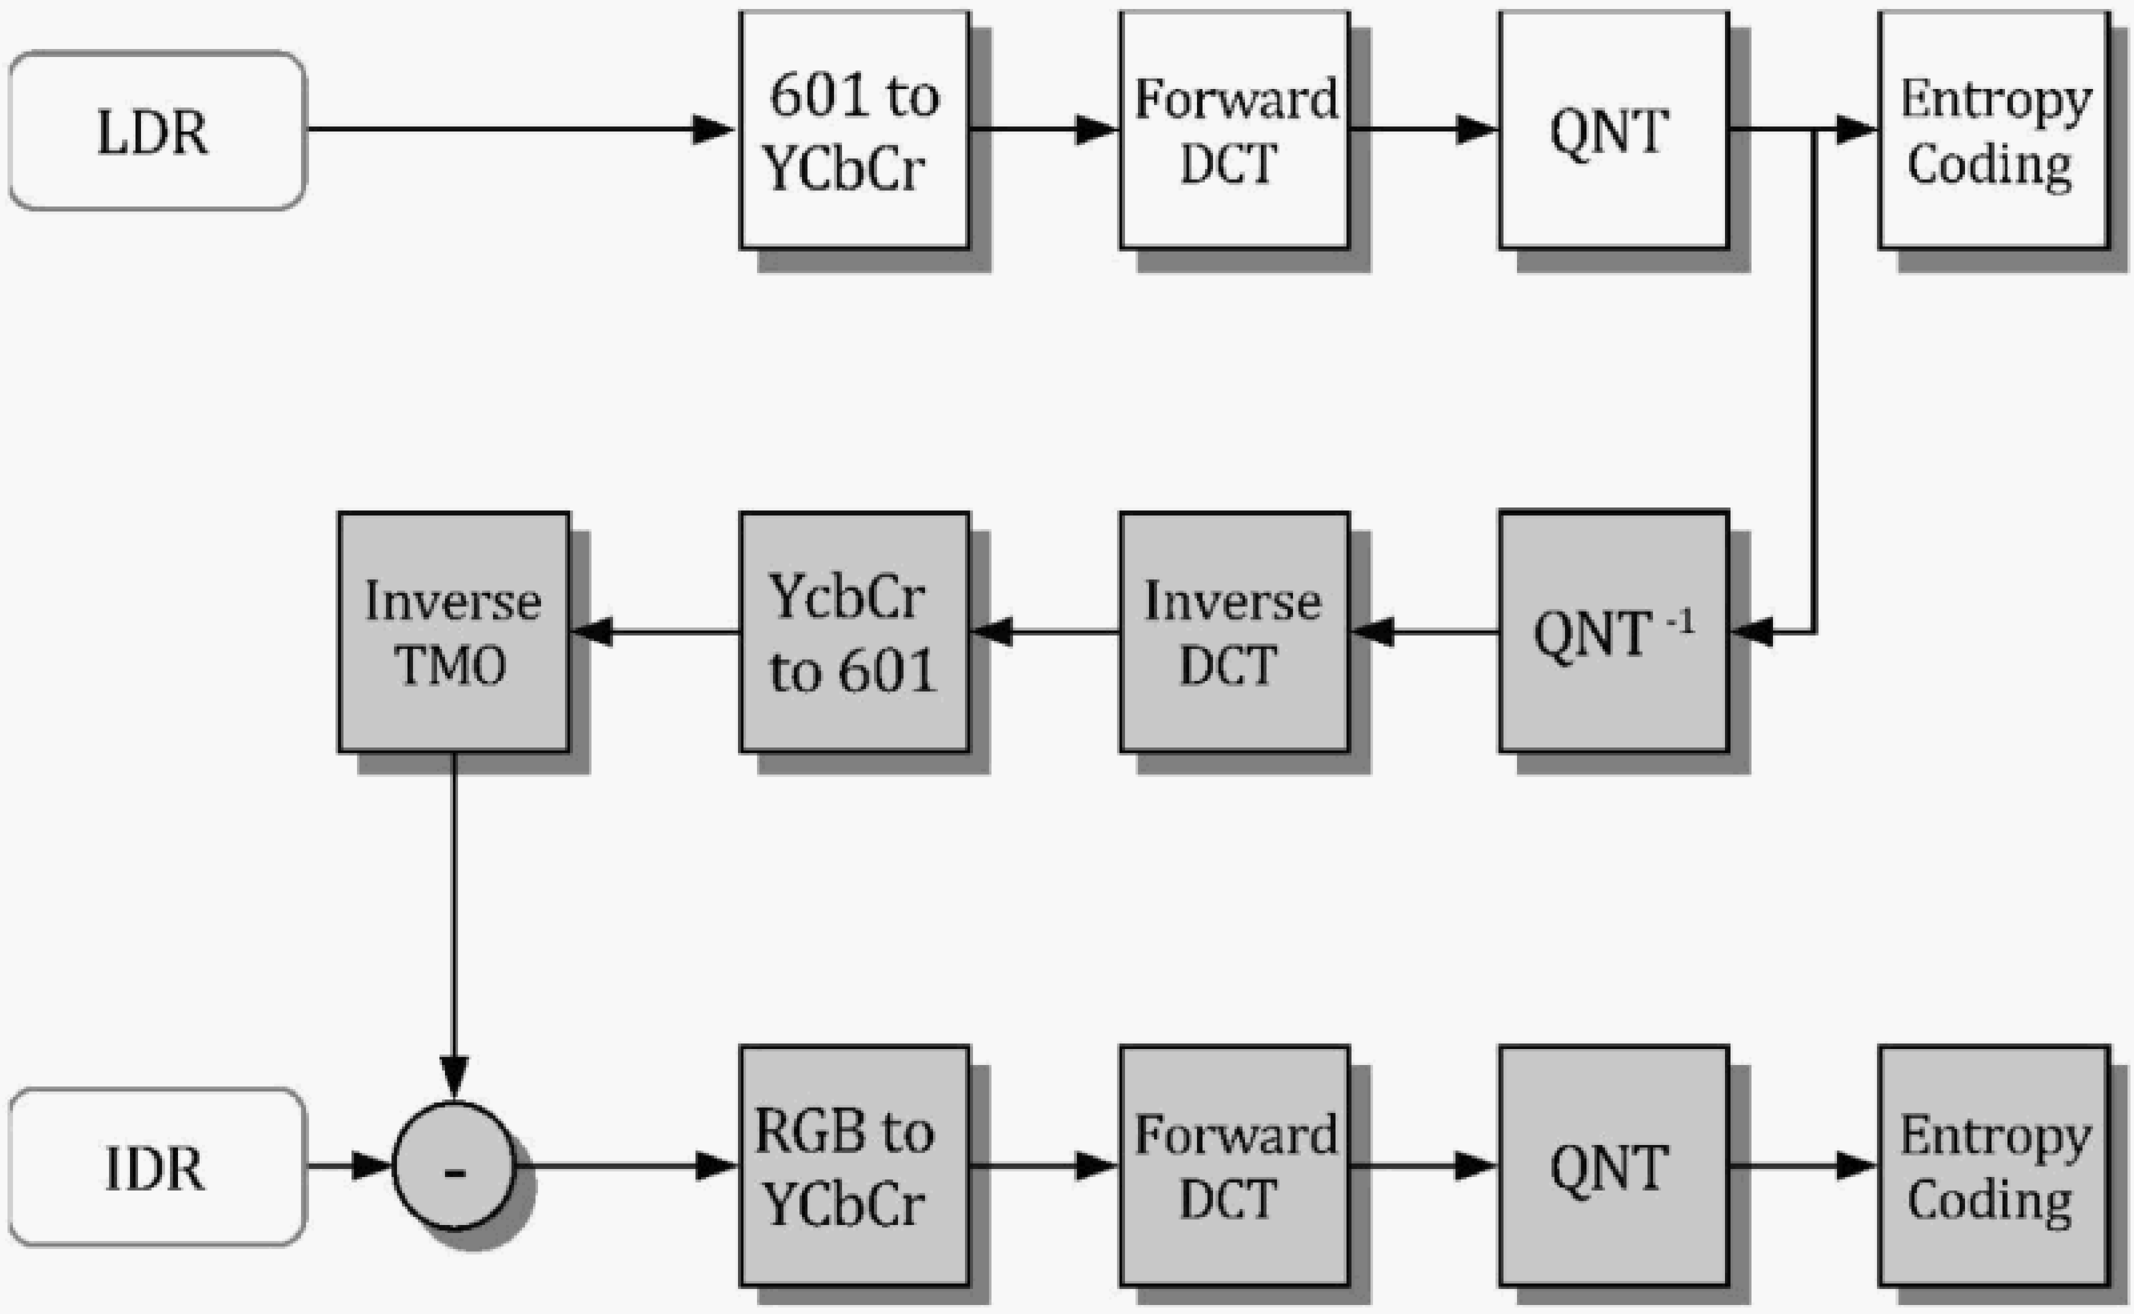

Figure A.2 — High level overview of the closed loop encoder

The **open loop** encoder only reproduces a smaller part of the decoding process and does not perform a full inverse DCT and inverse quantization step as the closed loop encoder. It is depicted in Figure A.3. Computation of the residual image is however abbreviated, and only the Base Transformation, base nonlinear point transformation, and Colour Transformation are reproduced at the encoder. That is, the error residuals collected in the residual image are only those generated by the nonlinear point transformation in the legacy coding path.

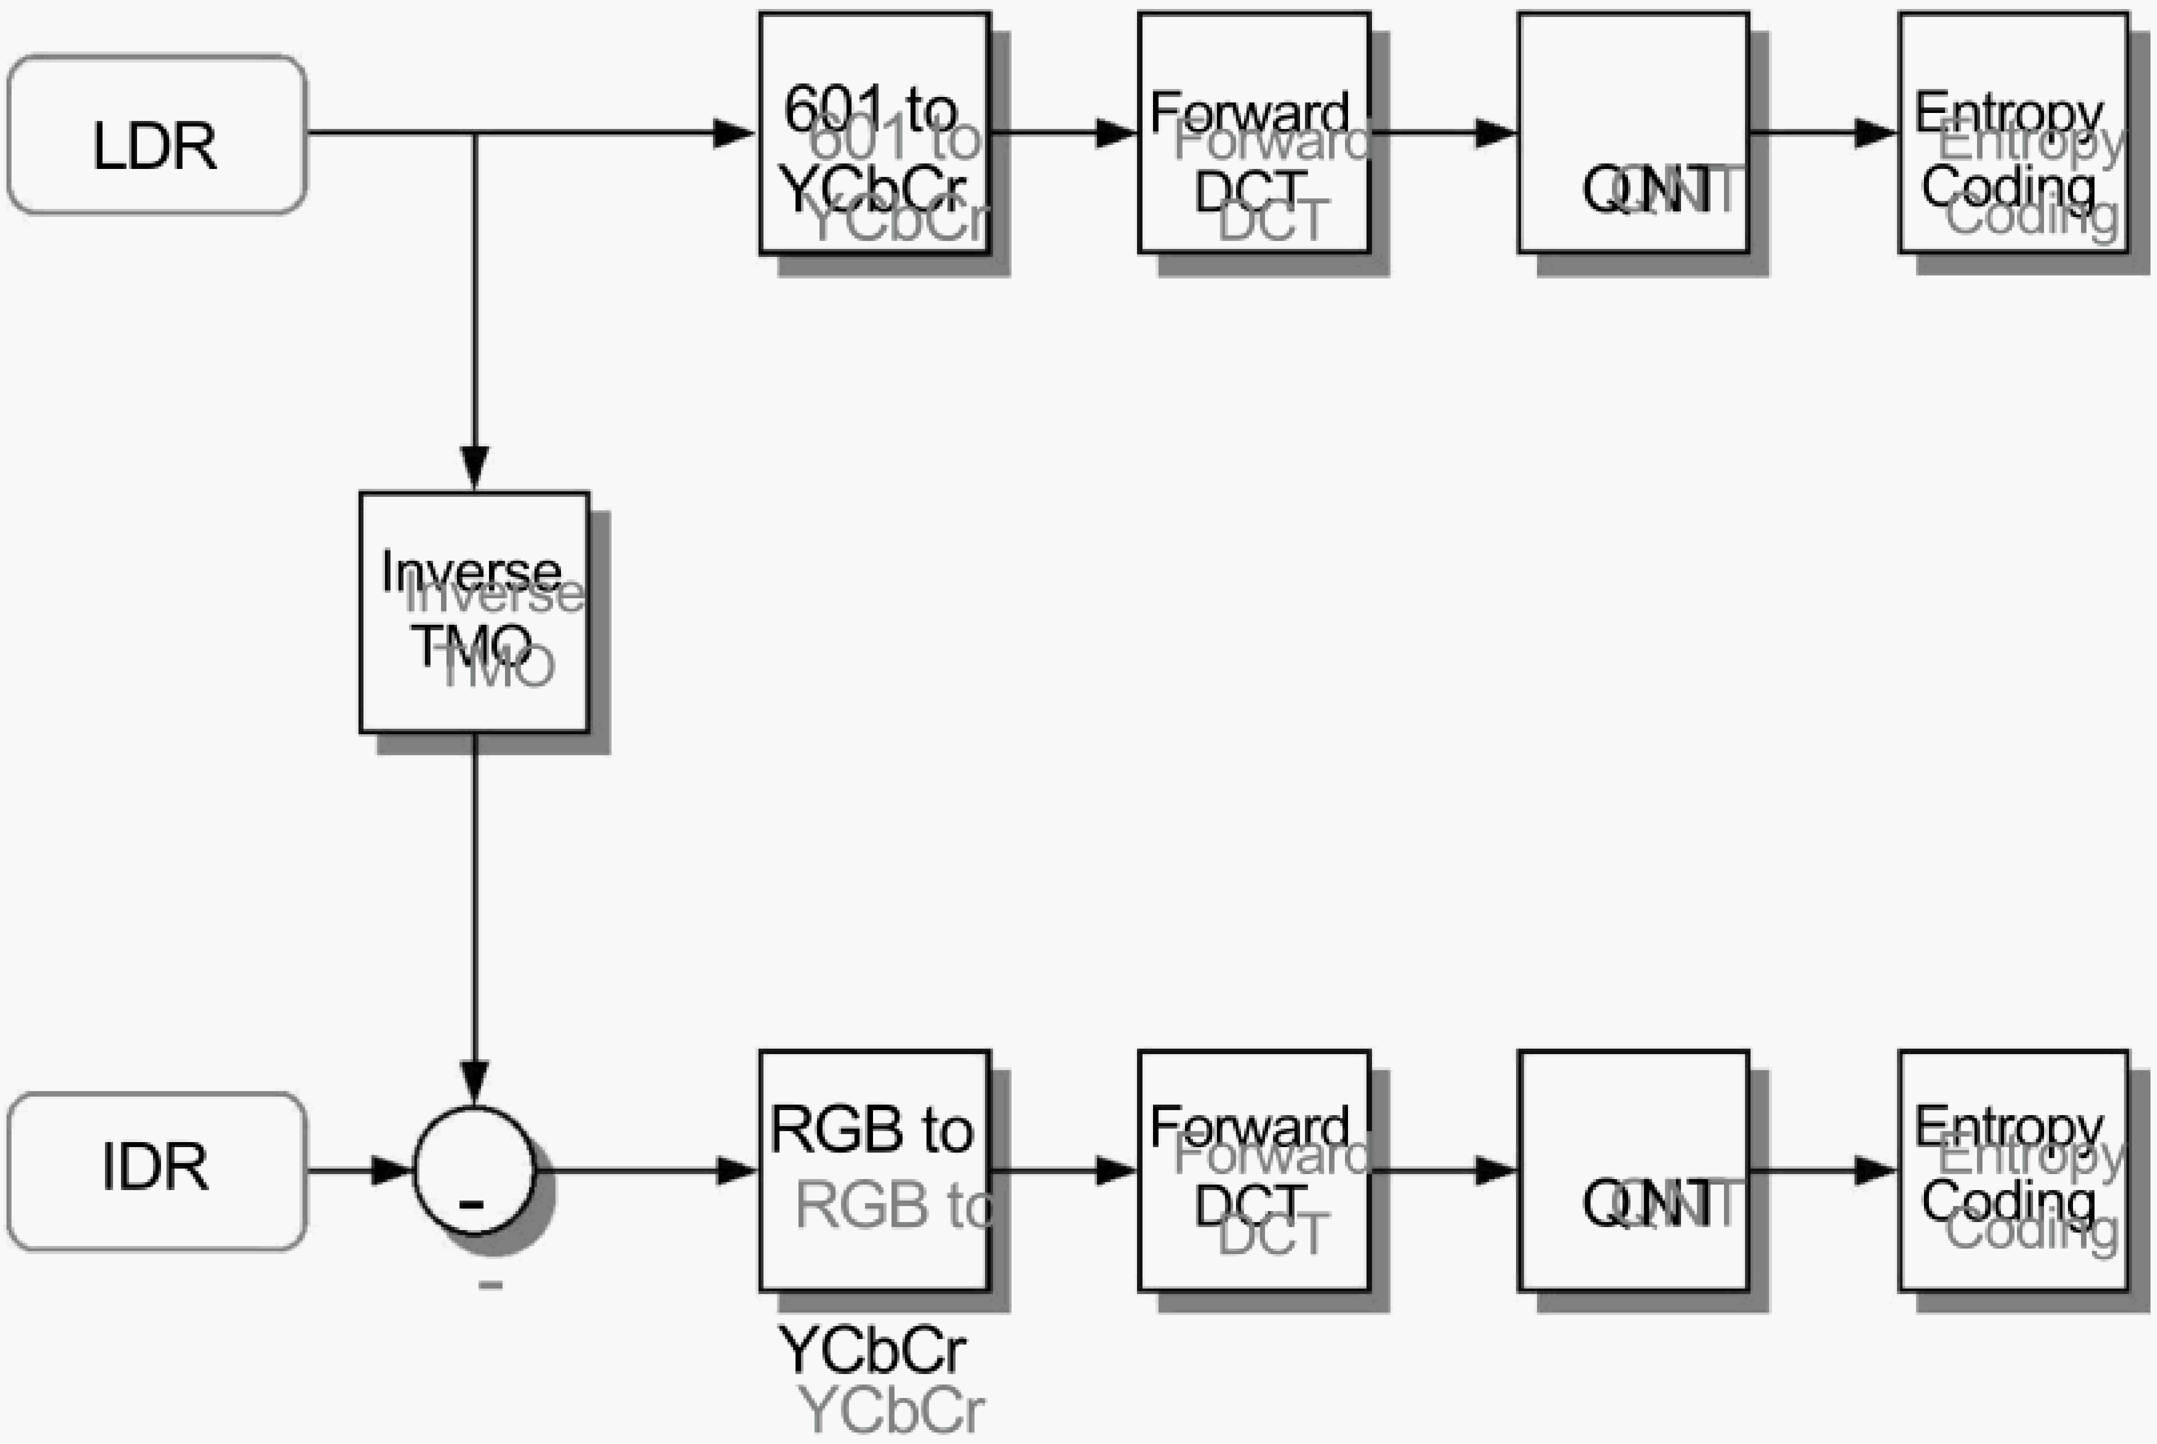

Figure A.3 — High level overview of the open loop encoder

Annex B
(normative)

Boxes

B.1 General

This Annex selects and refines a subset of the boxes defined in ISO/IEC 18477-3 for the purpose of intermediate range image coding. It lists those boxes of ISO/IEC 18477-3 that are required for this part of ISO/IEC 18477. All other boxes are optional and its interpretation is outside the scope of this part of ISO/IEC 18477. Other parts of ISO/IEC 18477 or other standards may define their meaning and decoders conforming to this part of ISO/IEC 18477 may ignore them.

Table B.1 lists the boxes required in this part of ISO/IEC 18477 that are paired with ISO/IEC 18477-3. Some of the boxes require additional specifications that are listed in subsequent clauses of this Annex.

Table B.1 — Boxes within this part of ISO/IEC 18477

Box Name	Box Type	Further definitions in which sub-clause of this part of ISO/IEC 18477
File Type box	0x66747970 (“ftyp”)	
Legacy Data Checksum box	0x4C43484B (“LCHK”)	
Residual Data box	0x52455349 (“RESI”)	
Residual Refinement box	0x5246494e (“RFIN”)	Annex D
Refinement Data box	0x46494e45 (“FINE”)	Annex D
Merging Specification box	0x53504543 (“SPEC”)	
Parametric Curve box	0x43555256 (“CURV”)	
Integer Table Lookup box	0x544f4e45 (“TONE”)	
Floating Point Table Lookup box	0x46544f4e (“FTON”)	
Fix-point Linear Transformation box	0x4D545258 (“MTRX”)	
Floating-point Linear Transformation box	0x46545258 (“FTRX”)	
Output Conversion box	0x4F434F4E (“OCON”)	B.2
Refinement Specification box	0x52535043 (“RSPC”)	
Base Nonlinear Point Transformation Specification box	0x4C505453 (“LPTS”)	B.3
Residual Nonlinear Point Transformation Specification box	0x5152505453 (“QPTS”)	B.4
Secondary Residual Nonlinear Point Transformation Specification box	0x5152505453 (“RPTS”)	B.5
Base Transformation box	0x4C545246 (“LTRF”)	B.6
Residual Transformation box	0x52545246 (“RTRF”)	B.7
Colour Transformation box	0x43545246 (“CTRF”)	B.8

B.2 Output conversion box

This mandatory box defines the final bitdepth of the reconstructed samples and defines the conversion process from the result of the base image/residual image merging process to the final output samples. It describes the final merging process and by that step B.10 of the algorithm described in A.1. This box is already defined in ISO/IEC 18477-3:2015, Annex B, though its application to this part of ISO/IEC 18477 further constraints the value of its fields.

This box shall never appear top level in the file, but it shall be a subbox of the Merging Specification box defined in ISO/IEC 18477-3:2015, Annex B. Exactly one Output Conversion box shall appear in the Merging Specification box if a Merging Specification box exists.

Table B.2 constraints the parameters of the Output Conversion box as applied in this part of ISO/IEC 18477.

Table B.2 — Parameter constraints of the Output Conversion box

Parameter	Constraints within this part of ISO/IEC 18477	Meaning
R_b	0..8	Number of additional bits available for high dynamic range data. The bit precision of the reconstructed high dynamic range image shall be computed as $8+R_b$.
Lf	0	This field indicates whether the compression is lossy or lossless. ISO/IEC 18477-6 only specifies lossy coding.
Oc	0	Shall be 0. This field indicates whether a conversion to floating point is desired, which is not within scope of this part of ISO/IEC 18477.
Ce	1	Shall be 1. This field indicates whether the output shall be clipped to range $[0,2^{8+R_b}-1]$. This is always the case for this part of ISO/IEC 18477.
Ol	0	Shall be 0. This field indicates whether an output lookup or point transformation is required. This is never the case for this part of ISO/IEC 18477.
to_0	0	Shall be 0. This field defines the output table for component 0 in other parts of the ISO/IEC 18477 family of standards.
to_1	0	Shall be 0. This field defines the output table for component 1 in other parts of the ISO/IEC 18477 family of standards.

Table B.2 *(continued)*

Parameter	Constraints within this part of ISO/IEC 18477	Meaning
to_2	0	Shall be 0. This field defines the output table for component 2 in other parts of the ISO/IEC 18477 family of standards
to_3	0	Shall be 0.

B.3 Base Nonlinear Point Transformation Specification box

This box defines the nonlinear point transformation between the output of the Base Transformation and the Colour Transformation. It, thus, defines step B4 in the decoder description in Annex A. At most, one Base Nonlinear Point Transformation Specification box shall exist as a sub-box of the Merging Specification box. It shall not appear at top-level of the file. The box layout of this box is that of the Nonlinear Transformation Specification box, defined in ISO/IEC 18477-3:2015, Annex B.

The nonlinear point transformation selected by this box shall either be an Integer Table Lookup box, a Floating Point Lookup box, or a Parametric Curve box. The corresponding boxes referenced by this box appear at top level of the ISO/IEC 18477-3 compliant file or as subbox of the Merging Specification box. The nonlinear point transformation itself is given by the process specified in ISO/IEC 18477-3:2015, Annex C. It requires four additional parameters, the input range R_w, R_e, and the output range R_t, R_f. The two value pairs shall be given as follows:

$R_w = 8 + R_h$ $\quad$ $R_e = 0$

$R_t = 8 + R_b$ $\quad$ $R_f = 0$

The value R_h is the number of refinement scans in the base decoding path and is found in the Refinement Specification box defined in ISO/IEC 18477-3:2015, Annex B. If the Refinement Specification box is absent, the inferred value of R_h is 0.

The value R_b is found in the Output Conversion box, where R_b+8 defines the total output precision of the image.

If this box does not exist, the implied nonlinear point transformation in the base decoding path is defined as if a Parametric Curve box with the identity function, i.e. with parameters t = 2 and e = 1, had been selected. That is, the input range is linearly scaled to the output range by the parameters R_w and R_t.

The type of this box shall be 0x4C52505453, ASCII encoding of “LPTS”. The box structure and layout does not deviate from that in ISO/IEC 18477-3:2015, Annex B, neither apply any restrictions to parameters of the box.

B.4 Residual Nonlinear Point Transformation Specification box

This box defines the nonlinear point transformation between the output of the residual DCT transformation and the Residual Transformation. It, thus, defines step B7 in the decoder description in Annex A. At most, one Residual nonlinear Point Transformation Specification box shall exist as a subbox of the Merging Specification box. It shall not appear at top-level of the file. The layout of this box is that of the Nonlinear Transformation Specification box, defined in ISO/IEC 18477-3:2015, Annex B. This box shall only be present if the Residual Data box is present.

The nonlinear point transformations selected by the td_i parameters this box, if it is present, shall be a Parametric Curve box with curve type t = 2 (identity) and rounding parameter e = 0. References to Integer or Floating Point Table Lookup boxes are not permitted. The corresponding boxes referenced by this box appear at top level of the ISO/IEC 18477-3 compliant file or as subboxes of the Merging

Specification box. The nonlinear point transformation itself is given by the process specified in ISO/IEC 18477-3:2015, Annex C. It requires four additional parameters, the input range R_w, R_e and the output range R_t, R_f. The two value pairs shall be given as follows:

$R_w = P + R_r \qquad R_e = 0$

$R_t = 8 + R_b \qquad R_f = 0$

The value R_r is the number of residual refinement scans in the base decoding path and is found in the Refinement Specification box defined in Table B.1. If the Refinement Specification box is absent, the inferred value of R_r is 0. The value of P is the bitdepth of the residual image; it is found in the frame header of the residual codestream (see Rec. ITU-T T.81 | ISO/IEC 10918-1:1994, Table B.2).

The value R_b is found in the Output Conversion box, where R_b+8 defines the total output precision of the image.

If this box does not exist, the implied nonlinear point transformation in the base decoding path is defined as if a Parametric Curve box with the identity function had been selected, i.e. with parameters t = 2 and e = 0. That is, the input range is linearly scaled to the output range by the parameters R_w and R_t.

The type of this box shall be 0x5152505453, ASCII encoding of “QPTS”. The box structure and layout does not deviate from that in ISO/IEC 18477-3:2015, Annex B.

NOTE The default rounding mode in the residual domain in the absence of this box is e = 0, unlike the default rounding mode in the base decoding path, which is e = 1. This deviation is intentional.

B.5 Secondary Residual Nonlinear Point Transformation Specification box

This box, if present, selects a secondary nonlinear point-transformation that is applied in the residual domain **after** the residual transformation. This box implements processing step B8d in the description of Annex A, i.e. it maps the inversely decorrelated residual error values Ω_i into the final residual error values Q_i that are added to the precursor image. If this box is not present, Q_i is identical to Ω_i. The layout of this box is that of the Nonlinear Transformation Specification box, defined in ISO/IEC 18477-3:2015, Annex B.

The nonlinear point transformation selected by this box shall only be a Parametric Curve box, with the Curve Type t = 5 (linear ramp) or t = 2 (identity) (see ISO/IEC 18477-3:2015, Annex B). The corresponding boxes referenced by this box appear at top level of the ISO/IEC 18477-3 compliant file or as subboxes of the Merging Specification box. The nonlinear point transformation itself is given by the process specified in ISO/IEC 18477-3:2015, Annex C. It requires four additional parameters, the input range R_w, R_e, and the output range R_t, R_f. The two values shall be given as follows:

$R_w = 8 + R_b \qquad R_e = 0$

$R_t = 8 + R_b \qquad R_f = 0$

The value R_b is found in the Output Conversion box, where R_b+8 defines the total output precision of the image.

If this box does not exist, the implied nonlinear point transformation in the residual decoding path is defined as if a Parametric Curve box with the identity function had been selected, i.e with parameters t = 2 and e = 0. That is, the input is identical to the output.

The type of this box shall be 0x52505453, ASCII encoding of “RPTS”. The box structure and layout does not deviate from that in ISO/IEC 18477-3:2015, Annex B.

B.6 Base Transformation box

This box defines the linear transformation between the output of the legacy entropy decoding process and the input of the base image nonlinear point transformation. It defines the transformation in step

B3 of the decoding process specified in Annex A. The box structure and layout is already defined in ISO/IEC 18477-3:2015, Annex B, though its purpose is refined here.

There shall be at most one Base Transformation box as subbox of the Merging Specification box. This box shall only exist if the number of components Nf is equal to 3, and it shall exist if the output bitdepth is larger than 8, i.e. the R_b parameter of the Output Conversion box is nonzero.

If the box does not exist, the inverse decorrelation transformation is the identity process if Nf is equal to 1. If the box does not exist and the number of components is equal to 3, the base transformation in step B3 of Annex A is defined by the Component Decorrelation Control Marker specified in ISO/IEC 18477-1:2015, Annex B, that is, it is the ICT if the Decorrelation Control Marker is absent or its cc value is equal to 1, it is the identity if the Decorrelation Control Marker exists and its cc value is 0. Details on the selection of the base transformation are specified in Table C.1. The linear transformations specified in Annex C require an additional level shift parameter R_s. The value of R_s for the base transformation shall be given as

$$R_s = 8 + R_h - 1$$

where R_h is the number of refinement scans included in the base decoding path. The value of R_h is found in the Refinement Specification subbox of the Merging Parameter box. It shall be 0 if no Refinement Specification box is present.

The type of the Base Transformation box shall be 0x4C545246, ASCII encoding of “LTRF”.

Table B.3 constraints the parameters of the Base Transformation box based on the parameters listed in ISO/IEC 18477-3:2015, Annex B.

Table B.3 — Parameter constraints for the Base Transformation box

Parameter	Constraints within this part of ISO/IEC 18477	Meaning
Xt	1, 2 or 5..15	Defines the linear transformation to be used as base transformation (see Table B.5 for the encoding of this field).
Re	0	Shall be 0.

B.7 Residual Transformation box

This box defines the linear transformation between the output of the nonlinear point transformation in the residual decoding path and the addition of the residual to the output of the colour transformation in the base decoding path. It defines the linear transformation in step B8 of the decoding process specified in Annex A. The box structure and layout is already defined in ISO/IEC 18477-3:2015, Annex B, though its purpose is refined here.

This box shall only exist as a subbox of the Merging Specification box specified in ISO/IEC 18477-3:2015, Annex B. It shall not appear top level. This box shall exist if and only if a Residual Data box is present at the top level of the file and the number of components Nf equals 3.

The linear transformations specified in Annex C require an additional level shift parameter R_s. The value of R_s for the residual transformation shall be given as

$$R_s = 8 + R_b - 1$$

where R_b+8 is the sample precision of the reconstructed IDR output image. The value of R_b can be found in the Output Conversion box (see B.2).

The type of this box shall be 0x52545246, ASCII encoding of “RTRF”.

Table B.4 constraints the parameters of the Residual Transformation box from the generic Linear Transformation Specification box defined in ISO/IEC 18477-3.

Table B.4 — Parameter constraints of the Residual Transformation box

Parameter	Constraints within this part of ISO/IEC 18477	Meaning
Xt	1, 2 or 5..15	Defines the linear transformation to be used as residual transformation (see Table B.5 for the encoding).
Re	0	Shall be 0.

Table B.5 — Encoding of the Xt parameter of the Base and Residual Transformation box

Value	Transformation to be used
0	Reserved for ITU | ISO/IEC purposes.
1	The identity transformation shall be used.
2	The ICT Transformation as specified in Annex C shall be used.
3	Reserved for ITU | ISO/IEC purposes.
4	Reserved for ITU | ISO/IEC purposes.
5..15	The free form transformation with offset shift defined by the Integer or Floating Point Linear Transformation box whose M value matches the value of Xt shall be used. The Integer or Floating Point Linear Transformation boxes are specified in ISO/IEC 18477-3:2015, B.2 and B.3 and their application and implementation are specified in Annex C.

B.8 Colour Transformation box

This box defines the linear transformation between the output of the nonlinear point transformation in the base decoding path and the addition of the inversely decorrelated transformed residual. It defines the linear transformation in step B4a of the decoding process specified in Annex A. The box structure and layout is already defined in ISO/IEC 18477-3:2015, Annex B, though its purpose is refined here.

This box shall only exist as a subbox of the Merging Specification box specified in ISO/IEC 18477-3:2015, Annex B and it may only exist if the number of components in the image Nf equals 3. It shall not appear top level. If this box does not exist, the colour transformation shall be the identity transformation; otherwise the Xt parameter of the Colour Transformation box specifies the transformation matrix to pick.

The linear transformations specified in Annex C require an additional level shift parameter R_S. The value of R_S for the colour transformation shall be given as

$R_S = -\infty$ (i.e. no level shift)

The type of this box shall be 0x43545246, ASCII encoding of “CTRF”.

Table B.6 constraints the parameters from that defined in ISO/IEC 18477-3.

Table B.6 — Parameter Constraints for the Colour Transformation box

Parameter	Constraints within this part of ISO/IEC 18477	Meaning
Xt	1 or 5..15	Defines the linear transformation to be used as base transformation (see Table B.7 for the encoding).
Re	0	Shall be 0.

Table B.7 — Encoding of the Xt parameter of the Colour Transformation box

Value	Transformation to be used
0	Reserved for ITU | ISO/IEC purposes.
1	The identity transformation shall be used.
2..4	Reserved for ITU | ISO/IEC purposes.
5..15	The free form transformation without offset shift defined by the Integer or Floating Point Linear Transformation box whose M value matches the value of Xt shall be used. The Integer or Floating Point Linear Transformation boxes are specified in ISO/IEC 18477-3:2015, B.2 and B.3 and their application and implementation are specified in Annex C.

Annex C
(normative)

Multi-component decorrelation

In this Annex and all of its subclauses, the flow charts and tables are normative only in the sense that they are defining an output that alternative implementations shall duplicate.

C.1 General

This Annex defines the multiple component decorrelation transformations available as base, residual, and colour transformations of the decoding process. They are selected by the Base Transformation box, the Residual Transformation box, or the Colour Transformation box. The base transformation replaces the multiple component decorrelation transformation of ISO/IEC 18477-1, the residual and colour transformation are new to this part of ISO/IEC 18477.

A multiple component decorrelation transformation takes one or three input components and generates from them one or three output components. A task dependent level shift by a value of 2^{R_S} ensures that chroma components are correctly centred and representable by unsigned values. The values of R_S that are applicable to the transformations are specified in B.6, B.7, and B.8.

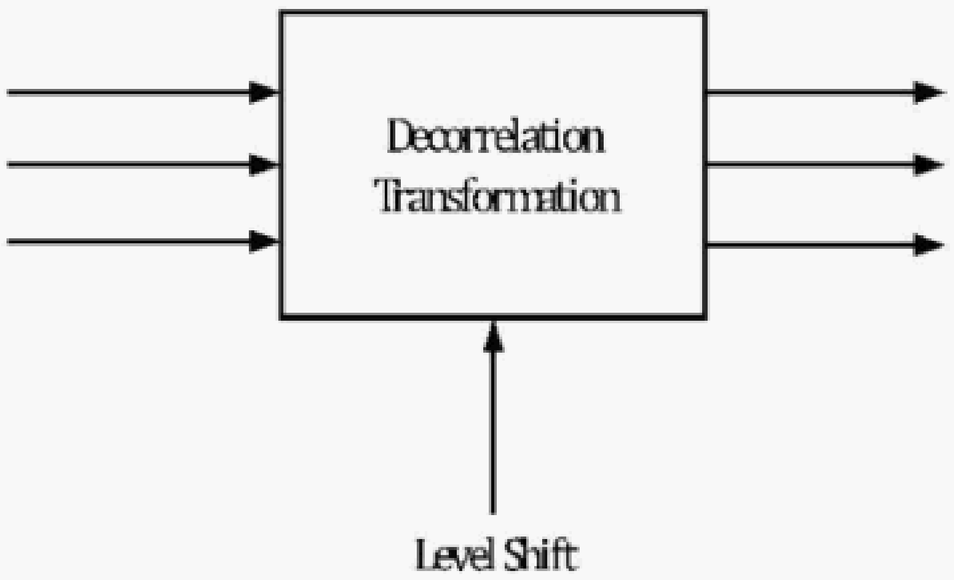

Figure C.1 — Input and output of a decorrelation transformation: Input components, output components, and the DC level shift

Table C.1 specifies the decorrelation transformation applicable to the samples reconstructed in the base decoding path.

Table C.1 — Selection of the Base Transformation

Nf Number of components	Xt field of the Base Transformation box	Component Decorrelation Marker present and cc = 0	Base Transformation as represented in block B3 in Annex A.
1	Base Transformation box shall not be present	ignored	Inverse Identity Transformation
3	Base Transformation box non-existing	yes	Inverse Identity Transformation
3	Base Transformation box non-existing	no	Inverse ICT
3	1	ignored	Inverse Identity Transformation
3	2	ignored	Inverse ICT

Table C.1 *(continued)*

Nf	Xt field of the Base Transformation box	Component Decorrelation Marker present and cc = 0	Base Transformation as represented in block B3 in Annex A.
Number of components			
3	5..15	ignored	Free form transformation from an Integer or Floating Point Linear Transformation box whose M value is identical to the value of LXt.
All other values			Reserved for ITU | ISO/IEC purposes

Table C.2 selects the residual transformation that decorrelates the residual data

Table C.2 — Selection of the Residual Transformation

Nf	Xt	R-Transformation as represented by block B8 specified in Annex A.
Number of components		
1	The Residual Transformation box shall not be present.	Inverse Identity Transformation
3	1	Inverse Identity Transformation
3	2	Inverse ICT
All other values		Reserved for ITU | ISO/IEC purposes

Table C.3 selects the colour transformation that transforms the precursor image into the target colourspace.

Table C.3 — Selection of the Colour Transformation

Nf	Xt	C-Transformation, as represented by block B4 (see Annex A).
Number of components		
1	The Colour Transformation box shall not be present.	Scaled Inverse Identity Transformation
3	1	Scaled Inverse Identity Transformation
3	5..15	Free form transformation defined by the Linear Transformation box whose M value matches Xt.
All other values		Reserved for ITU | ISO/IEC purposes

C.2 Irreversible Inverse Multi-component Transformation (ICT) (normative)

This transformation is it is identical to the ICT in ISO/IEC 18477-1 for 8-bit LDR data and can be understood as an inverse decorrelation from YCbCr to Rec. ITU-T BT.601 RGB.

In the following, let I_0, I_1, and I_2 be the input sample values reconstructed from the upsampled components 0, 1, and 2, and let L_0, L_1, and L_2 be the output sample value. Let R_S be the number of level shift scale bits defined in B.6, B.7, and B.8. Then:

$$L_0 = I_0 + 1{,}402\,00 \times (I_2 - 2^{R_S})$$

$$L_1 = I_0 - 0{,}714\,136\,285\,9 \times (I_2 - 2^{R_S}) - 0{,}344\,136\,286\,1 \times (I_1 - 2^{R_S})$$

$$L_2 = I_0 + 1{,}772 \times (I_1 - 2^{R_S})$$

C.3 Free-form Integer Multi-component Transformation (normative)

This transformation offers a generic linear transformation that can be applied as base, residual, or colour transformation. It is applied whenever a Transformation box from B.6, B.7, or B.8 selects an Integer Transformation box by matching its M value. Let I_0, I_1, and I_2 be the inputs to this transformation and O_0, O_1, and O_2 its output. Denote by a_{ij} the parameters in the matching Integer Transformation box. Then:

$$O_0 = [a_{11} \times I_0 + a_{12} \times (I_1 - 2^{R_s}) + a_{13} \times (I_2 - 2^{R_s})]/2^{13}$$

$$O_1 = [a_{21} \times I_0 + a_{22} \times (I_1 - 2^{R_s}) + a_{23} \times (I_2 - 2^{R_s})]/2^{13}$$

$$O_2 = [a_{31} \times I_0 + a_{32} \times (I_1 - 2^{R_s}) + a_{33} \times (I_2 - 2^{R_s})]/2^{13}$$

C.4 Free-form Floating Point Multi-component Transformation (normative)

This transformation offers a generic linear transformation that can be applied as base, residual, or colour transformation. It is applied whenever a Transformation box from B.6, B.7, or B.8 selects a Floating Point Transformation box by matching its M value. Let I_0, I_1, and I_2 be the inputs to this transformation and O_0, O_1, and O_2 its output. Denote by a_{ij} the parameters in the matching Integer Transformation box. Then:

$$O_0 = a_{11} \times I_0 + a_{12} \times (I_1 - 2^{R_s}) + a_{13} \times (I_2 - 2^{R_s})$$

$$O_1 = a_{21} \times I_0 + a_{22} \times (I_1 - 2^{R_s}) + a_{23} \times (I_2 - 2^{R_s})$$

$$O_2 = a_{31} \times I_0 + a_{32} \times (I_1 - 2^{R_s}) + a_{33} \times (I_2 - 2^{R_s})$$

C.5 Inverse Identity Transformation (normative)

The inverse identity transformation maps its input values to the output without applying an inverse linear decorrelation step. Denote the input value of the identity transformation with I_0 and the output value by L_0. The inverse identity transformation then sets

$$L_0 = I_0$$

If more than one component has to be transformed, the above equation(s) apply identically to all components.

C.6 Irreversible Forward Multi-Component Transformation (ICT) (informative)

This transformation is used to decorrelate the low dynamic range versions of the image before encoding the components by a Rec. ITU-T T.81 | ISO/IEC 10918-1 conforming encoder. It is identical to the forward ICT defined in ISO/IEC 18477-1. Let L_0 to L_2 be the input pixel values and R_s the number of level shift scale bits, then the output I_0 to I_2 shall be computed as follows:

$$I_0 = 2{,}990\ 0 \times L_0 + 0{,}58700 \times L_1 + 0{,}114\ 00 \times L_2$$

$$I_1 = -0{,}168\ 735\ 891\ 6 \times L_1 - 0{,}331\ 264\ 108\ 4 \times L_1 + 0.5 \times L_2 + 2^{R_s}$$

$$I_2 = 0{,}5\ L_2 - 0{,}418\ 687\ 589\ 2 \times L_1 - 0{,}081\ 312\ 410\ 85 \times L_2 + 2^{R_s}$$

C.7 Forward Identity Transformation (normative)

This transformation is used to generate the input of the legacy coding process when the number of components in the frame is 1, or the component decorrelation transformation is disabled. If L_0 is the input value, then this transformation sets the output value I_0 to

$$I_0 = L_0$$

If more than one component is to be transformed, then the above applies to all three components identically.

Annex D
(normative)

Entropy coding of refinement data

In this Annex and all of its subclauses, the flow charts and tables are normative only in the sense that they are defining an output that alternative implementations shall duplicate.

D.1 General (informative)

This Annex defines the procedures for decoding and encoding of images using an extended bitdepths in the DCT domain. If this coding method is used, the entropy coded data visible to legacy decoders describes an image with a precision of eight bits per sample, which is extended by a refinement coding procedure described in this Annex to up to 12 bits. Refinement coding provides a simple and efficient technique to extend the bit depths of the coded samples in a backwards compatible way even without requiring a residual codestream.

Refinement coding encodes additional least significant bits of the base or residual image by a mechanism that is closely related to the subsequent approximation coding method of the progressive scan mode of Rec. ITU-T T.81 | ISO/IEC 10918-1. One additional least significant bit is added per additional refinement scan.

Refinement coding is signalled by a Refinement Specification box in the Merging Specification box. This box includes two parameters, R_h and R_r where the first encodes the number of additional bits by which the base image is extended by refinement coding, and the latter specifies the number of bits by which the residual image is extended. It may take values from zero, refinement scan disabled, to four.

The entropy coded data of each refinement scan is encapsulated into boxes, a single box per scan and thus per added bit. The box type is selected according to the image type the refinement scan applies to: Refinement Data boxes contain entropy coded data of refinement scans extending the base image, Residual Refinement boxes entropy coded data extending the bit-precision of the residual image. Table D.1 provides an overview which boxes carry which data.

Table D.1 — Refinement Scan Types and Encapsulating box types

Refinement of which data	Box Type Used	Number Refinement Scans specified in	Scan Type
Base image	Refinement Data box	R_h	Refinement (modified successive approximation scan)
Residual image	Residual Refinement Data box	R_r	Refinement (modified successive approximation scan)

If the number of refinement scans is larger than one, multiple (Residual) Refinement Data boxes shall be written, one per scan over the data. The Box Instance Number **En**, see ISO/IEC 18477-3:2015, Figure A.1 disambiguates then between the scans, i.e. the first scan has a value of En = 1, the next one En = 2, and so on. See ISO/IEC 18477-3:2015, Annex A for details on the Box syntax. The number of (Residual) Refinement boxes shall be given by the R_h or R_r parameter of the Refinement Specification box.

An encoder creating four refinements over the base image data would place the first refinement scan into a Refinement Data box with En = 1, the second into a box with En = 2, and so on. However, because entropy coded data of individual refinement scans may still require more than 64K for storage, each of the Refinement Data boxes of the individual scans may be broken up into several APP_{11} marker segments. The first of the segments contributing to **the same refinement scan** will have Z = 1, the second Z = 2, the third Z = 3, and so on. See ISO/IEC 18477-3:2015, Annex A for more details on the box syntax.

If the refinement scan is enabled by setting $R_h \neq 0$ or $R_r \neq 0$ in the Refinement Specification box, the quantization and inverse quantization process defined in Rec. ITU-T T.81 | ISO/IEC 10918-1 have to be modified to include the additional bits, and the entropy coding and decoding procedures includes additional scans over the image to encode and decode such bits. The entropy coded data of each scan is located in a separate Refinement Data box or Residual Refinement Data box, as specified in ISO/IEC 18477-3:2015, B.12 and B.13.

The DCT transformation and multi-component decorrelation transformation remain unchanged except that the DC level shift has to be modified to include the additional bits. That is, while Rec. ITU-T T.81 | ISO/IEC 10918-1:1994, F.1.1.3 and F.2.1.5 specify a level shift of $S_{dc} = 2^7$ for a sample precision of eight bit, the level shift is replaced by a shift of $S_{dc} = 2^{7+h}$ if h refinement scans are included in the decoding process, where h is either R_r or R_h depending on whether the legacy or the residual codestream is refined.

A decoder that chooses to ignore refinement scans will continue to operate with the level shift as indicated in the legacy standard and will be able to reconstruct a visible image, but will not take advantage of the enhanced bit precision.

Additional care must be taken to avoid overflow conditions in the DCT as the range of the input and output data is extended.

With h refinement scans present, the scans over the base or residual image contained in the legacy codestream or the Residual Data box operate as if the point shift values Al and Ah of the scan headers were replaced by Al+h and Ah+h. That is, the legacy coding passes decode by a quantization step size that is by a factor of 2^h larger than the indicated quantization step size.

Refinement coding then decodes into the h least significant bits of the quantized DCT coefficients as if the scans using the legacy scan modes were the first passes of a progressive JPEG coding mode. The refinement scans then operate as subsequent approximation scans extending the legacy scans, by the very same mechanism defined in Rec. ITU-T T.81 | ISO/IEC 10918-1:1994, Annex G.

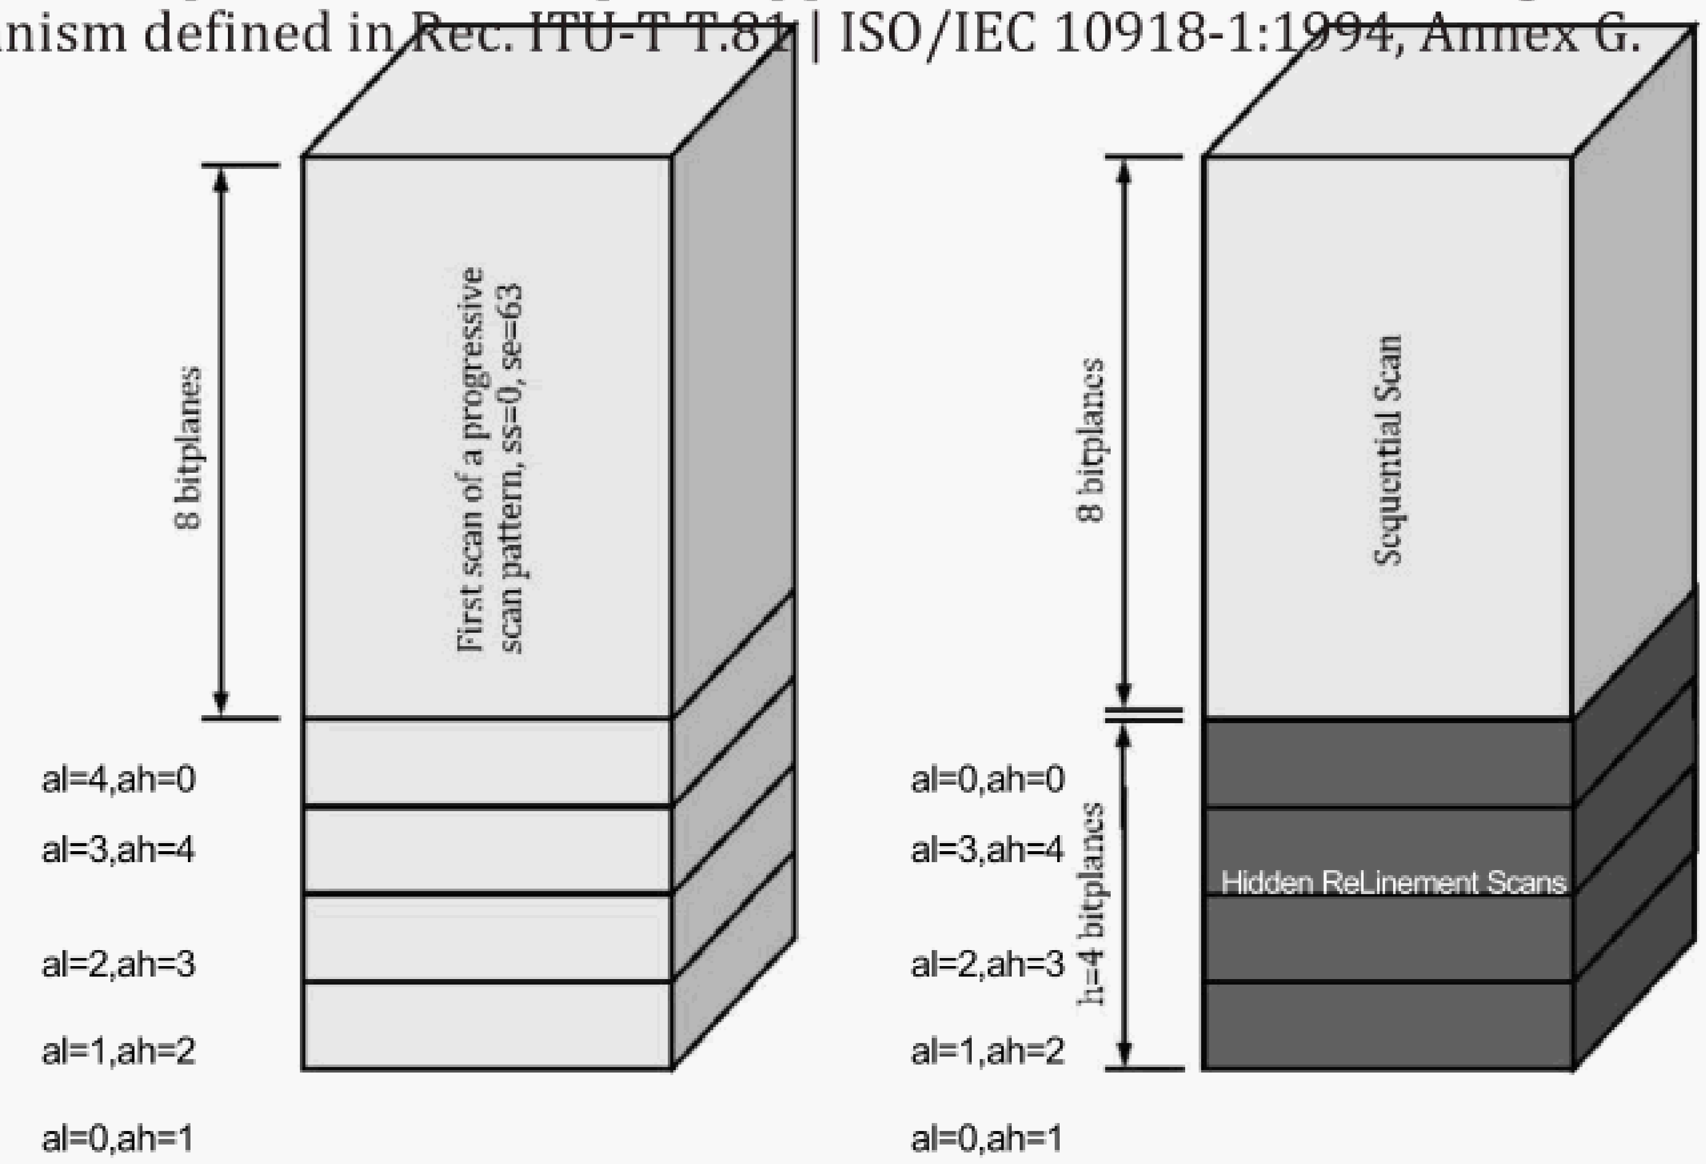

Figure D.1 — Progressive Coding with 12 bitplanes and Refinement Coding with $R_h = 4$ refinement scans

Progressive Coding with 12 bitplanes and refinement coding with four refinement scans do not differ significantly. The progressive scan signals additional LSBs by setting al = 0 in the first scan, whereas sequential coding does not signal the presence of refinement scans in the codestream and leaves al = 0. Note that the actual scan pattern and entropy coding procedure for the first eight bitplanes of the progressive coding example is exactly identical to the sequential scan with eight bitplanes, and that refinement coding uses the progressive successive approximation coding. The entropy coded data will thus be identical, for both the first eight bitplanes and the remaining four LSBs. The difference is that the four LSBs remain invisible to legacy decoders if refinement coding is used instead of a 12-bit progressive mode.

Refinement coding is thus nothing but the successive approximation coding of the legacy standard, except that the very first coding pass (or passes) of legacy data do not indicate the necessary upshift of bits to make room for additional least significant bits. Bits included in the legacy scan pattern are thus reconstructed to values by a factor of 2^h larger in the presence of h refinement scans.

D.2 Modifications of the inverse quantization process for refinement decoding (normative)

The regular dequantization process in the absence of refinement scans multiplies the entropy decoded DCT coefficient value $S_{i,j}$ at block position i,j with the quantization matrix entry $q_{i,j}$ to get the DCT coefficient $C_{i,j}$. In the presence of a residual scan, h additional bits $T_{i,j}$ are decoded by the refinement scan that shall be included as h least significant bits in the reconstruction of the DCT coefficient $C_{i,j}$ as follows:

$$C_{i,j} = (2^h \times S_{i,j} + T_{i,j}) \times q_{i,j}$$

NOTE 1 While the decoding process will ensure that $T_{0,0}$ will always be positive or zero, all other $T_{i,j}$ may be positive, negative, or zero.

NOTE 2 An equivalent implementation strategy for the above inverse quantization procedure is to decode the data from the legacy scans represented by the entropy coded data following the SOS markers as if the point transformation parameter of the legacy scans would be set to Al+h (or 0+h) instead of Al (or zero). Al is the point transformation parameter of the scan header, see of Rec. ITU-T T.81 | ISO/IEC 10918-1:1994, A.4 and B.2.3. Data decoded by the refinement scan uses an unmodified decoding procedure and is then automatically placed in the least significant bits of $S_{i,j}$.

D.3 Modifications of the quantization process for refinement coding (informative)

In the presence of refinement scans, the quantization procedure generates from the DCT coefficients $C_{i,j}$ two data sets, the legacy quantized data $S_{i,j}$ which will be encoded by the entropy coding methods provided by Rec. ITU-T T.81 | ISO/IEC 10918-1, and the refinement data $T_{i,j}$ which undergoes refinement coding. $S_{i,j}$ and $T_{i,j}$ are computed from the DCT coefficients $C_{i,j}$ and the quantization matrix $q_{i,j}$ as follows:

$$X_{i,j} = \mathrm{sign}(C_{i,j}) \times \left\lfloor \left|C_{i,j}\right| / q_{i,j} + 0{,}5 \right\rfloor$$

$$S_{0,0} = \left\lfloor X_{0,0} / 2^h \right\rfloor \qquad T_{0,0} = X_{0,0} - 2^h \times S_{0,0}$$

$$S_{i,j} = \mathrm{sign}(X_{i,j}) \times \left\lfloor \left|X_{i,j}\right| / 2^h \right\rfloor \qquad T_{i,j} = X_{i,j} - 2^h \times S_{i,j} \text{ for } (i,j) \neq (0,0)$$

NOTE The division and rounding algorithm for computing $S_{i,j}$ generates a "fat zero" for AC coefficients, i.e. a dead zone of twice the size of a regular quantization bucket. DC coding does not follow this convention to avoid drift errors when encoding. The rounding conventions picked here are intentionally identical to the rounding procedure used when encoding data with the successive approximation of the progressive scan defined in Rec. ITU-T T.81 | ISO/IEC 10918-1, and this is, in fact, one possible implementation strategy for refinement coding: The legacy codestream is created by an encoder as if the successive approximation parameter would be Al+h (or h) instead of Al (or zero), the remaining bits are then encoded by the refinement scan defined in this Annex.

D.4 Decoding process for entropy coded refinement data (normative)

This subclause defines the decoding procedure for the refinement data $T_{i,j}$ that extends the precision of the DCT coefficients by h bits. Refinement data is encapsulated in Refinement Data boxes or Residual Refinement boxes, which are embedded in the legacy codestream by means of the mechanisms defined in ISO/IEC 18477-3:2015, Annex B. Each such box contains exactly one scan over the data. A JPEG XT image using h refinement scans will thus contain h Refinement boxes.

The contents of the Refinement Data box or Residual Refinement box consists of marker segments that define the Huffman tables for refinement data coding, followed by a scan header and entropy coded data. The syntax of the contents of the Refinement Data box follows Rec. ITU-T T.81 | ISO/IEC 10918-1, except that neither a frame header nor DQT marker segments shall be present. Restart markers, if indicated in the legacy codestream, shall however be included, following the syntax defined in the legacy standard, and a DHT marker segment defining the Huffman code table required to decode the data shall be present as well. The scan included in this stream is decoded by a modification of the successive approximation decoding algorithm using Huffman decoding defined in Rec. ITU-T T.81 | ISO/IEC 10918-1:1994, Annex G, **regardless of the frame type indicated in the legacy or residual codestream** (i.e. regardless of whether the frame type is baseline, extended sequential, or progressive).

[Tables/ misc]	Scan header	ECS $_0$	[RST $_0$]	...	ECS $_{last}$	[RST $_{last}$]

Figure D.2 — Syntax of the refinement codestream

The entropy coding algorithm used for refinement decoding is identical to the successive approximation decoding defined in Rec. ITU-T T.81 | ISO/IEC 10918-1:1994, Annex G, except that the scan coefficients ZZ'(K) have to be initialized before starting the first refinement scan by the entropy decoded data of the legacy or residual codestream upshifted by h bits, i.e.

$$ZZ'(K) = 2^h \times ZZ(K)$$

where ZZ(k) is the value of the legacy/residual decoded data at the k-th position of the zig-zag scan defined in Rec. ITU-T T.81 | ISO/IEC 10918-1:1994, Figure 5. The refinement data $T_{i,j}$ after decoding is then defined by the output ZZ'(K) of all successive approximation scans included in the Refinement Data or Residual Refinement boxes:

$$T_{(i,j)}(k) = ZZ'(K) - 2^h \times ZZ(K)$$

NOTE An alternative implementation strategy would first decode the legacy portion of the codestream using either the baseline, sequential, or progressive decoding procedure, but with a point transformation/successive approximation parameter of Al+h (or h) for each scan included in the legacy codestream. The refinement scan extends these scans on the same scan buffer by using a successive approximation scan on the same scan buffer while using the successive approximation parameter exactly as indicated in the scan headers included in the Refinement boxes. Note further that the refinement scan always uses a successive approximation scan pattern, regardless of the frame type of the legacy codestream.

D.5 Encoding of refinement data (informative)

This subclause specifies an encoding process of refinement data that is compatible to the decoding process defined in D.4. In the first step, the scan pattern ZZ'(K) is initialized with the quantized data encoded in the legacy codestream, $S_{i,j}$, plus the refinement data $T_{i,j}$:

$$ZZ'(K) = 2^h \times ZZ(K) + T_{(i,j)}(k)$$

where ZZ(K) is the legacy data at the k-th position in the zig-zag scan represented by the legacy codestream. Then, one or several successive approximation scans following the specifications in Rec. ITU-T T.81 | ISO/IEC 10918-1:1994, Annex G encode the h least significant bits of the data, i.e. the Al parameter of the scans generated by the refinement scans will vary between h-1 and zero.

The resulting entropy coded data of each scan, together with the scan headers defining the scan pattern and the DHT marker segments defining the Huffman tables, but without any frame header or DQT markers, form the contents of the Refinement Data or Residual Refinement boxes. The output of each scan, together with the scan header and the DHT marker segments are encapsulated in such a box, and the box is split into one or multiple JPEG XT Extension marker segments by the mechanism described in ISO/IEC 18477-3:2015, Annex B.

NOTE An alternative implementation would first use a legacy scan process with a point transformation of Al+h (or h) bits to encode the LDR data, and would then run a successive approximation scan on the same data buffer with an unmodified parameter of Al. Note further that the refinement encoding procedure always uses the successive approximation scan of the progressive coding mode, regardless of the frame type of the legacy codestream.

D.6 Modifications of the inverse DCT in the presence of refinement scans (normative)

The Inverse DCT Transformation (IDCT) takes integer samples $Q_{i,j}$ in the form of an 8 × 8 block and generates from that integer samples $Y_{i,j}$ also organized in an 8 × 8 block. The input samples $Q_{i,j}$ are integers generated by the dequantization procedure described in Rec. ITU-T T.81 | ISO/IEC 10918-1. The following modifications apply to the inverse DCT transformation specified in Rec. ITU-T T.81 | ISO/IEC 10918-1:1994, A.3:

— The DC Level Shift specified in Rec. ITU-T T.81 | ISO/IEC 10918-1:1994, A.3.1 shall be 2^{P+h-1} instead of 2^{P-1}, i.e. the level shift is scaled by the number of refinement scans.

Annex E
(normative)

Profiles

This part of ISO/IEC 18477 does not currently define any profiles. A file compliant to this part of ISO/IEC 18477 shall be indicated by a an entry CL_i in the compatibility list of the File Type box with the value 0x69726670, ASCII encoding for “irfp”. The File Type box is defined in ISO/IEC 18477-3:2015, Annex B.

Bibliography

[1] ISO/IEC 10646-1, *Information Technology — Universal multiple-octet coded character set (UCS) — Part 1: Architecture and basic multilingual plane*

[2] ISO/IEC/IEEE 60559, *Information technology — Microprocessor systems — Floating-point arithmetic*

[3] IEC 61966-2-1, *Multimedia systems and equipment — Colour measurement and management — Part 2-1: Colour management — default RGD colour space —sRGB*

[4] Rec. ITU-T T.86 | ISO/IEC 10918-4, *Information technology — Digital compression and coding of continuous-tone still images — Registration of JPEG profiles, SPIFF profiles, SPIFF tags, SPIFF colour spaces, APPn markers, SPIFF compression types, and Registration Authorities*

[5] Rec. ITU-T T.871 | ISO/IEC 10918-5, *Information technology — Digital compression and coding of continuous-tone still images — JPEG File Interchange Format*

ICS 35.040
Price based on 34 pages

This page deliberately left blank

British Standards Institution (BSI)

BSI is the national body responsible for preparing British Standards and other standards-related publications, information and services.

BSI is incorporated by Royal Charter. British Standards and other standardization products are published by BSI Standards Limited.

About us

We bring together business, industry, government, consumers, innovators and others to shape their combined experience and expertise into standards -based solutions.

The knowledge embodied in our standards has been carefully assembled in a dependable format and refined through our open consultation process. Organizations of all sizes and across all sectors choose standards to help them achieve their goals.

Information on standards

We can provide you with the knowledge that your organization needs to succeed. Find out more about British Standards by visiting our website at bsigroup.com/standards or contacting our Customer Services team or Knowledge Centre.

Buying standards

You can buy and download PDF versions of BSI publications, including British and adopted European and international standards, through our website at bsigroup.com/shop, where hard copies can also be purchased.

If you need international and foreign standards from other Standards Development Organizations, hard copies can be ordered from our Customer Services team.

Subscriptions

Our range of subscription services are designed to make using standards easier for you. For further information on our subscription products go to bsigroup.com/subscriptions.

With British Standards Online (BSOL) you'll have instant access to over 55,000 British and adopted European and international standards from your desktop. It's available 24/7 and is refreshed daily so you'll always be up to date.

You can keep in touch with standards developments and receive substantial discounts on the purchase price of standards, both in single copy and subscription format, by becoming a BSI Subscribing Member.

PLUS is an updating service exclusive to BSI Subscribing Members. You will automatically receive the latest hard copy of your standards when they're revised or replaced.

To find out more about becoming a BSI Subscribing Member and the benefits of membership, please visit bsigroup.com/shop.

With a Multi-User Network Licence (MUNL) you are able to host standards publications on your intranet. Licences can cover as few or as many users as you wish. With updates supplied as soon as they're available, you can be sure your documentation is current. For further information, email bsmusales@bsigroup.com.

Revisions

Our British Standards and other publications are updated by amendment or revision.

We continually improve the quality of our products and services to benefit your business. If you find an inaccuracy or ambiguity within a British Standard or other BSI publication please inform the Knowledge Centre.

Copyright

Useful Contacts:

Customer Services
Tel: +44 845 086 9001
Email (orders): orders@bsigroup.com
Email (enquiries): cservices@bsigroup.com

Subscriptions
Tel: +44 845 086 9001
Email: subscriptions@bsigroup.com

Knowledge Centre
Tel: +44 20 8996 7004
Email: knowledgecentre@bsigroup.com

Copyright & Licensing
Tel: +44 20 8996 7070
Email: copyright@bsigroup.com

BSI Group Headquarters

389 Chiswick High Road London W4 4AL UK

bsi.