
**Information technology — JPEG 2000
image coding system —**

**Part 3:
Motion JPEG 2000**

**AMENDMENT 2: Motion JPEG 2000
derived from ISO base media file format**

*Technologies de l'information — Système de codage d'image JPEG
2000 —*

Partie 3: Motion JPEG 2000

*AMENDEMENT 2: Motion JPEG 2000 dérivée du format de fichier de
média de base de l'ISO*



PDF disclaimer

This PDF file may contain embedded typefaces. In accordance with Adobe's licensing policy, this file may be printed or viewed but shall not be edited unless the typefaces which are embedded are licensed to and installed on the computer performing the editing. In downloading this file, parties accept therein the responsibility of not infringing Adobe's licensing policy. The ISO Central Secretariat accepts no liability in this area.

Adobe is a trademark of Adobe Systems Incorporated.

Details of the software products used to create this PDF file can be found in the General Info relative to the file; the PDF-creation parameters were optimized for printing. Every care has been taken to ensure that the file is suitable for use by ISO member bodies. In the unlikely event that a problem relating to it is found, please inform the Central Secretariat at the address given below.

IECNORM.COM : Click to view the full PDF of ISO/IEC 15444-3:2002/Amd 2:2003

© ISO/IEC 2003

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
Case postale 56 • CH-1211 Geneva 20
Tel. + 41 22 749 01 11
Fax + 41 22 749 09 47
E-mail copyright@iso.org
Web www.iso.org

Published in Switzerland

Contents

Page

Foreword	iv
Introduction	v
1 Scope	1
2 Normative references	1
3 Terms and definitions	1
4 Compatibility and Technology derivation	2
4.1 Family Members	2
4.2 ISO Media file Inheritance and Compatibility	2
4.3 JP2 Inheritance and Compatibility	2
4.4 Conformance	2
4.5 Profiles and Levels	3
4.6 Visual Composition	3
4.7 Box Order	5
5 File Identification	5
6 Required Additions	5
6.1 Sample Description Box	5
7 Template Fields Used	7
Annex A (normative) File and Codestream profiles	9
A.1 Profile Introduction	9
A.2 Motion JPEG 2000 Simple Profile	9
Annex B (informative) Guidelines for use of the JPEG 2000 Codec	10
B.1 Introduction	10
B.2 Frequency Weighting for Motion Sequences	10
B.3 Encoder sub-sampling of components	11
Annex C (informative) Indicating Sub-sampling Chroma Offset	12
Annex D (informative) Field Structures for Interlace	14
Bibliography	17

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.

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 2.

The main task of the joint technical committee is to prepare International Standards. Draft International Standards adopted by the joint technical committee are circulated to national bodies for voting. Publication as an International Standard requires approval by at least 75 % of the national bodies casting a vote.

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.

Amendment 2 to ISO/IEC 15444-3:2002 was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 29, *Coding of audio, picture, multimedia and hypermedia information*.

Introduction

This document is the consequent revision of the Motion JPEG 2000 specification, based on the common text of the MP4 and MJ2 formats, which is called the ISO Base Media File Format; it will be consolidated subsequently into a new edition of ISO/IEC 15444-3:2002.

IECNORM.COM : Click to view the full PDF of ISO/IEC 15444-3:2002/Amd 2:2003

Information technology — JPEG 2000 image coding system —

Part 3: Motion JPEG 2000

AMENDMENT 2: Motion JPEG 2000 derived from ISO base media file format

1 Scope

This Amendment specifies the use of the wavelet-based JPEG 2000 codec for the coding and display of timed sequences of images (motion sequences), possibly combined with audio, and composed into an overall presentation. In this specification, a file format is defined, and guidelines for the use of the JPEG 2000 codec for motion sequences are supplied.

2 Normative references

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

ISO 639-2:1998, *Codes for the representation of names of languages — Part 2: Alpha-3 code*

ISO/IEC 14496-1:2001, *Information technology — Coding of audio-visual objects — Part 1: Systems* [particularly the syntax description language (SDL), clause 14]

ISO/IEC 15444-12, *Information technology — JPEG 2000 image coding system — Part 12: ISO base media file format* (technically identical to ISO/IEC 14496-12)

ITU-T Rec.T.800 | ISO/IEC 15444-1, *Information technology — JPEG 2000 image coding system: Core coding system*

3 Terms and definitions

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

3.1

motion sequence

timed sequence of JPEG 2000 images

4 Compatibility and Technology derivation

4.1 Family Members

This is a stand-alone specification; it defines the file format for MJ2. However, it stands as a member of a family of specifications with common formatting.

The other family members include:

- The JPEG 2000 single image format, JP2.
- The ISO Base Media File Format, on which the MP4 format is based.
- The QuickTime file format, on which the ISO Base Media format is based.

These specifications share a common definition for the structure of a file (a sequence of objects, called boxes here and atoms in QuickTime), and a common definition of the general structure of an object (the size and type).

All these specifications require that readers ignore objects that are unrecognizable to them.

This specification takes precedence over those from which it inherits, in any case where there are differences or conflicts; however no such conflicts are known to exist.

4.2 ISO Media file Inheritance and Compatibility

The Motion JPEG 2000 file format is defined as derived from the ISO Base Media file format. Notwithstanding anything in that base specification, hint tracks are not a normative part of this specification. Sections 7 (Streaming Support) and 10 (RTP Hint Track Format), though compatible with this specification, do not form normative parts of this specification. They may be used as a compatible, optional, extension, but are not required for compatibility with this specification. There may be license implications in the use of this or other compatible extensions to this format.

4.3 JP2 Inheritance and Compatibility

The still image format, JP2, defines a number of boxes. The signature box from that specification shall be present. If the JP2 specification requires a particular position (e.g. first in the file), that positioning shall be followed here.

It is permissible under this specification to make a file that adheres to both this specification and the JP2 specification. In that case:

- 1) The compatibility list shall include all the compatible brands.
- 2) The objects (boxes or atoms) required by the JP2 specification shall also be present.
- 3) The objects (boxes or atoms) optional in the JP2 specification may also be present.

A still image reader, reading a file which contains both a presentation (conformant to this specification) and a still image, would 'see' only the still image. Likewise a motion reader would 'see' only the presentation. A more powerful reader may display both, or offer the user a choice.

The JP2 specification includes an optional IPR (Intellectual Property Rights) box which is therefore also optional in this specification. Among other issues this addresses unique identification and protection of content.

4.4 Conformance

Implementations of Motion JPEG 2000 decoders shall support JPEG 2000 image sequences, as well as raw and twos-complement audio if audio output is available. They may also support compressed audio, using MP4 formats, or other track types from MPEG-4. The support of such MPEG-4 tracks is not required; however,

readers shall not fail if they are present. If MPEG-4 composition (BIFS) is used, then the simple composition used in this specification should also be set up in such a way that a reader not implementing BIFS will display a suitable result.

Files conformant with this specification shall contain at least one Motion JPEG 2000 video track. They may contain more video tracks, uncompressed audio, or compressed MP4 audio.

4.5 Profiles and Levels

There are two tools for profiling Motion JPEG 2000 files.

The first consists of the optional specification of tools and levels of the JPEG 2000 coding system (codestream features). These are indicated in the optional sample description extension JP2 Profile Box (see below 5).

The second tool allows a file overall to be identified as belonging to a definition which forms a proper subset of the general specification. Such definitions might restrict such features as:

- the use of data references, and multiple files;
- the layout order of the boxes, and the data within the boxes (e.g. that data is in time order and interleaved);
- the use of profiles of the JPEG 2000 codestream;
- the existence of other tracks, and their format (e.g. audio, MPEG-7, etc.).

The conformance to these restricted profiles is indicated in the file type box by the addition of the compatible profiles as brands within the compatibility list. Annex A defines the available profiles in this specification.

4.6 Visual Composition

Composition of multiple image sequences in a 2D environment can be achieved by using multiple video tracks which overlap in time. Their composition is defined by the following structures:

- The matrix in the track header specifies their positioning and scaling.
- The layer field in the track header specifies the front-to-back ordering of the tracks.
- The graphics mode and opcolor fields in the video media header are used to specify the ways in which each track is composited onto the existing image (this compositing is performed from back to front).

Applications requiring more complex compositing may use the BIFS system from MPEG-4, optionally. The matrix, graphics mode, and layers should be setup so that a reader not implementing BIFS displays the desired result. Matrix values which occur in the headers specify a transformation of video images for presentation. The point (p,q) is transformed into (p', q') using the matrix as follows:

$$\begin{pmatrix} p & q & 1 \end{pmatrix} * \begin{vmatrix} a & b & u \\ c & d & v \\ x & y & w \end{vmatrix} = \begin{pmatrix} m & n & z \end{pmatrix}$$

$$m = ap + cq + x; \quad n = bp + dq + y; \quad z = up + vq + w;$$

$$p' = m/z; \quad q' = n/z$$

The coordinates {p,q} are on the decompressed frame, and {p', q'} are at the rendering output. Therefore, for example, the matrix {2,0,0, 0,2,0, 0,0,1} exactly doubles the pixel dimension of an image. The co-ordinates transformed by the matrix are not normalized in any way, and represent actual sample locations. Therefore {x,y} can, for example, be considered a translation vector for the image.

The co-ordinate origin is located at the upper left corner, and X values increase to the right, and Y values increase downwards. {p,q} and {p',q'} are to be taken as absolute pixel locations relative to the upper left hand corner of the original image (after scaling to the size determined by the track header's width and height) and the transformed (rendering) surface, respectively.

Each track is composed using its matrix as specified into an overall image; this is then transformed and composed according to the matrix at the movie level in the MovieHeaderBox. It is application-dependent whether the resulting image is 'clipped' to eliminate pixels, which have no display, to a vertical rectangular region within a window, for example. So for example, if only one video track is displayed and it has a translation to {20,30}, and a unity matrix is in the MovieHeaderBox, an application may choose not to display the empty "L" shaped region between the image and the origin.

All the values in a matrix are stored as 16.16 fixed-point values, except for u, v and w, which are stored as 2.30 fixed-point values. For upwards compatibility into the MPEG-4 BIFS (scene composition) system, matrices used here restrict (u,v,w) to be (0,0,1), for which the hex values are (0,0,0x40000000). This permits the simple composition used here to be mapped into BIFS if a scene later requires full scene management.

The values in the matrix are stored in the order {a,b,u, c,d,v, x,y,w}.

Tracks are composed to the presentation surface from back (highest layer number) to front (lowest layer number), against an indeterminate initial colour. There are various composition modes available; the backmost (first-rendered) track would normally use 'copy' as the initial image is indeterminate. Subsequent layers can then be composed on top in a variety of ways. The following table details the composition modes available. Note that (currently) only the 'transparent' mode uses the opcolor field.

Table 1 — Graphics Composition Modes

Mode	Code	Description
Copy	0x0	Copy the source image over the destination
Transparent	0x24	Replace the destination pixel with the source pixel if the source pixel isn't equal to the opcolor. (Also known as 'blue-screen').
Alpha	0x100	Replace the destination pixel with a blend of the source and destination pixels, with the proportion controlled by the alpha channel. The alpha channel is applied to all channels.
Pre-multiplied black alpha	0x102	Pre-multiplied with black means that the colour components of each pixel have already been blended with a black pixel, based on their alpha channel value. Effectively, this means that the image has already been combined with a black background, which must be removed before composition.
Component alpha	0x110	One or more alpha channels are present, which are applied to individual colour channels, and the image must be composed channel-by-channel

Images are only alpha-composed if both the graphics composition mode requests alpha composition, and the images contain alpha channels, as declared by the Channel Definition Box inside the JP2 Header Box. Therefore the graphics mode can be used to prevent alpha composition of an image with alpha channels, if that is desired.

If there is a single alpha channel applied to the entire image, then the value of the graphics must be 'Alpha' if that channel is a straight 'Opacity' channel, and must be 'Pre-multiplied black alpha' if that channel is a 'Pre-multiplied' opacity channel. If there are one or more alpha channels in the image which are applied to individual channels and not to the whole image, and alpha composition is desired, then the 'Component alpha'

value must be used for the graphics mode. Support of 'Component alpha' composition is optional in Part 3 of this specification.

The alpha blending formulas are defined in Part 1 of this specification.

NOTE Use of the "transparent" opcode may yield unexpected results when the image codestreams are compressed in a non-reversible fashion, or are subject to scaling in quality or resolution, either during or after content production. Such operations are not guaranteed to preserve individual sample values precisely.

4.7 Box Order

All JPEG 2000 files start with a signature box, therefore this rule is added to the Box Order rules:

- 1) The JP2 Signature Box and File Type Box **shall** occur first and second in the file.

5 File Identification

In the file type compatibility box, the brand shall be 'mjp2' for files conforming to this specification, and 'mjp2' shall be a member of the compatibility list.

See Annex A for a complete list of all profile brand names.

The preferred file extension is '.mj2'. The MIME type video/mjp2 is used, as defined in the appropriate RFC.

6 Required Additions

6.1 Sample Description Box

6.1.1 Definition

Box Types: 'mjp2', 'raw', 'twos'
 Container: Sample Table Box ('stb')
 Mandatory: Yes
 Quantity: Exactly one

A Motion JPEG 2000 visual sample entry shall contain a JP2 Header Box from the JPEG 2000 Part 1 specification; however, the sRGB YCC enumerated colour space (codepoint 18) from Part 2 may also be used to identify the colour space used, in addition to the Part 1 enumerated colour spaces (such as sRGB and greyscale). If the JP2 Header Box indicates the presence of alpha channels, then the 'depth' field in the VisualSampleEntry must also indicate their presence, with the value 0x20. Similarly if the JP2 Header Box defines a monochrome image with no alpha, the 'depth' field must contain the value indicating grayscale (0x28). Otherwise, the 'depth' field must declare colour images (0x18).

If two fields are present in the samples, the JP2 Header Box applies to the complete image, not to each field individually. Therefore the height as declared in the JP2 Header Box and the VisualSampleEntry applies to the entire de-interlaced image.

NOTE This means that you may not be able to construct a legitimate JP2 file by composing a single field's codestream together with the JP2 header box found in the sample description.

The sample format for Motion JPEG 2000 data is a set of boxes. Currently this specification permits only JP2 Codestream Boxes ('jp2c') as defined in the JP2 specification. If there is no Field Coding Box present, or the field count is 1, the sample shall contain precisely one codestream box. If the field count is 2 then there shall be two codestream boxes. Other boxes, if present in the sample, shall be ignored. The last (or only) box in the sample may have a value of 0 for its length field, indicating it extends to the end of the sample, as indicated by the sample size given in the sample size table. The actual codestreams presented to the decoder are formed

by concatenating the contents of the JP2 Prefix Box, if any, in the sample description before each codestream presented in the jp2c box(es) in the samples. If field coding is used, the same prefix is concatenated before both fields. Typically, the prefix will contain a JPEG 2000 main header; however, this is not required in the general case, though specific profiles may limit the use of the prefix box.

If the codestreams used in a sequence conform to a specific profile of the JPEG 2000 coder, a JP2 Profile Box may be used to indicate such conformance.

The visual sample entry may optionally contain a field-ordering box (see below). If `fieldcount` is 2, each field will be half the height of the overall image, as declared in the 'height' field of the sample description. To be precise, if the height field contains the value H , then the field with the topmost scanline has $((H+1) \div 2)$ lines, and the other field has $(H \div 2)$ lines. The utility of the prefix box may be diminished when field coding is used, especially when H is odd.

The original format of the material (interlace or progressive) may be documented by the Original Format Box. The `original_fieldcount` must take the value 1 (progressive) or 2 (interlaced). When the value is 2, the `original_fieldorder` documents whether the topmost line came from the earlier or later field. The documentation of the original format is independent of whether the material has been coded as progressive (frame-based) or interlaced, which is documented by the field coding box. This box is informative only and may assist readers in display or transcoding. For example, material that was originally interlaced but is encoded frame-based and must be displayed on an interlaced display can be positioned so that the display interlace matches the original material's interlace.

The values present in the `VisualSampleEntry`, its constituent boxes including the JP2 Header Box, and the codestreams that these boxes describe, must agree, to the extent that the format and precision of fields allow. This agreement includes, but is not limited to, width and height information, and the resolution declaration (within the accuracy permitted by the different representations). Files with conflicts are non-conforming and readers may attempt to decide which values are correct, or reject the file.

The fields `horizresolution` and `vertresolution` document the highest resolution component of the image (which is typically, but not required to be, the luminance, in a sub-sampled image).

For audio tracks, in the formats defined here (with code points 'raw' and 'twos'), the data is stored as uncompressed samples. If stereo is stored, the data consists of interleaved left/right samples. The raw format uses offset-binary; for 8-bit samples, values range from 0 to 255, with 128 indicating silence. For 'twos', 8-bit values range from -128 to 127, with 0 being silence. Values of 16 bits are similarly derived (with the bytes in network byte order, or big-endian format).

6.1.2 Syntax

```
// Visual Sequences
class MJ2SampleEntry() extends VisualSampleEntry ('mjp2'){
    JP2HeaderBox();
    FieldCodingBox(); // optional
    MJP2ProfileBox(); // optional
    MJP2PrefixBox(); // optional
    MJP2SubSamplingBox(); // optional
    MJP2OriginalFormatBox(); // optional
}

// Field-Based Coding
class FieldCodingBox() extends Box('fiel'){
    int(8) fieldcount;
    int(8) fieldorder; // both storage and temporal order
}

class MJP2OriginalFormatBox() extends Box('orfo'){
    int(8) original_fieldcount;
    int(8) original_fieldorder;
}
```

```

class MJP2ProfileBox() extends FullBox('jp2p', 0, 0){
    unsigned int(32)[] compatible_brands;
}

class MJP2PrefixBox() extends Box('jp2x'){
    int(8)[] data; // the data is the initial codestream part
}

class MJP2SubSamplingBox () extends Box('jsub'){
    unsigned int(8) horizontal_sub;
    unsigned int(8) vertical_sub;
    unsigned int(8) horizontal_offset;
    unsigned int(8) vertical_offset;
}

// Audio Sequences

class MJ2AudioSampleEntry() extends AudioSampleEntry (AudioFormat){
}

```

6.1.3 Semantics

`AudioFormat` is either 'raw' or 'twos'.

`Compressorname` the value "\017Motion JPEG 2000" is recommended (\017 is 15, the length of the string as a byte)

`depth` takes one of the following values

- 0x18 – images are in colour with no alpha
- 0x28 – images are in grayscale with no alpha
- 0x20 – images have alpha (gray or colour)

`compatible_brands` is a list, filled to the end of the containing box, of JPEG 2000 profiles, to which the associated codestreams conform; see the JPEG 2000 specification for the defined values.

`Horizontal_sub` and `vertical_sub` indicate whether the chroma components of a YCbCr encoding were downsampled in the codestream; the value indicates the number of luminance samples to a single chroma sample in the given direction. This can assist decoders in memory allocation, or in using optimized sub-sampled display interfaces.

`Horizontal_offset` and `vertical_offset` specify the offset of the first chroma sample from the first luminance sample, as measured on the sample grid. If a CRG marker is present in the codestream, these values take precedence over those in the codestream. See Annex C for example values.

`fieldcount` specifies the number of fields in the samples, and shall be 1 or 2

`fieldorder` describes the order of the two fields, and is only relevant if `fieldcount` equals 2:

- 0 field coding unknown
- 1 field with the topmost line is stored first in the sample; fields are in temporal order
- 6 field with the topmost line is stored second in the sample; fields are in temporal order

`original_fieldcount` specifies the number of fields in original material before encoding, and shall be 1 or 2

`original_fieldorder` describes the order of the two fields, and is only relevant if

`original_fieldcount` equals 2:

- 0 field coding unknown
- 11 topmost line came from the earlier field
- 16 topmost line came from the later field

7 Template Fields Used

Motion JPEG 2000 uses simple composition, as defined above. Therefore the following fields are required to be correctly set in conforming files:

- 1) For Visual composition: `matrix` in the Movie Header Box and Track Header Box; `layer`; `graphicsmode` and `opcolor` in the Video Media Header Box; and `depth` in a VisualSampleEntry;

- 2) For Audio Composition: `volume` in the Movie Header Box and Track Header Box; `balance` in the Sound Media Header Box;
- 3) For variable-rate playback, `rate` in the Movie Header Box;
- 4) To describe the audio and visual sample correctly, `horizresolution`, and `vertresolution` in a `VisualSampleEntry`; and `channelcount`, `samplesize` and the `samplerate` in an `AudioSampleEntry`.

IECNORM.COM : Click to view the full PDF of ISO/IEC 15444-3:2002/Amd 2:2003

Annex A (normative)

File and Codestream profiles

A.1 Profile Introduction

This Annex normatively specifies the ways in which Motion JPEG 2000 files can be profiled.

There are three areas for which profiles and levels can be declared:

- a) The optional still image which may occur at the top-level of a Motion JPEG 2000 file.
- b) The characteristics of the individual frames (codestreams) which form a motion sequence.
- c) The overall characteristics of the Motion Presentation itself.

The profile for the optional still image is declared by including the correct 'compatible brands' in the top-level File Type Box. These profiles (brands) are specified for JP2 files, and that specification shall be followed here.

Profile code-points should be used to profile the images that form the frames of motion sequences. These code-points are placed within the JP2 Profile Box within the sample description.

The characteristics of the motion presentation are declared by placing the brands of the compatible profiles in the top-level file-type box. This specification defines two profiles:

- 1) Unrestricted. The brand 'mjp2' indicates unrestricted conformance to this specification and shall be placed in the top-level File Type Box in all files.
- 2) Simple. The brand 'mj2s' indicates the simple Motion JPEG 2000 profile, as defined below.

A.2 Motion JPEG 2000 Simple Profile

File conforming to the simple profile have the following characteristics:

- 1) Exactly one video track is present.
- 2) At most a single audio track, using only 8 or 16-bit raw audio, is present.
- 3) Each track shall have exactly one sample description, used by all samples.
- 4) The sample rate of the audio, if present, shall not exceed 48 kHz.
- 5) The frame rate of the video shall not exceed 30 frames per second.
- 6) The video codestream profile shall be profile 0 for both the motion sequence and the still image, if present.
- 7) The file is self-contained; no data references are used, and therefore all media data is contained within the single file.
- 8) The media data in the Media Data Box(es) is placed within the box(es) in temporal order.
- 9) If more than one track is present, the media data for the tracks is interleaved, with a granularity no greater than the greater of (a) the duration of a single 'sample' (in file format terms) or (b) one second.
- 10) The transformation matrices used are restricted to uniform scaling and rotation by multiples of 90°.

Annex B (informative)

Guidelines for use of the JPEG 2000 Codec

B.1 Introduction

Certain artifacts have been observed when using the 15444-1 codec to create motion sequences. Experiments have been performed to identify the cause of these artifacts and indicate encoding strategies that would be likely to mitigate them. This Annex is intended to review the results of the experiments that were performed and recommend possible methods that may be employed. These methods are not necessarily the best or only approach to addressing these observed artifacts. This Annex is intended as a point of reference for individuals implementing 15444-3 (Motion JPEG 2000).

B.2 suggests particular frequency weighting for motion sequences. B.3 addresses sub-sampling issues.

B.2 Frequency Weighting for Motion Sequences

The objective of employing frequency weighting to motion images is to explore methods that might improve subjective quality. Tuning of weighting values for chroma component is an important area to investigate. An experiment showed that there are some flickering artifacts when displaying a moving sequence of images although individual pictures do not show an annoying artifact. The frequency weighting was observed to reduce these artifacts in all test sequences, especially at low to middle bitrate (for example 0.25 to 0.50bpp for CIF). A finer and better adjustment of visual weights to improve the image sequence's quality may provide additional improvements.

For colour images the frequency weighting tables of the Y, Cr, and Cb components should differ in order to take advantage of the properties of the human visual system. For example, it is usually desirable to emphasize the luminance component more than the chrominance components.

Table B.1 specifies a set of CSF weights which were designed for the luminance component based on the CSF value at the mid-frequency of each subband. The other two tables (Table B.2 and Table B.3) specify sets of CSF weights for Cb and Cr. Note that the tables are intended for a 5-level wavelet decomposition.

The table does not include the weight for the lowest frequency subband, nLL, which is always 1. Levels 1, 2, ..., 5 denote the subband levels in low to high frequency order. (HL, LH, HH) denotes the three frequency orientations within each subband.

Table B.1 — CSF weights for Luminance

Level	HL	LH	HH
1	1.000000	1.000000	1.000000
2	1.000000	1.000000	1.000000
3	0.999994	0.999994	0.999988
4	0.837755	0.837755	0.701837
5	0.275783	0.275783	0.090078

Table B.2 — CSF weights for Cb

Level	HL	LH	HH
1	0.812612	0.812612	0.737656
2	0.679829	0.679829	0.567414
3	0.488887	0.488887	0.348719
4	0.267216	0.267216	0.141965
5	0.089950	0.089950	0.027441

Table B.3 — CSF weights for Cr

Level	HL	LH	HH
1	0.856065	0.856065	0.796593
2	0.749805	0.749805	0.655884
3	0.587213	0.587213	0.457826
4	0.375176	0.375176	0.236030
5	0.166647	0.166647	0.070185

B.3 Encoder sub-sampling of components

In the most of all cases, the format of motion sequences is either YcbCr (BT-601-5 for standard TV and BT-709-4 for HDTV) or YUV. In these, colour components are already sub-sampled and the sub-sampled signal may be the input of the Motion JPEG 2000 encoder. Although the encoder sub-sampling of colour components is not recommended in 15444-1, the sub-sampling may be commonly used in Motion JPEG 2000. Therefore, CSF weights also should be weighted directly to sub-sampled colour components.

Sub-sampling of chroma components in JPEG 2000 can be achieved in two ways: (a) the highest-frequency sub-bands may be omitted, and (b) if no ICT is used, then the codestream permits the components to differ in size.

Annex C (informative)

Indicating Sub-sampling Chroma Offset

This annex provides informative material for the Sub-sampling Box 'jsub' (see 6.1), with offset values for common sub-sampling formats.

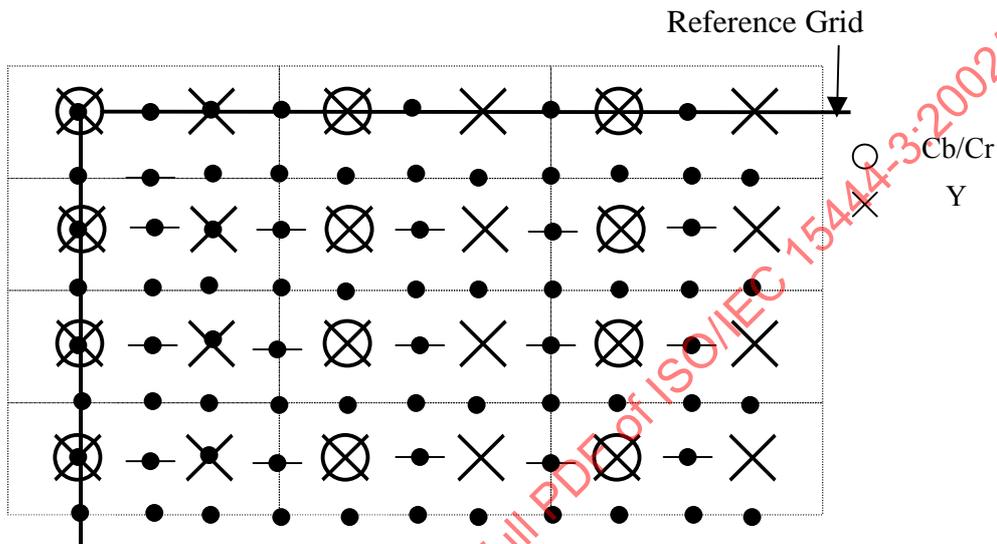


Figure C.1 — 4:2:2 format (co-sited)

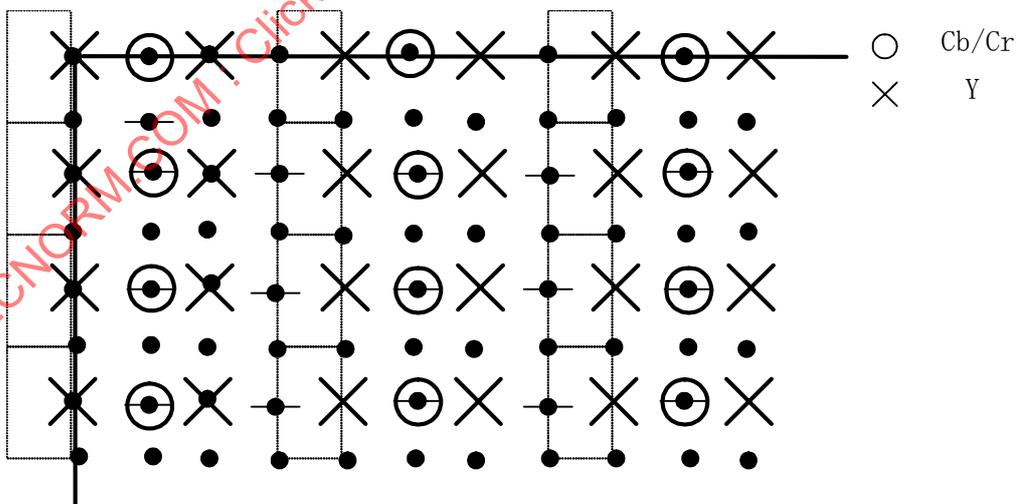


Figure C.2 — 4:2:2 format (centered)

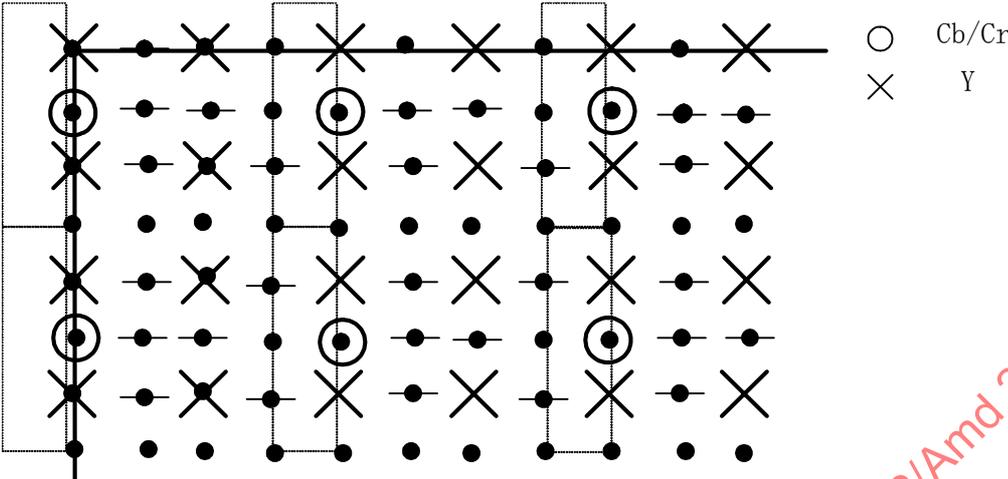


Figure C.3 — 4:2:0 format (co-sited)

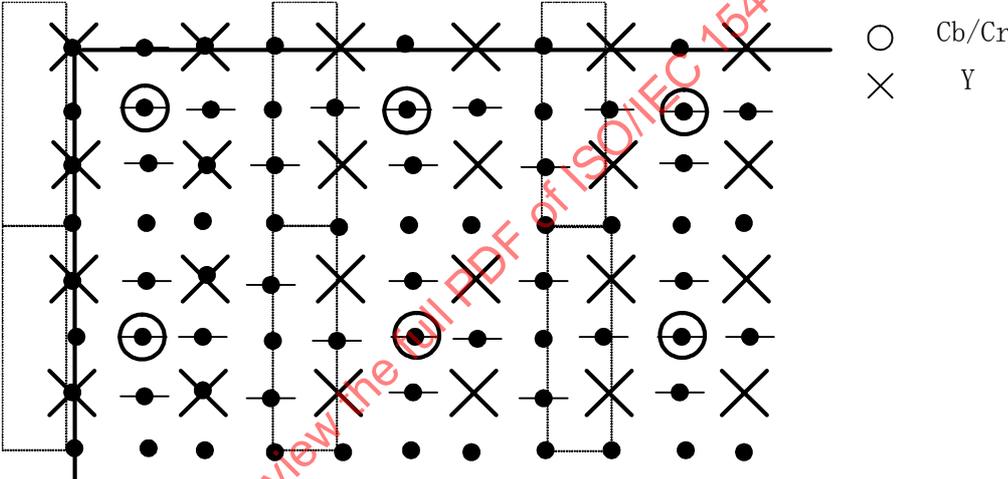


Figure C.4 — 4:2:0 format (centered)

Table C.1 — Chroma Phase Values

	Figure C.1	Figure C.2	Figure C.3	Figure C.4
Horizontal-offset	0	1	0	1
Vertical-offset	0	0	1	1