



INTERNATIONAL TELECOMMUNICATION UNION

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

T.805

(01/2012)

SERIES T: TERMINALS FOR TELEMATIC SERVICES
Still-image compression – JPEG 2000

Information technology - JPEG 2000 image coding system: Compound image file format

CAUTION !

PREPUBLISHED RECOMMENDATION

This prepublication is an unedited version of a recently approved Recommendation. It will be replaced by the published version after editing. Therefore, there will be differences between this prepublication and the published version.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications, information and communication technologies (ICTs). The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

Compliance with this Recommendation is voluntary. However, the Recommendation may contain certain mandatory provisions (to ensure, e.g., interoperability or applicability) and compliance with the Recommendation is achieved when all of these mandatory provisions are met. The words "shall" or some other obligatory language such as "must" and the negative equivalents are used to express requirements. The use of such words does not suggest that compliance with the Recommendation is required of any party.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the Recommendation development process.

As of the date of approval of this Recommendation, ITU [had/had not] received notice of intellectual property, protected by patents, which may be required to implement this Recommendation. However, implementers are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database at <http://www.itu.int/ITU-T/ipr/>.

© ITU 2012

All rights reserved. No part of this publication may be reproduced, by any means whatsoever, without the prior written permission of ITU.

Information technology - JPEG 2000 image coding system: Compound image file format

Summary

This Recommendation | International Standard defines a normative but optional file format for storing compound images using the JPEG 2000 file format family architecture. This format is an extension of the JP2 file format defined in Recommendation ITU-T T.800 | ISO/IEC 15444-1 Annex I and uses boxes defined for both the JP2 file format and the JPX file format defined in Recommendation ITU-T T.801 | ISO/IEC 15444-2 Annex M. This Recommendation | International Standard is useful for applications storing multiple pages, images with mixed content, and/or images that need more structure than provided in JP2.

Applications that implement this file format shall implement it as described in this Recommendation | International Standard. This Recommendation | International Standard:

- specifies a binary container for multiple bi-level and continuous-tone images used to represent a compound image,
- specifies a mechanism by which multiple images can be combined into a single compound image, based on the Mixed Raster Content model,
- specifies a mechanism for grouping multiple images in a hierarchy of layout objects, pages and page collections,
- specifies a mechanism for storing JPEG 2000 and other compressed image data formats,
- specifies a mechanism by which metadata can be included in files specified by this Recommendation | International Standard.

1

CONTENTS

Page

Foreword	5
1 Scope	6
2 Normative references	6
2.1 Identical Recommendations International Standards	6
2.2 Additional references	7
3 Terms and Definitions	8
4 Abbreviations	9
5 General arrangement	9
5.1 Mixed Raster Content Model	9
5.2 File Elements and Structure	10
5.2.1 Page Collections	11
5.2.2 Pages	13
5.2.3 Layout Objects	13
5.2.4 Thumbnails	17
5.2.5 Contiguous and Fragmented Codestreams	18
5.2.6 Shared Data	18
5.2.7 Metadata	18
5.3 Hidden Text Metadata	18
5.4 JPM Use Scenarios	19
5.4.1 Self-Contained Files and Files with External References	19
5.4.2 Files prepared for the Internet	19
5.4.3 Page Collection Example	19
5.4.4 Page Collections in a Client / Server Context	20
5.4.5 Example of a Client/Server Interaction	20
5.4.6 Example of Scanner Capture	21
Annex A Compound Image File Structure	22
A.1 File Identification	22
A.2 File Organization	22
A.3 Box Definition	24
A.4 Boxes Used in a Compound Image File	24
Annex B Box Definitions	27
B.1 File Level Boxes	27
B.1.1 JPEG 2000 Signature box	27
B.1.2 File Type box	27
B.1.3 JP2 Header box (superbox) after File Type box	27
B.1.4 Compound Image Header box	27
B.1.5 Data Reference box	29
B.1.6 Page Collection box (superbox)	30
B.1.6.1 Primary Page Collection Locator box	30
B.1.6.2 Page Table box	31
B.1.7 Shared Data Entry box	32
B.1.8 Shared Data Reference box	33
B.2 Page Level Boxes	33
B.2.1 Page box (superbox)	33
B.2.1.1 Page Header box	34
B.3 Layout Object Level Boxes	35
B.3.1 Layout Object box (superbox)	35
B.3.1.1 Layout Object Header box	35
B.4 Object Level Boxes	36
B.4.1 Object box (superbox)	36
B.4.1.1 Object Header box	37
B.4.1.2 Object Scale box	38
B.4.1.3 JP2 Header box (superbox) in Object box	39

B.5	JP2 Codestream Element Boxes	39
B.5.1	Fragment Table box (superbox).....	39
B.5.1.1	Fragment List box	39
B.5.2	Media Data box.....	40
B.5.3	Contiguous Codestream box	40
B.6	General / Common Boxes	41
B.6.1	Base Colour box (superbox)	41
B.6.1.1	Base Colour Value box	41
B.6.2	JP2 Header box (superbox).....	42
B.6.2.1	Image Header box	43
B.6.2.2	Colour Specification box	44
B.6.2.2.1	EP field format for the CIELab colourspace.....	48
B.6.3	Label box	49
B.6.4	Cross-Reference box.....	49
B.6.5	Hidden Text Metadata box (superbox)	50
B.6.6	HTX Reference box	50
B.6.7	Free box	51
Annex C	Metadata	52
C.1	Adding intellectual property rights information in JPM	52
C.2	Adding vendor specific information to the JPM file format	52
C.2.1	XML boxes	52
C.2.2	UUID boxes	52
C.2.3	UUID Info boxes (superbox)	52
Annex D	Profiles	53
D.1	JPM Profiles.....	53
D.1.1	Web Profile	53
D.2	Decompression Profiles	54
Annex E	Guidelines for Constructing URLs for JPM Files	55
E.1	Pages and Layout Objects	55
E.2	Metadata boxes	55
E.3	Labels.....	55
E.4	Page Collections	55
E.5	Page Thumbnails.....	56
E.6	Document Thumbnail	56
E.7	Byte Ranges	56
Annex F	Hidden Text and Annotations Storage.....	57
F.1	Storage of HTX in JPM	57
F.2	Compression of HTX.....	57
Annex G	Hidden Text and Annotations Types and Elements	58
G.1	Overview.....	58
G.2	Types.....	59
G.2.1	Shape.....	59
G.2.2	Coordinates	59
G.2.3	Percentage.....	59
G.2.4	Angle.....	59
G.2.5	Resolution	60
G.3	Common Attributes.....	61
G.3.1	Core Attributes.....	61
G.3.2	Position Attributes	62
G.4	Elements.....	63
G.4.1	HTX	63
G.4.2	Parameter	63
G.4.3	Hidden Text	64
G.4.4	Region.....	64
G.4.5	Paragraph	66
G.4.6	Line	67
G.4.7	Word	68
G.4.8	Alternative Word	69

G.4.9	Character.....	70
G.4.10	Alternative Character.....	70
G.4.11	Snippet.....	72
G.4.12	Annotations.....	72
G.4.13	Area.....	73
Annex H Hidden Text and Annotations Schema		74
H.1	XML Schema.....	74
H.1.1	Version and Encoding.....	74
H.1.2	Schema.....	74
Annex I Hidden Text and Annotations Examples.....		75
I.1	Example 1	75
I.2	Example 2	77
I.3	Example 3	77

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.

ISO/IEC 15444-6 was prepared by Joint Technical Committee ISO/IEC JTC 1, Information technology, Subcommittee SC 29, Coding of audio, picture, multimedia and hypermedia information.

ISO/IEC 15444 consists of the following parts, under the general title Information technology — JPEG 2000 image coding system:

- Part 1: Core coding system
- Part 2: Extensions
- Part 3: Motion JPEG 2000
- Part 4: Conformance testing
- Part 5: Reference software
- Part 6: Compound image file format
- Part 8: Secure JPEG 2000
- Part 9: Interactivity tools, APIs and protocols
- Part 10: Extensions for three-dimensional data
- Part 11: Wireless
- Part 12: ISO base media file format
- Part 13: An entry level JPEG 2000 encoder
- Part 14: XML structural representation and reference

INTERNATIONAL STANDARD

ISO/IEC 15444-6:201X (E)

Rec. ITU-T T.805 (201XE)

ITU-T RECOMMENDATION

Information technology – JPEG 2000 image coding system: Compound image file format

1 Scope

This Recommendation | International Standard defines a normative but optional file format for storing compound images using the JPEG 2000 file format family architecture. This format is an extension of the JP2 file format defined in Recommendation ITU-T T.800 | ISO/IEC 15444-1 Annex I and uses boxes defined for both the JP2 file format and the JPX file format defined in Recommendation ITU-T T.801 | ISO/IEC 15444-2 Annex M. This Recommendation | International Standard is useful for applications storing multiple pages, images with mixed content, and/or images that need more structure than provided in JP2.

Applications that implement this file format shall implement it as described in this Recommendation | International Standard. This Recommendation | International Standard:

- specifies a binary container for multiple bi-level and continuous-tone images used to represent a compound image,
- specifies a mechanism by which multiple images can be combined into a single compound image, based on the Mixed Raster Content model,
- specifies a mechanism for grouping multiple images in a hierarchy of layout objects, pages and page collections,
- specifies a mechanism for storing JPEG 2000 and other compressed image data formats,
- specifies a mechanism by which metadata can be included in files specified by this Recommendation | International Standard.

2 Normative references

The following Recommendations and International Standards contain provisions which, through reference in this text, constitute provisions of this Recommendation | International Standard. At the time of publication, the editions indicated were valid. All Recommendations and Standards are subject to revision, and parties to agreements based on this Recommendation | International Standard are encouraged to investigate the possibility of applying the most recent edition of the Recommendations and Standards listed below. Members of IEC and ISO maintain registers of currently valid International Standards. The Telecommunication Standardization Bureau of the ITU maintains a list of currently valid ITU-T Recommendations.

2.1 Identical Recommendations | International Standards

- ISO 5807:1985, *Information processing – Documentation symbols and conventions for data, program and system flowcharts, program network charts and system resources charts*
- ISO 8601:2000, *Data elements and interchange formats – Information interchange – Representation of dates and times*
- ISO/IEC 646:1991, *Information technology – ISO 7-bit coded character set for information interchange*
- ISO/IEC 8859-1:1998, *Information technology – 8-bit single-byte coded graphic character sets – Part 1: Latin alphabet No. 1*
- ISO/IEC 11578:1996, *Information technology – Open Systems Interconnection – Remote Procedure Call (RPC)*
- ISO 3166-1:1997, *Codes for the representation of names of countries and their subdivisions – Part 1: Country codes*
- ISO 3166-2:1998, *Codes for the representation of names of countries and their subdivisions – Part 2: Country subdivision code*
- Recommendation ITU-T T.4, *Standardization of Group 3 facsimile terminals for document transmission*

- Recommendation ITU-T T.6, *Facsimile coding schemes and coding control functions for group 4 facsimile apparatus*
- Recommendation ITU-T T.42, *Continuous-tone colour representation method for facsimile*
- Recommendation ITU-T T.44 | ISO/IEC 16485:2000, *Information technology – Mixed Raster Content (MRC)*
- Recommendation ITU-T T.44 Amendment 1, *Accommodation of new Annex B*
- Recommendation ITU-T T.45, *Run-length colour encoding*
- Recommendation ITU-T T.81 | ISO/IEC 10918-1:1994, *Information technology – Digital compression and coding of continuous-tone still images: Requirements and guidelines*
- Recommendation ITU-T T.82 | ISO/IEC 11544:1993, *Information technology – Coded representation of picture and audio information – Progressive bi-level image compression*
- Recommendation ITU-T T.83 | ISO/IEC 10918-2:1995, *Information technology – Digital compression and coding of continuous-tone still images: Compliance testing*
- Recommendation ITU-T T.84 | ISO/IEC 10918-3:1997, *Information technology – Digital compression and coding of continuous-tone still images: Extensions*
- Recommendation ITU-T T.84 Amendment 1 | ISO/IEC 10918-3:1997/Amd 1:1999, *Provisions to allow registration of new compression types and versions in the SPIFF header*
- Recommendation ITU-T T.86 | ISO/IEC 10918-4:1999, *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 (REGAUT)*
- Recommendation ITU-T T.87 | ISO/IEC 14495-1:2000, *Information technology – Lossless and near-lossless compression of continuous-tone still images – Baseline*
- Recommendation ITU-T T.88 | ISO/IEC 14492:2001, *Information technology – Lossy/lossless coding of bi-level images*
- Recommendation ITU-T T.89, *Application profiles for Recommendation T.88 – Lossy/lossless coding of bi-level images (JBIG2) for facsimile*
- Recommendation ITU-T T.800 | ISO/IEC 15444-1:2000, *Information technology – JPEG 2000 image coding system – Part 1: Core coding system*
- Recommendation ITU-T T.801 | ISO/IEC 15444-2, *Information technology – JPEG 2000 image coding system – Part 2: Extensions*
- Recommendation ITU-T T.803 | ISO/IEC 15444-4, *Information technology – JPEG 2000 image coding system – Part 4: Conformance testing*

2.2 Additional references

- IEEE Std. 754-1985 R1990, *IEEE Standard for Binary Floating-Point Arithmetic*
- IETF RFC 1766, *Tags for the Identification of Languages*, March 1995
- IETF RFC 2279, *UTF-8, A transformation format of ISO 10646*, January 1998
- IETF RFC 1950, *ZLIB Compressed Data Format Specification version 3.3*, May 1996
- IETF RFC 1951, *DEFLATE Compressed Data Format Specification version 1.3*, May 1996
- IETF RFC 2045, *Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies*
- IETF RFC 2396, *Uniform Resource Identifiers (URI): Generic Syntax*, August 1998
- ICC.1:1998-09, *International Color Consortium, File Format for Color Profiles*
- IEC 61966-2-1:1999-10, *Multimedia systems and equipment – Colour measurement and management – Part 2-1: Colour management – Default RGB colour space – sRGB*
- IEC 61966-2-1/Amd.1: *Multimedia systems and equipment – Colour measurement and management – Part 2-1: Colour management – Default RGB colour space – sRGB*
- W3C, *Extensible Markup Language (XML) 1.0, Fifth Edition*, <<http://www.w3.org/TR/REC-xml>>
- W3C, *Cascading Style Sheets, level 1 (CSS1) Specification*, <<http://www.w3.org/pub/TR/REC-CSS1>>
- W3C, *Cascading Style Sheets, level 2 (CSS2) Specification*, <<http://www.w3.org/TR/REC-CSS2>>
- W3C, *HTML 4.01 Specification*, <<http://www.w3.org/TR/html401>>
- W3C, *XHTML 1.0 Extensible HyperText Markup Language, Second Edition*, <<http://www.w3.org/TR/xhtml1>>
- W3C, *XML Schema Part 0: Primer, Second Edition*, <<http://www.w3.org/TR/xmlschema-0>>
- W3C, *XML Schema Part 1: Structures, Second Edition*, <<http://www.w3.org/TR/xmlschema-1>>

3 Terms and Definitions

For the purposes of this Recommendation | International Standard, the following definitions apply.

- 3.1 base colour:** The colour of an object for which no image data is available.
- 3.2 BasePage:** The original state of the page before it is rendered with layout objects.
- 3.3 box:** A portion of the file format defined by a length and a unique box type. Boxes of some types may contain other boxes.
- 3.4 component:** A two-dimensional array of samples.
- 3.5 compound image:** An image that may contain scanned images, synthetic images or both, and that preferably requires a mix of continuous tone and bi-level compression methods.
- 3.6 file format:** A codestream or codestreams and additional support and information not explicitly required for decoding of the codestream or codestreams. Examples of such support data include text fields providing security and historical information, data to support the placement of multiple codestreams within a given data file, and data to support exchange between platforms or conversions to other file formats.
- 3.7 fragment:** A portion of the codestream for an image. Subclause 5.2.5 describes fragment usage.
- 3.8 JP2 file:** The name of a file in the file format described in Recommendation ITU-T T.800 | ISO/IEC 15444-1. Structurally, a JP2 file is a contiguous sequence of boxes.
- 3.9 JPM file:** The name of a file in the file format described in this International Standard. A JPM file can contain one or more pages, composed from one or more layout objects, each of which is composed from at most two objects. Structurally, a JPM file is a contiguous sequence of boxes.
- 3.10 JPX file:** The name of a file in the file format described in Recommendation ITU-T T.801 | ISO/IEC 15444-2. Structurally, a JPX file is a contiguous sequence of boxes.
- 3.11 layout object:** An entity that comprises at most two paired objects or MRC layers.
- 3.12 main page collection:** The main page collection contains all pages and page collections in a file.
- 3.13 mask object:** An object that is used to select the samples of a corresponding image object that are to be imaged on a page.
- 3.14 metadata:** Additional data associated with the image data beyond the image data.
- 3.15 MRC:** Mixed Raster Content; a multi-layer imaging model described in Recommendation ITU-T T.44 | ISO/IEC 16485.
- 3.16 object:** An image that is part of a layout object; an MRC layer.
- 3.17 page:** The largest collection of layout objects that can be imaged independently of any other layout objects; a canvas or frame for imaging.
- 3.18 page collection:** A collection of pages logically grouped together in a JPM file. Each page must be contained in at least one page collection.
- 3.19 primary page collection:** A page collection which provides back and forward navigation in the main document associated with a page.
- 3.20 PageImage:** The image created by rendering the BasePage with the layout objects. The PageImage_k is the images created by rendering the BasePage with the first k layout objects.
- 3.21 profile:** A subset of all possible field values in a file.
- 3.22 superbox:** A box that itself contains a contiguous sequence of boxes (and only a contiguous sequence of boxes).
- 3.23 hidden text:** symbolic representation for the characters and words found in an image
- 3.24 annotation:** particular region of a page in a JPM document that has associated a URL reference, a note or a highlight
- 3.25 hidden text XML:** XML data which describe hidden text and annotations for a single page in a JPM file and which conform to the schema in Annex H
- 3.26 compressed hidden text XML:** hidden text XML data compressed using the mechanisms defined in F.2
- 3.27 hidden text UUID box:** UUID box containing compressed hidden text XML
- 3.28 hidden text XML Schema:** XML Schema for hidden text XML, as defined in H.1

4 Abbreviations

For the purposes of this Recommendation | International Standard, the following abbreviations apply. The abbreviations defined in Recommendation ITU-T T.800 | ISO/IEC 15444-1 Clause 4 also apply.

CCITT	International Telegraph and Telephone Consultative Committee (organ of the ITU up to 28 February 1993)
DPI	Dots per inch
IPR	Intellectual Property Rights
JP2	JPEG 2000 File Format defined in Recommendation ITU-T T.800 ISO/IEC 15444-1
JPX	JPEG 2000 File Format defined in Recommendation ITU-T T.801 ISO/IEC 15444-2; JPEG 2000 File Format Extended
JPM	JPEG 2000 File Format defined in this International Standard; JPEG 2000 File Format - Multilayer
MRC	Mixed Raster Content
UUID	Universal Unique Identifier
HTX	Hidden Text XML

5 General arrangement

The purpose of this clause is to give an overview of this International Standard. Terms defined in previous clauses in this International Standard will also be introduced. (Terms defined in Clauses 3 and 4 in Recommendation ITU-T T.800 | ISO/IEC 15444-1 continue to apply in this International Standard.) Throughout this International Standard, text formatted as a NOTE in the following form is informative only:

NOTE - Informative text appears here.

This International Standard defines a file format for storing compound images using the JPEG 2000 file format family architecture. A compound image file contains multiple images, both contiguous tone and bi-level, together with composition models describing how the individual images are combined to generate the compound image. This International Standard is based on the multi-layer Mixed Raster Content (MRC) imaging model, defined in Recommendation ITU-T T.44 | ISO/IEC 16485.

This International Standard defines a member of the JPEG 2000 file format family that enables the efficient processing, interchange and archiving of raster-oriented pages containing a mixture of multi-level and bi-level images. This efficiency is realized by representing the mixed-content image using multiple layers, determined by image type, and applying image specific encoding, spatial and colour resolution processing. A rasterized page may contain one or more image types, such as: multi-level continuous-tone or palettized (contone) content usually associated with naturally occurring images; bi-level detail associated with text and line-art; and multi-level colours associated with the text and line-art. This International Standard makes provisions for processing, interchange, and archiving of these image types in multiple layers and defines composition models which regenerate the desired image.

5.1 Mixed Raster Content Model

A file that conforms with this International Standard contains one or more pages. The PageImage associated with a page is generated by combining the page's layout objects with the BasePage.

The BasePage is the initial PageImage before any layout objects have been rendered. The BasePage has the same width and height as the page and is either transparent or filled with a single colour. BasePage is defined in 5.2.2.1. The Layout Objects are applied sequentially, in an order defined by the Layout Object Identifier field in the Layout Object Header boxes, to the BasePage to create the final PageImage. The Layout Object with Layout Object Identifier value of 0 is the page thumbnail and is not used in creating the PageImage.

Associated with each Layout Object is a mask M and an image I . The mask M is an opacity image and has only one component; I can be greyscale or colour, with one or more components. M and I are defined in Subclause 5.2.3. Both M and I have the same width and height as the page.

The following equations show the model for combining the BasePage and a sequence of n layout objects with non-zero Layout Object Identifier values to create the final PageImage. We will use the notation $N[c][x,y]$ to refer to the sample value at position (x,y) in component c of an image N .

$$\text{PageImage}_0[c][x,y] = \text{BasePage}[c][x,y] \quad (1)$$

$$\text{PageImage}_m[c][x,y] = \frac{s_m - M_m[c][x,y] \times \text{PageImage}_{m-1}[c][x,y] + M_m[0][x,y] \times I_m[c][x,y]}{S_m}, \quad m = 1, \dots, n \quad (2)$$

$$\text{PageImage}[c][x,y] = \text{PageImage}_0[c][x,y]$$

where $M_m[c][x,y]$ and $I_m[c][x,y]$ are the image sample values of component c at position (x,y) of the m^{th} Layout Object's mask M and image I respectively, and $s_m = 2^{\text{bit depth of } M_m} - 1$. M , I and BasePage are defined in 5.2.3 and 5.2.2.1. PageImage[c][x,y] is the final page image, obtained after combining all n layout objects associated with the page.

NOTE - Figure 1 shows a simple example of a PageImage constructed from a BasePage with a single solid colour and two Layout Objects. Note that white in the masks M_0 and M_1 denotes a value of 0.

NOTE - In a JPM file, the mask and image objects in a layout object typically have different spatial resolutions. Therefore, they must be scaled to the same resolution and to the page resolution before they are combined to create a page image according to Equations 1-3. In a JPM file, the size of the mask, image and page are specified in page grid units, but their resolutions may not be specified. The scaling of the mask and image is specified by a scale factor. The method of scaling is not specified in this International Standard, although informative guidelines for scaling of the mask and image are provided. The page image would then be scaled to the resolution of the output device for rendering on a display or printer. As described here, there would be two separate scaling operations: in the first, the mask and image of successive layout objects are scaled to a common resolution and combined on the page; in the second, the resulting page image is scaled to the device resolution. In practice, these two scaling operations are often combined into a single, device-dependent and implementation-specific scaling operation.




Figure 1 — Example compound document

5.2 File Elements and Structure

The files that conform to the format defined in this International Standard are called JPM files. At its core, a JPM file is a sequence of pages, where each page in turn is a sequence of layout objects. A layout object normally consists of a mask object and an image object. Mask and image objects are composited to build up the final page image according to equations 1-3. The key elements or boxes of a JPM file are: page collections, page, layout object, and object. An object points to its image data directly via a Contiguous Codestream box or indirectly via a Fragment Table box. Like all members of the JPEG 2000 file format family, a JPM file begins with a JPEG 2000 Signature Box and a File Type Box. A Compound Image Header Box, containing general information about the file, is then followed by Page Collection, Page, Fragment Table, Contiguous Codestream, Media Data and general metadata boxes. Refer to Annex B for constraints on the location of boxes in a JPM file.

This list illustrates the hierarchical relationship between the key elements in a JPM file. A particular order of these boxes is not implied. Full details of all the boxes may be found in Annex B.

- JPEG 2000 Signature
- File Type
- Compound Image Header
- Page Collection
- Page
 - Layout Object
 - Mask Object
 - Image Object
- Fragment Table
- Codestream Fragment
- Contiguous Codestream

5.2.1 Page Collections

A JPM compatible file consists of a sequence of pages, each represented by a Page box which occurs at the top level of the file and each of which can be rendered independently of any other page. Page collections are used to logically group pages in a JPM file. Page collections can be logically nested, so that a page collection can itself consist of one or more page collections and/or one or more pages. Page collections referred to from other page collections are called subsidiary page collections. All pages in a JPM file must be pointed to by at least one page collection.

A page can be said to be contained in a page collection, but this does not mean that the Page box is located within the Page Collection box. It is not. Page boxes and Page Collection boxes both occur at the top level of the file and are not contained within other boxes.

A JPM file contains one page collection known as the main page collection which is used to locate all pages of the file. Any additional page collections in a JPM file are logically nested within the main page collection (see 5.2.1.3). A Page Collection box contains an optional Label box, optional metadata (XML and/or UUID) boxes, and a Page Table box that contains the locations of the pages and page collections belonging to the page collection.

It is recommended that optimized files have the main page collection located near the beginning of the file. While Page boxes and Page Collection boxes occur at the top level of the file, they may be located in external files. This case may be viewed as equivalent to being at the top level of the file.

Each page in a JPM file has a primary page collection. The purpose of a primary page collection is to enable navigation in the primary document to which the page belongs using the sequential order within the primary page collection, thus providing support for previous page and next page commands.

Every Page Collection box contains a Page Table box. The Page Table box entries point to the locations of the Page and Page Collection boxes within the page collection. A flag for each entry specifies whether the location is that of a Page box or Page Collection box, as well as indicating whether the box contains a thumbnail or metadata.

By walking the tree of pages and logically nested page collections in a page's primary page collection, all pages in the page's primary document can be reached. Every page (with the exception of a self-contained JPM file containing only a single page) has a Primary Page Collection Locator or PPCLoc box. This box points to the primary page collection of the page and provides an index, Plx, into that page collection's page table where the page is referenced. Then the next page and previous page can be found by walking one page forward or backward from the current page.

NOTE - Multiple page collections can exist in a JPM file. Some may have functions other than basic navigation. A table of figures could point to those pages containing figures. A section table or chapter table might point to only the first pages of sections or chapters. Any page collections of this sort must be auxiliary page collections, since they provide redundant pointers to pages and page collections and are sparse rather than comprehensive (see 5.2.1.3).

NOTE - Figure 2 illustrates a logical grouping of page collections PC and pages P in a JPM file. PCa is the main page collection and the primary page collection for pages P0, P8 and P9 and page collections PCb, PCc and PCe. In a JPM file, the Page Table box of the Page Collection box for PCa would reference the Page boxes for P0, P8 and P9, and the Page Collection boxes for PCb, PCc and PCe. The Primary Page Collection Locator boxes in these Page and Page Collection boxes would reference the Page Collection box for PCa.

PCb is a subsidiary page collection of PCa and the primary page collection for pages P1, P6 and P7 and for page collection PCd. PCd is the primary page collection for pages P2, P3, P4 and P5. PCc is the primary page collection for pages P10 and P11.

PCe is an auxiliary page collection, which references page collection PCc and page P5.




Figure 2 — Example of page collections and pages

5.2.1.1 Main Page Collection

Each JPM file has one main page collection. The purpose of the file's main page collection is to comprehensively list all pages (and subsidiary page collections) within a JPM file. This allows random seeking to any of the pages in the file.

The Page boxes for these pages occur at the top level of the file format, as do Media Data boxes containing fragments of codestreams. Since these Media Data boxes may be large and since they would likely be interspersed with Page boxes, the Page boxes might be widely separated in the file. In a JPM file containing a client copy of several browsed pages of a server's JPM file, each successively viewed page and some portion of its codestream fragments would be appended in turn to the bottom of the file. An update to the main page collection allows each of these new pages to be located.

The main page collection comprehensively references all pages and page collections by means of a hierarchical arrangement. Some pages may be referenced from a page collection beneath the main page collection, but all pages and page collections are part of the tree structure of the main page collection.

A JPM file optimized for browsing would have the main page collection near the front of the file. On the other hand, a client copy created during a browsing session would likely have the main page collection appended to the end of the file each time it is modified. The old main page collection's Page Collection box could then be left in place but have its box type changed from "pcol" to "free". A later garbage collection step could delete these Free boxes. Because of cases like this, the file format has a pointer to the main page collection included in the Compound Image Header box near the top of the file. This makes it easy to locate the main Page Collection box.

5.2.1.2 Primary Page Collection

Each page or page collection has a primary page collection. By way of distinction, each JPM file has a main page collection. A primary page collection is a property of a page or page collection, not a property of a JPM file.

The primary page collection of a page is the page collection where a JPM reading application would find the "next page" and "previous page" for a current page.

A Primary Page Collection Locator or PPCLoc box must appear in all Page boxes and all Page Collection boxes, with the one exception detailed in the next paragraph. The PPCLoc box provides a pointer to the primary page collection of the page or page collection. This backward pointer enables a comprehensive tree walk to find all the pages and page collections in the file.

The PPCLoc box is optional only in the case of a single-page, self-contained JPM file (i.e. one which has no external references). In this case, it appears that the page has no primary page collection, but in fact the file's main page collection functions as the page's primary page collection.

The primary page collection of a page or page collection may not be in the local JPM file in which that Page box or Page Collection box is located. The local file may, for example, have three single pages, each of which has been copied from one of three different remote files. Keeping the original primary page collection pointers for these pages pointing to the original remote files allows a user to keep browsing the original source document via next page and previous page commands.

The main page collection for a file may be a copy of a subsidiary page collection on a remote server, for instance, in which case it would have the parent page collection on that remote server as its primary page collection. When the main page collection does not have a primary page collection, then the main Page Collection Box shall not contain a PPCLoc box.

5.2.1.3 Auxiliary Page Collection

Auxiliary page collections are page collections which provide redundant pointers to pages already pointed to in the logical tree structure of the main page collection. Examples of auxiliary page collections could include a List of Figures or a List of Tables. Auxiliary page collections still appear in the logical tree structure of the main page collection so that

they can be located, but a flag is used to mark them as auxiliary. This way, a decoding application knows they are not to be used to perform a comprehensive tree walk through all the pages referenced in the main page collection.

Auxiliary page collections appear in the main page collection to assist an application in locating them within the file, but they are not part of the logical tree that is walked to comprehensively locate all pages. They instead provide a redundant means of reaching selected pages and thus should be ignored when trying to determine the natural page order of the file. Auxiliary page collections can appear down in the hierarchy of the main page collection; they need not occur at the top level.

Auxiliary page collections should be labeled by means of a Label box in order to be useful to a receiving application. If they are labeled, then a decoding application could present the label to the end user and offer such options as "next page in List of Figures" and "previous page in List of Figures."

As an example, if page 17 appears in the "List of Figures" page collection, the decoding application would return to that page collection to get the next or previous page containing a figure, but would return to the primary page collection for page 17 (by means of the PPCLoc box in page 17's Page box) to find the next page or previous page.

5.2.2 Pages

A page in a JPM file is represented by a Page box, a superbox that consists of a Page Header box, containing general information about the page, a Page Collection Locator box, containing the location of the page's primary page collection, an optional Base Colour box, which describes the base page colour, optional Metadata boxes, and Layout Object boxes, one for each layout object on the page. Full details of these boxes may be found in Annex B.2.

Pages occur at the top level of the file format. This means incremental updates to the file may be accomplished by simply appending new pages to the end of the file. Page boxes and boxes containing codestreams or portions of codestreams may be intermixed in the file.

With the exception of document thumbnails, each image in a JPM file is logically associated with at least one page.

5.2.2.1 Base Page

The BasePage is the initial PageImage before any layout objects have been rendered. Let page_width and page_height be the width and height of the page respectively, as signalled in the Page Header box.

Let spc be the colourspace in which the I images are to be combined to generate a PageImage. BasePage is an image with dimensions page_width and page_height and colourspace spc.

If the PColour field of Page Header box is 0 then the BasePage is transparent and contains the sample values of any underlying image, converted to the colourspace spc.

If the PColour field of the Page Header box is 1 then every BasePage sample contains the representation of white in the colourspace spc.

If the PColour field of the Page Header box is 2 then every BasePage sample contains the representation of black in the colourspace spc.

If the PColour field of the Page Header box is 255 then there is a mandatory Base Colour box within the Page box and every BasePage sample contains the spc representation of the colour indicated by the Base Colour box.

5.2.3 Layout Objects

Within a Page Box, there are as many Layout Object boxes as there are layout objects on the page. A Layout Object box is a superbox that consists of a Layout Object Header box, containing general information about the layout object, optional boxes containing metadata associated with the layout object, and one or two Object boxes - either an image Object box and/or mask Object box, or a combined image/mask Object box.

Each Object box contains an Object Header box identifying whether the object represents an image, a mask or a combined image/mask, specifying the location of any codestream associated with the object and containing positioning information for the object.

An image object typically has a codestream associated with it, but may not, in which case the NoCodestream field in the image Object header is set to 1. An image object may also have an associated constant colour or image base colour. If an image object does not have an associated codestream, then it must have a defined image base colour. A mask object or image/mask object, if present, must have a codestream associated with it and have NoCodestream=0.

If an object has an associated codestream then the Object Header box is followed by an optional Object Scale box and a JP2 Header box containing boxes describing the object data: an Image Header, Colour Specification, optional Bits Per Component, Palette, Component Mapping and Channel Definition boxes.

Associated with each layout object is a single component opacity mask M and an image I, each with the same width and height as the containing page.

The I images for the layout objects to be combined to generate a PageImage must all use the same colourspace and bit-depth. The colourspace to be used for the I images of layout objects may be decided by the implementor and is not defined by this International Standard.

The general methods for generating the mask M and the image I associated with a layout object are described in 5.2.3.2 and 5.2.3.4.

5.2.3.1 describes the special case of generating a mask M and an image I from a JBIG 2 codestream with an associated Recommendation ITU-T 45 encoding of colour tags.

5.2.3.1 Colour Tagged JBIG 2 Layout Objects

If the Layout Object box contains an image Object box with a compression type of Recommendation ITU-T T.45 (Run Length) coding then it must also contain a mask Object box with a compression type of Recommendation ITU-T T.88 (JBIG 2), and both must have the NoCodestream field in the Object Header boxes set to 0. In addition, the following fields must be the same for the image and mask Objects:

- the OVoff and OHoff fields in the Object Header boxes;
- the VRN, VRD, HRN, VRD fields in the Scale boxes;
- the HEIGHT and WIDTH fields in the Image Header box in the JP2 Header boxes.

In 5.2.3.3.1, *m_orig* is the image obtained by decompressing the JBIG 2 codestream associated with the mask Object box.

In 5.2.3.4.1, *i_orig* is the image obtained by applying the colour tags associated with the image Object box to the symbol occurrences in the JBIG 2 mask codestream as described in ITU-T T.45.

5.2.3.2 Mask M for a Layout Object

If the Layout Object box does not contain a mask Object box or a combined image/mask Object box, then the mask M for the layout object is defined to have a bit-depth of 1 and to have value 1 at all sample locations, i.e. $M[0][x,y] = 1$.

Subclauses 5.2.3.3.1 - 5.2.3.3.5 define M when a codestream is associated with a mask Object box or a combined image/mask Object box in the Layout Object box.

1. The first stage in generating M, described in 5.2.3.3.1, is to decode the mask object, converting to a bit-depth of 8 if the mask has a lower bit-depth, and ensuring that the samples use a "min is white" representation.
2. The second stage, described in 5.2.3.3.2, is to scale the mask.
3. The third stage, described in 5.2.3.3.3, is to clip the mask from the top and from the left.
4. The fourth stage, described in 5.2.3.3.4, is to position the mask on the page.
5. The fifth and final stage, described in 5.2.3.3.5, is to clip the mask to the layout window.

These stages are described individually for clarity and an efficient JPM decoder may be able to combine one or more of these stages.

Figure 3 illustrates the five stages for an example mask. Note that white in the images denotes a value of 0.




Figure 3 — Example of generating M

5.2.3.2.1 Decoding the Mask

If a mask Object box is used then let m_orig be the image obtained by decompressing the associated codestream, otherwise let m_orig be the image containing only the opacity component of the image obtained by decompressing the associated codestream. Let m_orig_width and m_orig_height be the width and height of m_orig and let b be the bit-depth of m_orig .

If b is 1 then, unless indicated otherwise by the codestream, samples of value 0 are assumed to represent white while samples of value 1 are assumed to represent black.

If b is greater than 1 then, unless indicated otherwise by the codestream, samples of value 0 are assumed to represent black while samples of value $2^b - 1$ are assumed to represent white.

Let $grey$ be an image with bit-depth $g = \max(8, b)$ and the same width and height as m_orig . $grey$, defined as follows:

$$grey[0][x, y] = \begin{cases} 2^b - 1 - m_orig[0][x, y] \times \frac{2^g - 1}{2^b - 1} & \text{if 0 represents black in } m_orig \\ m_orig[0][x, y] \times \frac{2^g - 1}{2^b - 1} & \text{otherwise} \end{cases} \quad (4)$$

5.2.3.2.2 Scaling the Mask

If an Object Scale box is contained in the mask Object box or combined image/mask Object box then let m_VRN , m_VRD , m_HRN and m_HRD be the scaling factors contained in this box, otherwise let m_VRN , m_VRD , m_HRN and m_HRD all be 1. Let m_scaled be the result of scaling $grey$ vertically by the ratio $\frac{m_VRN}{m_VRD}$ and horizontally by the

ratio $\frac{m_HRN}{m_HRD}$, generating an image with bit-depth g and dimensions:

$$m_scaled_width = \left\lfloor \frac{m_HRN}{m_HRD} \times m_orig_width \right\rfloor \quad \text{and} \quad m_scaled_height = \left\lfloor \frac{m_VRN}{m_VRD} \times m_orig_height \right\rfloor \quad (5)$$

NOTE - The method for scaling a binary or grey mask up or down is implementation specific. For example, the scaling may use nearest neighbour or bilinear interpolation.

5.2.3.2.3 Clipping the Mask

Let m_OVoff and m_OHoff be the vertical and horizontal offsets in the mask Object box or combined image/ mask Object box.

Let $m_clipped$ be an image with bit-depth g , width $m_width = \max(0, m_scaled_width - m_OHoff)$ and $m_height = \max(0, m_scaled_height - m_OVoff)$. $m_clipped$ is defined as follows:

$$m_clipped[0][x,y] = m_scaled[0][x + m_OHoff, y + m_OVoff] \quad (6)$$

5.2.3.2.4 Positioning the Mask

Let $page_width$ and $page_height$ be the width of the page containing the layout object, as signalled in the Page Header box. Let $LVoff$ and $LHoff$ be the layout vertical and horizontal offsets in the Layout Object Header box.

Let m_pos be an image with a width of $page_width$ and height of $page_height$, defined as follows:

$$m_pos[0][x,y] = \begin{cases} m_clipped[0][x - LHoff, y - LVoff] & 0 \leq x - LHoff < m_width \text{ and } 0 \leq y - LVoff < m_height \\ 2^g - 1 & \text{otherwise} \end{cases} \quad (7)$$

5.2.3.2.5 Window Clipping the Mask

Let M be an image with a width of $page_width$ and height of $page_height$. Let $LWidth$ and $LHeight$ be the layout object width and layout object height as specified in the Layout Object Header box. M is defined as follows:

$$M[0][x,y] = \begin{cases} m_pos[0][x,y] & 0 \leq x - LHoff < m_width \text{ and } 0 \leq y - LVoff < m_height \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

5.2.3.3 Image I for a Layout Object

Let spc be the colourspace in which the I images are to be combined to generate a $PageImage$.

If the Layout Object box contains an image in either an image Object box or a combined image/mask Object box and also contains a Base Colour box, then let $base$ be the representation in colour space spc of the colour indicated by the Base Colour box; otherwise let $base$ be the representation in colour space spc of the colour black.

If the Layout Object box does not contain an image Object box or a combined image/mask Object box, then the image I for the layout object is defined to be an image in colour space spc where every sample has the colour $base$.

If the Layout Object does contain an image Object box or a combined image/mask Object box, but with no associated codestream, then the image I for the layout object is also defined to be an image in colour space spc where every sample has the colour $base$.

Subclauses 5.2.3.4.1 - 5.2.3.4.5 define I when a codestream is associated with a image Object box or a combined image/mask Object box in the Layout Object box.

1. The first stage in generating I , described in 5.2.3.4.1, is to decode the image object and convert the image to the spc colour space.
2. The second stage, described in 5.2.3.4.2, is to scale the image.
3. The third stage, described in 5.2.3.4.3, is to clip the image from the top and from the left.
4. The fourth stage, described in 5.2.3.4.4, is to position the image on the page.
5. The fifth and final stage, described in 5.2.3.4.5, is to clip the image to the layout window.

These stages, as with those to generate the mask M , are described individually for clarity and an efficient JPM decoder may be able to combine one or more of these stages.

5.2.3.3.1 Decoding the Image

If an image Object box is used then let i_orig be the image obtained by decompressing the associated codestream, otherwise let i_orig be the image containing all but the last component of the image obtained by decompressing the associated codestream. Let i_orig_width and i_orig_height be the width and height of i_orig .

Let $conv$ be the image obtained by converting $orig$ to the colour space spc .

5.2.3.3.2 Scaling the Image

If an Object Scale box is contained in the image Object box or combined image/mask Object box then let i_VRN , i_VRD , i_HRN and i_HRD be the scaling factors contained in this box, otherwise let i_VRN , i_VRD , i_HRN and i_HRD all be 1. Let i_scaled be the result of scaling $conv$ vertically by the ratio $\frac{i_VRN}{i_VRD}$ and horizontally by the ratio

$\frac{i_HRN}{i_HRD}$, generating an image with dimensions:

$$i_scaled_width = \left\lfloor \frac{i_HRN}{i_HRD} \times i_orig_width \right\rfloor \quad \text{and} \quad i_scaled_height = \left\lfloor \frac{i_VRN}{i_VRD} \times i_orig_height \right\rfloor \quad (9)$$

NOTE - An image may be scaled up using bilinear interpolation or, if it's a palette image, using nearest neighbour interpolation. An image may be scaled down using averaging followed by subsampling, or if it's a palette image, using subsampling. If the image is compressed using JPEG 2000, an implementation may also choose to use the progressive-byresolution decompression feature to obtain a power-of-two reduced resolution image for scaling.

5.2.3.3.3 Clipping the Image

Let i_OVoff and i_OHoff be the vertical and horizontal offsets in the image Object box or combined image/mask Object box.

Let $i_clipped$ be an image with bit-depth g , width $i_width = \max(0, i_scaled_width - i_OHoff)$ and $i_height = \max(0, i_scaled_height - i_OVoff)$. $i_clipped$ is defined as follows:

$$i_clipped[0][x,y] = i_scaled[0][x + i_OHoff, y + i_OVoff] \quad (10)$$

5.2.3.3.4 Positioning the Image

Let $page_width$ and $page_height$ be the width of the page containing the layout object, as signalled in the Page Header box. Let $LVoff$ and $LHoff$ be the layout vertical and horizontal offsets in the Layout Object Header box.

Let i_pos be an image with a width of $page_width$ and height of $page_height$, defined as follows:

$$i_pos[0][x,y] = \begin{cases} i_clipped[0][x - LHoff, y - LVoff] & 0 \leq x - LHoff < i_width \text{ and } 0 \leq y - LVoff < i_height \\ base[c] & \text{otherwise} \end{cases} \quad (11)$$

5.2.3.3.5 Window Clipping the Image

Let I be an image with a width of $page_width$ and height of $page_height$. Let $LWidth$ and $LHeight$ be the layout object width and layout object height as specified in the Layout Object Header box. I is defined as follows:

$$I[0][x,y] = \begin{cases} i_pos[0][x,y] & 0 \leq x - LHoff < i_width \text{ and } 0 \leq y - LVoff < i_height \\ 0 & \text{otherwise} \end{cases} \quad (12)$$

5.2.4 Thumbnails

Thumbnails are an optional feature which allow an overview of a document or a page to be presented to an end user. There are two types of thumbnails: document thumbnails and page thumbnails.

A document thumbnail offers an overview representation of an entire document, much as the cover or spine of a book serves to do. It may simply be a repetition of the page thumbnail for the title page. Document thumbnails may be JP2 compatible or not.

A JPM file with a document thumbnail is created by adding a JP2 Header box for the thumbnail image immediately after the File Type box. The codestream for the thumbnail image must be contained in the first Contiguous Codestream box following the JP2 Header box.

NOTE - A JP2-compatible document thumbnail can be used in the construction of a JPM file that also conforms to the JP2 file format defined in Recommendation ITU-T T.800 | ISO/IEC 15444-1. A JP2 compatible document thumbnail must use JPEG 2000 as the coder and appears to a JP2 reader to be a valid JP2 file. The JP2 compatibility flag is set in the File Type box by adding the string "jp2\040" to the compatibility string. A JP2 reader shall then look for a JP2 Header box and jump over any following boxes until it finds a Contiguous Codestream box. It would then access the image represented by the data in that box. When a document thumbnail is present, the file would contain a JP2 Header Box, and a Contiguous Codestream Box near the beginning of the file, somewhere after the Compound Image Header Box and before the first Page Box, usually before the Page Collection Box.

A page thumbnail is simply the layout object in a page with layout object identifier 0. A page thumbnail layout object must be the first layout object in a page. Its presence is signalled by a flag in the page table which indicates that a given page pointed to from the page table contains a page thumbnail. Page thumbnails are optional.

NOTE - Page thumbnails allow an overview of each individual page to be displayed. This can be used during navigation, for example to render a small icon near each node of a tree-like table of contents. It can also be used to allow the display of a more useful initial image onto a new page than simply the page's background colour while waiting for additional content to arrive.

5.2.5 Contiguous and Fragmented Codestreams

The codestream data for the objects of a page may either be a single contiguous codestream, or be composed of one or more codestream fragments. A codestream consisting of fragments is accessed via a Fragment Table box. A Fragment List box lists the locations in the file, or external pointers to another file, for the codestream fragments. Each codestream fragment is contained within a Media Data box. If the codestream is not fragmented then the object shall contain the location of a Contiguous Codestream box containing the codestream data. Full details of these boxes may be found in Annex B.5.

The fragment table mechanism permits any codestream to be broken up into an ordered list of fragments, each pointed to by its own offset and data reference. A data reference indexes into the data reference table and produces a string listing the filename or URI of an external file where the fragment is to be found. If the data reference is 0, the fragment is located in the current file. This mechanism allows fragments to be placed in an optimal position to support streaming of the data to a client application from a server, or to allow progressive refinement of multiple codestreams simultaneously.

Fragments may start or stop at any byte boundary of the codestream. The length of each fragment is specified by the length parameter retrieved from the fragment table. If a fragment does not end at a point in the codestream where a full code block or full raster line can be decoded, the decoding application holds this additional information until the next fragment has been retrieved. The next fragment of the byte stream is then appended to the residual data from the prior fragment and then decoding is resumed. Codestreams having any compression type may be decoded in this way.

Fragment table entries always point to codestreams. All header information required to decode the codestream, such as width and height, compression type or bit depth is stored in fields in the Object box.

5.2.6 Shared Data

A file can contain multiple boxes or fragments containing identical data. File size can be reduced by sharing these boxes, including by reference a box or fragment that occurs in one part of a JPM file in a different part of the same file or in a different JPM file. The two mechanisms for sharing boxes in a JPM file are Shared-Data Reference box and Cross-Reference box. Wherever a Shared-Data Reference or a Cross-Reference box occurs in the file, the data it references is treated as if it exists at the point in the file where the reference is located. The reference can be by means of either an identifier, in the case of a Shared-Data Reference box (Annex B.1.8), or a data reference including an offset and length, in the case of a Cross-Reference box (Annex B.6.4).

A Shared-Data Reference box in a file uses an identifier that is defined in the same file; therefore this mechanism shall only be used to share data within a file. The Cross-Reference box uses a data reference and therefore may be used to share data across files. In the case of a reference to an offset from the beginning of the file and length, the offset applies before any substitution implied by the reference occurs. A shared data substitution would not be understood by a JP2-only reader, so that a JP2-conforming JPM file must include directly any box whose presence in the file is necessary for JP2 conformance.

5.2.7 Metadata

This International Standard allows metadata to be attached to any structural element of a JPM file, including the file or document itself, page collections, pages, layout objects and objects. This is indicated in the box definitions in Annex B by the inclusion of a "metadata" box in the data field of the container box. Whenever there is a reference to a metadata in a box definition, it is understood that the box may contain an XML box or a UUID box with metadata. These "metadata" boxes are defined in Annex C.

5.3 Hidden Text Metadata

Hidden text metadata is data representing the text, text elements and text flow associated with an image. In the context of this standard, hidden text is associated with a particular region of a page in a JPM document.

Common uses for hidden text include text searching and highlighting, cut-and-paste, and text-to-speech processing. Hidden text describes the flow of the text on a page as well as the text elements.

JPM allows a rich, multiple content-type representation of a document. Each region of a page may be encoded with a compression technique best suited to its characteristics. In regions containing text, high fidelity reproduction of the source image is retained by not replacing the text regions with a character-based rendition through OCR, but rather by using advanced coding methods such as JBIG2. Even OCR results with a 99 percent accuracy contain substantial numbers of errors per page which require expensive human labour to correct. The searchable nature of a character-based rendition can be obtained instead by associating hidden "dirty OCR" results with the corresponding text image. This standard defines a format for hidden text metadata.

A key issue with hidden text is capturing the ambiguities seen by the OCR engine in a way that allows properly-constructed search engines to find whether and where a given word might be present in a text image. Properly captured, this information provides nearly as much searching precision as an approach using human-corrected "clean OCR" data, but at much lower cost. Search results are most useful where there are fewer false positives to weed through. Intelligent

search engines can take account of such data as confidence and alternate characters or alternate words to appropriately alter the ranking of search hits on less certain characters.

In many cases, true ambiguity exists in the image and it would confuse a human observer as well. In these cases, saving confidence values for characters and their alternatives or describing several alternative parsings of a string of characters into words can amount to saving the state of the OCR process to allow the problem to be revisited in a later stage, perhaps by a different engine or by access to first a general dictionary and then a set of more specialized dictionaries.

As a last step, when a person is presented with the search results, they can dismiss a given search hit by comparison to the actual image data for a character or word. For this purpose (and to allow later-stage OCR processes to resume analysis on the image), bounding box rectangles can be defined for all the elements of the hidden text such as characters, words, lines, paragraphs and regions. By indicating a container relationship among these items, intelligent navigation and text selection can occur at character, word, line, paragraph boundaries. A reading order through these rectangles can be defined for what was in the image just a random placement of unrelated glyphs.

While it is primarily designed for use by machines such as search engines, the hidden text can also serve as a crude (if "dirty") or adequate (if "clean") alternate representation for an image region to allow it to display on character-based devices (such as mobile phones) or small-area graphics devices (such as PDAs).

Annotations are added to the document typically with a WYSIWYG editor to indicate URL references, notes, and to highlight key sections of the document text. Each annotation is associated with a particular region of a page in a JPM document.

XML is used for hidden text and annotations because it is a format widely used to store structured information, and can be machine processed.

5.4 JPM Use Scenarios

5.4.1 Self-Contained Files and Files with External References

This Subclause is informative section. Subclause 5.2 outlines the structure of a JPM file. In a local file, typically all image content and parameter boxes are contained in a single file for ease of tracking and maintenance, and because local storage is available and local bandwidths are high.

The multiple layout object model is a method of representing and constructing document images, which allows access paradigms based on the multiple objects in addition to providing compression advantages due to applying the optimal compression method to each object. In addition, document images have an internal structure that implies a reading order, unlike photographic images, so that breaking them into layout objects can bring streaming advantages where only the small area of the page (a column, say, or a paragraph) that is of interest is brought over and refined.

The fragment table and cross reference features of the JPEG 2000 family of file formats provide enormous flexibility in the construction of files, whose contents are distributed across multiple files, some of which may be located remotely on networks such as the Internet. This subclause examines some of the file organisations enabled by these features.

Useful JPM files tend to have structural information boxes at the top and large codestream objects either at the bottom of the file or remotely located or both. This allows rapid parsing and seeking among the component parts of the file by first getting the necessary structural information and then seeking to the offsets specified in the local or remote files.

Structural information includes such information as Page Table boxes, Page Collection boxes, Label boxes and Fragment Table boxes.

Page Table boxes are a key structural feature. They refer to Page boxes via page references. These references contain an offset, length and data reference for the Page Box associated with each page in the complete file. Each of these data references may be zero, indicating a page reference to a Page box in the current file, or non-zero, indicating that the data reference value is to be used as an index into the Data Reference box. The Data Reference box contains an enumerated list of strings, each string being a file name or a URI pointing to an Internet file.

5.4.2 Files prepared for the Internet

JPM files prepared for use on the Internet will typically already have been decomposed into layers and perhaps also into separate layout objects. These preparation steps each give a measure of improved streaming performance in browsing remote JPM files.

Such files would likely have structural information placed near the top of the file, with codestreams occurring later. Codestreams and subsets of codestreams are accessible directly by seeking to a point in the file indicated by entries in the fragment table associated with that codestream.

5.4.3 Page Collection Example

The main page collection of a very large multi-volume set of books such as an encyclopedia should not be printed page individually; such a data structure would be too large to be brought over to a client before even the first page can be

displayed. Instead, the main page collection might serve as what amounts to a table of contents and point to each of the front matter pages individually, then to page collections that each cover entire volumes: Volume A, Volume B and so on. Each of these "volume" page collections would reference "section" page collections that would reference "article" page collections, and so on. Each of these page collections could have an appropriate label applied by means of Label boxes in their Page Collection boxes.

A search of metadata for a term occurring on a few of the pages returns a new "search results" JPM file containing a single main "search results" page collection pointing to a handful of pages. The first entry in this page collection box would point directly to the first search result page. This would actually point to one of the pages in the encyclopedia over on the server via an external reference. By navigating the "search results" page collection, the application can let the user go to next search hit and previous search hit. If the user then finds the article of interest and wants to continue reading onto a page back in the remote encyclopedia where no search hit occurred, they would then need a next page function as well.

The client software would go to the PPCLoc box found in that page's Page box. This entry has a pointer to the primary Page Collection box back in the main JPM file. This primary Page Collection box will then be used by the client software to first find the current page's entry in that page collection by means of the PIx portion of the PPCLoc box. The PIx + 1 entry in that page collection would be the next page. By that means, the previous and next pages can be found.

5.4.4 Page Collections in a Client / Server Context

The JPM file format places Page boxes and Page Collection boxes at the file level (not within other superboxes) to assist with incremental browsing of pages. New pages or page collections drawn over from a server may simply be appended to the end of the local (incomplete) copy of the file.

In a client/server environment such as the web, it would be common for a user of a client application to browse only certain pages of a multi-page JPM file on a server. An initial link could point the client application to a specific page in the server's JPM file. The client application would create a new local JPM file and could choose to continuously update it as portions of the server file are retrieved. This would ensure that the browser cache always contains a valid JPM file when browsing ceases.

This local file would have a main page collection. The main page collection would initially contain just one page. As the user begins fetching additional pages, these would be appended to the local file and the main page collection's page table updated to point to each new local copy of a server page. At the end of browsing, a stand-alone file having no clear connection to the server file could exist on the client.

A better way of performing this browsing step would be to copy the server's main page collection to the client prior to downloading any pages. This page collection would be fixed up to point to the server copies of the pages by creating a data reference to the server URL. Then, as the first page is brought down to the client, the main page collection is updated to point to the local page, while continuing to point to the server for all other pages.

5.4.5 Example of a Client/Server Interaction

In a client/server system, a user might want to browse a multi-page JPM file on a server, and to begin on a page in the middle of that remote file. Perhaps a search engine found a hit on the JPM file metadata associated with that page. The URI returned by the search engine would point to the JPM file and could contain a construction such as <http://webimaging.org/test.jpm?page=page17>. This would instruct the JPM decoder subsystem of the browser to abort the download of the full JPM file as soon as the key structural boxes at the top of the box have completed downloading. Then, byte serving protocols would be used to retrieve only the ranges of bytes needed to render the data in the Page box associated through the Label box to label "page17." Each Object box in each Layout Object box of a given Page box may contain a fragment table reference to a Fragment Table Box. By retrieving only enough of each layout object's Fragments to render a rough initial version of the page, a full rendering of all the layout objects can be deferred. The end user could in parallel begin moving the mouse over the browser's display window to indicate which layout objects might better be refined next, based on their having been indicated of interest.

During this interactive browsing session, fragments are streaming over to the client copy of the JPM file in an order that is quite different than the order in the server copy. Page 17 comes first in the client copy, where page 1 was more likely first in the server copy. Similarly, the layout objects within page 17 will have a different order in the client version of the page, since the user's mouse movements dictated their retrieval. The client version can nevertheless form a perfectly equivalent surrogate for the server version because the page tables and fragment tables in the client copy will have started out pointing entirely to fragments on the server. As fragments and Page boxes stream partially over to the client, the appropriate entries in the Page Table boxes and Fragment Table boxes are updated to point to the appropriate segments of the client copy of the file whether it is just in the file cache or explicitly saved.

When substantial changes are made to the fragment table, then the fragment table pointer can be changed to point to a new, post-pended fragment table at the end of the file, and the old fragment table can be turned into a Free box.

All the fragments of the data in the file could be retrieved either from another file located on the web server external to the server copy of the JPM file or from a file located on another web server, or from an external file located on the client

machine. The fragment reference mechanism supports all of these sources of data fragments external to the main JPM file.

Similarly, the cross reference mechanism provides a means for JPM boxes (as opposed to data fragments) to be located in any of those possible places as well. The Cross Reference box has an offset, length and data reference that points to an internal or external location so that a box can be found. For cross references internal to the main file, Shared Data Reference boxes permit repeated boxes to appear only once, with the second and subsequent instances being included by reference to the first instance. The key difference between fragment tables and cross references is that cross references point to the beginning of the referenced box, whereas fragment table entries point directly to the first byte of the codestream fragment, beyond the beginning of any containing Media Data box or other box.

5.4.6 Example of Scanner Capture

In addition to supporting a variety of streaming and remote browsing behaviors, the file format can support datastreams coming from high-speed document scanners, where the images are not yet optimised for viewing but must meet other stringent constraints, such as timing. These files are not heavily processed at scan time, but rather capture all the available image data (perhaps multiple images per page) coming from the scanner and capture software and save this data to make it available to downstream post-processing software that can convert it to a web-optimised form.

Scanned files may have metadata specific to the scanning process. This could include scanner metrics such as histograms, skew angle measurements, or cropping rectangle specifications. Other metadata might describe the scanner model or date and time of scanning and so on.

Scanned files have not been prepared for efficient browsing over a network. They are intended to involve minimal encoder processing effort in order to meet the speed constraints imposed by high-speed scanning.

JPM files coming from a scanning subsystem would typically not have been decomposed into layers or layout objects, but may instead include an entire high-quality compressed colour or greyscale image suitable for use by a subsequent decomposition system. If the scanning system also is capable of producing a bitonal (depth 1 bit) image natively, it may wish to include both this image and the full colour image in the JPM file to assist downstream decomposition processing. Two different scanned images of the same content which have not been pre-processed according to the MRC model might be placed into a JPM file as follows: place the colour image as a layout object on one page and the bitonal image as a layout object on another page and associate the two pages by placing them together in a page collection. Appropriate Label Boxes could clarify the relationship between the two images and indicate that they represent the same page.

Annex A

Compound Image File Structure

(This annex forms an integral part of this Recommendation | International Standard)

This Annex describes the structure and organization of a JPM file. A JPM file uses the file format architecture specified in Recommendation ITU-T T.800 | ISO/IEC 15444-1. Therefore, a JPM file is a contiguous sequence of boxes. This Annex defines a box and describes the notation used for the box definitions in this International Standard. This Annex also lists the boxes that are used in a JPM file.

A.1 File Identification

The brand shall be 'jpm\040' for files that are completely defined by this Part of this International Standard; brand is defined in Recommendation ITU-T T.800 | ISO/IEC 15444-1 Annex I. JPM files are files that conform to this Part of this International Standard. JPM files can be identified using several mechanisms. When stored in traditional computer file systems, files that conform to this Part should have the file extension.jpm (readers should allow mixed case). On Macintosh file systems, the type code should be 'jpm\040'. The MIME type is image/jpm.

A.2 File Organization

A JPM file uses the file format architecture specified in Recommendation ITU-T T.800 | ISO/IEC 15444-1. The binary structure of the file is a contiguous sequence of boxes. The start of the first box shall be the first byte of the file and the end of the last box shall be the last byte of the file. The binary structure of a box in a JPM file is identical to that defined in the JP2 file format (Recommendation ITU-T T.800 | ISO/IEC 15444-1 Annex I.4).

This International Standard defines boxes mandatory to a JPM file, as well as several optional boxes. Other Recommendations | International Standards may define other boxes that may also be found within a JPM file. In all cases, the information contained within a JPM file shall be in the box format; byte-streams not in the box format shall not be found in a JPM file.

A JPM file may be self-contained in that it contains all the data needed to composite the page or pages in the file. A JPM file may also reference images and data in external files. References to the content of an external file are by means of the Data Reference box (Annex B.1.5) or Cross Reference boxes (Annex B.6.4).

Schematically, the hierarchical organization of boxes in a JPM file is shown in Figure A-1. Boxes with dashed borders are optional in conforming JPM files. However, an optional box may define mandatory boxes within that optional box. In that case, if the optional box exists, those mandatory boxes within the optional box shall exist. This illustration specifies only the containment relationship between the boxes in the file. A particular order of those boxes in the file is not generally implied. The definition of a box will include whether or not that box is required to be found at a specific location within the file or another box.




Figure A-1 — Conceptual structure of a JPM file

Copyright International Telecommunication Union (ITU) 2012. All rights reserved.

The JPEG 2000 Signature Box identifies the file as being part of the JPEG 2000 family of file formats.

The File Type box specifies file type, version and compatibility, including specifying if this file is a conforming JPM file or if it may be read, at least in part, by a conforming JP2 or JPX reader.

An optional JP2 Header box may follow the File Type box, describing a document thumbnail image for the JPM file.

The Compound Image Header box contains global information that describes the Compound Image file.

The Page Collection box contains a Page Table box for locating the individual pages of a JPM file.

The Page box describes the page onto which its constituent layout objects are composited using the imaging model. It may also contain references to metadata associated with the page. The Page box also contains one or more layout objects.

A Compound Image File must contain a JPEG 2000 Signature box, a File Type box and a Compound Image Header box. The file may or may not contain image data. Figure A-1 shows the data stored with the Compound Image File.

The JPM format allows the codestreams in the image layers to be non-contiguous. This enables interleaving codestream fragments among images and inter-image progressive rendering, where multiple images are rendered progressively at the same time.

The JPM format also allows the codestreams in the image layers to be stored outside but referenced from the Compound Image File. This means that the images used in a compound image can be stored separately from their parent compound image or JPM file.

A.3 Box Definition

Physically, each element in a JPM file is encapsulated within a binary structure called a box. That binary structure is defined in Recommendation ITU-T T.800 | ISO/IEC 15444-1 Annex I.4.

Annex B defines the boxes used in this Part of this International Standard. Each box is described using the notation defined in Recommendation ITU-T T.800 | ISO/IEC 15444-1 Annex I.3.6.

The following additional information is provided in the box definitions.

Box type:
Container:
Mandatory:
Quantity:
Location:

Box type is the value of the TBox field of the box. Mandatory indicates whether the box is required to be present in the box that contains it. Container specifies the box that contains the box being defined. Quantity is the number of boxes being defined that can occur in the Container. Location describes where in the Container the box being defined can occur.

- Where not specified, all integer values are assumed to use the big-endian byte order.
- Some boxes have values assigned to groups of bits within a field. In some cases there are bits, denoted by "x", that are not assigned a value for any field within a parameter. The fields shall contain a value of zero for all such bits. The decoder shall ignore these bits.

A.4 Boxes Used in a Compound Image File

Table A-1 lists the boxes used in a JPM file and defined or referenced within this International Standard.

Table A-1 — Boxes defined or referenced within this International Standard

Box names	Types	Superbox	Comments (Informative)
Base Colour	'bclr' (0x6263 6C72)	Yes	This box contains all the information specifying the colour of a page or of an object for which
Base Colour Value	'bcvl' (0x6263 766C)	No	This box specifies the component values for the colour of a page or of an object for which no image data is available.
Colour Specification	'colr' (0x636F 6C72)	No	This box specifies the colourspace of an image.
Compound Image Header	'mhdr' (0x6D68 6472)	No	This box contains general information about the compound image file, such as profile and version.
Contiguous Codestream	'jp2c' (0x6A70 3263)	No	This box contains a valid and complete JPEG 2000 codestream.
Cross-Reference	'cref' (0x6372 6566)	No	This box specifies that a box found in another location (either within the JPM file or within another file) should be considered as if it was directly contained at this location in the JPM file.
Data Reference	'dtbl' (0x6474 626C)	No	This box contains a set of pointers to other files or data streams not contained within the JPM file itself.
File Type	'ftyp' (0x6674 7970)	No	This box specifies file type, version and compatibility information, including specifying if this file is a conforming JPM file.
Fragment List	'flst' (0x666C 7374)	No	This box specifies a list of fragments that make up one particular codestream within the file.
Fragment Table	'ftbl' (0x6674 626C)	Yes	This box describes how a codestream has been split up or fragmented and then stored within the file.
Free	'free' (0x6672 6565)	No	This box contains data that is no longer used and may be overwritten when the file is updated.
Hidden Text Metadata	'htxb' (0x68747862)	Yes	This optional box contains hidden text and annotations.
HTX Reference Box	'phtx' (0x70687478)	No	This optional box can be used to point to Hidden Text Metadata box contents at top file level.
Image Header	'ihdr' (0x6968 6472)	No	This box specifies the size of the image and other related fields.
JP2 Header	'jp2h' (0x6A70 3268)	Yes	This box contains a series of boxes that contain header-type information about a file containing a single image.
JPEG 2000 Signature	'JP\040\040' (0x6A50 2020)	No	This box uniquely defines the file as being part of the JPEG 2000 family of files.
Label	'lbl\040' (0x6C62 6C20)	No	This box specifies a textual label for a Page box or a Page Collection box.
Layout Object	'lobj' (0x6C6F 626A)	Yes	This box contains the information and image data needed to composite a mask-image object pair.
Layout Object Header	'lhdr' (0x6C68 6472)	No	This box describes the properties of the layout object and assigns each a unique identifier within the page.

Table A-1 — Boxes defined or referenced within this International Standard (continued)

Box names	Types	Superbox	Comments (Informative)
Media Data	'mdat' (0x6D64 6174)	No	This box contains generic media data, which is referenced through the Fragment List box.
Object	'objc' (0x6F62 6A63)	Yes	This box contains all the image data and information for one object in a layout object.
Object Header	'ohdr' (0x6F68 6472)	No	This box describes the properties of an object, identifying it as a mask, an image or a combined mask/image object, and assigns each object a unique identifier within the file.
Object Scale	'scal' (0x7363 616C)	No	This box describes the scaling for an object before applying it to the page.
Page	'page' (0x7061 6765)	Yes	This box contains all the information needed to image a page including references to codestream data.
Page Collection	'pcol' (0x7063 6F6C)	No	This box groups together the locations of a set of pages so that they are treated as related and associated with each other.
Page Header	'phdr' (0x7068 6472)	No	This box describes the properties of a page and gives the number of layout objects in the page.
Page Table	'pagt' (0x7061 6774)	No	This box gives the locations of the pages in a page collection and enables random access of pages in a file.
Primary Page Collection Locator	'ppcl' (0x7070 636C)	No	This box gives the location of the primary page collection for the page collection.
Shared Data Entry	'sdat' (0x7364 6174)	Yes	This box contains a box that can be referenced by an identifier from multiple places within a file.
Shared Data Reference	'sref' (0x7372 6566)	No	This box can be used to insert a box in the file by reference to a previous occurrence of the box in the same file.
UUID	'uuid' (0x7575 6964)	No	This box contains vendor specific information.
UUID Info	'uinf' (0x7569 6E66)	Yes	This box contains additional information associated with a UUID.
XML	'xml\040' (0x786D 6C20)	No	This box contains vendor specific information in XML format.

Annex B

Box Definitions

(This annex forms an integral part of this Recommendation | International Standard)

This Annex defines the boxes that are used in a conforming JPM file and that shall be interpreted by all conforming readers. Each of these boxes conforms to the standard box structure as defined in Annex A.3.

B.1 File Level Boxes

B.1.1 JPEG 2000 Signature box

Box type: 'jp\040\040' (0x6A502020)
Container: File Mandatory: Yes
Quantity: Exactly one
Location: First box in file

The format and structure of the JPEG 2000 Signature box is defined in Recommendation ITU-T T.800 | ISO/IEC 15444-1 Annex I.5.1.

B.1.2 File Type box

Box type: 'ftyp' (0x66747970)
Container: File Mandatory: Yes
Quantity: Exactly one
Location: Immediately after the JPEG 2000 Signature box

The format and structure of the File Type box is defined in Recommendation ITU-T T.800 | ISO/IEC 15444-1 Annex I.5.1.

The BR field in the File Type box shall be 'jpm\040' for files that are completely defined by this International Standard. In addition, a file that conforms to this International Standard shall have at least one CLi field in the File Type box, and shall contain the value 'jpm\040' in one of the CLi fields in the File Type box.

B.1.3 JP2 Header box (superbox) after File Type box

Box type: 'jp2h' (0x6A703268)
Container: File Mandatory: No
Quantity: At most one
Location: Anywhere after the File Type box

The format and structure of the JP2 Header box is defined in Annex B.6.2.

The first file level JP2 Header box following the File Type box describes a document level thumbnail of the JPM file. The first contiguous codestream following a file level JP2 Header box shall contain the associated thumbnail codestream. If the JP2 Header box and the Contiguous Codestream box are JP2 compatible, then the JPM file is JP2 compatible and 'jp2\040' occurs in the compatibility list in the File Type box. See 5.2.5 for more details.

B.1.4 Compound Image Header box

Box type: 'mhdr' (0x6D686472)
Container: File Mandatory: Yes
Quantity: Exactly one
Location: Anywhere after the File Type box

The Compound Image Header box contains general information about the compound image. It is recommended that this box be placed immediately after the File Type box and any file level JP2 Header box.

The type of a Compound Image Header box shall be 'mhdr' (0x6D686472). This box contains the following fields.