



ISO Base Media File Format and Apple HEVC Stereo Video

Format additions

Version 1.0

June 9, 2025

Note: The information contained within this document is preliminary and is subject to change.

- Introduction3**
- References3**
- Stereo Video4**
 - Stereoscopic, Stereopsis and Stereo Media4
 - Stereoscopic Video Tracks4
 - Multiview Video Tracks and MV-HEVC Compression.....5
- Video Extended Usage5**
 - Video Extended Usage Box Hierarchy6
 - 1. Video Extended Usage ('vexu') box7
- Use of other signaling extensions22**
 - Horizontal field-of-view box22
- Auxiliary Video Track Handler Type23**
- Spatial Audio23**
 - Spatial Audio Technologies24
- Timed Metadata and Spatial Media24**
 - Caption parallax timed metadata items24
- Conclusion25**
- Document Revision History26**

Introduction

This document describes Apple extensions of, or specialized use of, the ISO Base Media Format (a.k.a. ISOBMFF) to support spatial media. These extensions also apply to the QuickTime File Format. Spatial media is intended to produce a richer experience for the user; whether a richer audio experience, a richer visual experience, or a combination of the two.

Like the QuickTime File Format (QTFF) upon which it is based, the ISOBMFF format is meant to serve as a container of media using tracks and movie-level structures. The movie format of that media continues to evolve from the earliest “postage stamp” (i.e., very low resolution) video with one- and two-channel uncompressed or barely compressed audio to modern formats performing remarkable levels of visual compression for 4K and even 8K video with very rich multichannel, ambisonic and object-based audio encoding. This is coupled with facilities to carry and present captions such as WebVTT text tracks or closed captions embedded with video. Static and timed metadata can be carried to augment the presented media. Still other kinds of media tracks have been supported and will likely get added.

To support media that delivers rich spatial experiences, the QTFF and ISOBMFF foundations are being extended with new media formats, with extensions to supported media formats, and with new constructs to inform relationships among the new and earlier supported media. Some of these extensions are specific to their spatial nature, whereas others are fundamental and used by the former. This is all intended to be done in a way—where possible—so an existing ISOBMFF or QTFF player or processor can interact with the spatial media possibly in a reduced but compatible form while allowing new playback or processing to take fuller advantage of the newly afforded richness.

This document describes new and updated file format structures to support spatial media. Some of these structures are accessible through Apple AVFoundation and CoreMedia framework interfaces and those serve as the preferred alternative to direct structural access when running on a platform with Apple frameworks available. Those reading or writing the format directly—pursuant to relevant licensing—should however be able to accomplish their goals with the structural descriptions in the following sections.

Another consideration for ISOBMFF is that it is used in a fragmented movie form for HTTP based delivery technologies such as HTTP Live Streaming. The support in standalone MP4 files and fragmented MP4 resources is much the same.

Note: The words “may”, “should”, and “shall” are used in the conventional specification sense, that is, respectively, to denote permitted, recommended, or required behaviors.

References

[QTFF] QuickTime File Format (QTFF), 2016

[ISOBMFF] ISO/IEC 14496-12:2020 ISO Base Media File Format

[ISONALU] ISO/IEC 14496-15:2019 “Carriage of network abstraction layer (NAL) unit structured video in the ISO base media file format”

[HEVC] ISO/IEC 23008-2:2020 “High efficiency video coding”

[METADATA] “Video Contour Map Payload Metadata within the QuickTime Movie File Format—Format Additions”

Stereo Video

Stereoscopic, Stereopsis and Stereo Media

Just as stereo audio indicates different audio for the left and the right ear, visual media can be stereoscopic in which a view is available to be presented to the left eye and another view is available to be presented simultaneously to the right eye. The presentation of both the left and right views allows for an effect known as *stereopsis* which can be defined as:

the perception of depth produced by the reception in the brain of visual stimuli from both eyes in combination; binocular vision.

The production and display of this is sometimes referred to in cinema as *3D*, and the implementation and storage of the views can vary. This cinema use of *3D* should be distinguished from *3D rendering* involving a framework like Apple's Metal framework, where geometry, materials, lighting and cameras are modeled and rendered by a GPU or CPU. In the latter case, such three-dimensional rendering might produce a view as seen from the left eye and a simultaneous view seen from the right eye and therefore be stereoscopic. Rendering of a scene might however produce a single view that is not stereoscopic. This is sometimes called *monoscopic*.

Stereoscopic media can also be captured photographically, where two cameras are offset horizontally to produce a video where the left-eye view and the right-eye view are each encoded. In this case, there's not necessarily any modeling of the scene or any GPU rendering. Playback presents the left captured view to the viewer's left eye and the right captured view to the viewer's right eye. These left and right captured views might also have been processed. Although storage strategies can vary, this document describes how to store stereoscopic content using standardized ISO/ITU formats and some extensions to QTFF/ISOBMFF.

Stereoscopic Video Tracks

QTFF/ISOBMFF standalone and fragmented movies can include a single video track associated with both the left and the right eye. This video track's access units carry both a base and secondary layer that correspond to a primary stereo eye view (left or right) and the complementary stereo eye (i.e., right if the primary is the left, left if the primary is the right).

The expectation is that both stereo-eye views (i.e., the left-eye view and right-eye view) will be shown to both the left and right eyes simultaneously. Both stereo views are available and synchronized according to the movie timeline. When played, stereopsis is achieved.

Note: The ability to produce a movie with just one stereo eye video track (whether left or right) can be useful in production workflows. Two tracks in the same movie or in two movies might be useful. These might be combined into a new movie either by encoding with both views in one video track or less commonly by carrying two video tracks. Though potentially applicable to QTFF/ISOBMFF, it is not a described use case here.

This document introduces a `VisualSampleEntry` extension that can signal among other things whether the associated video track is stereoscopic and which stereo eyes are carried in that video track. This new signaling is referred to as *video extended usage* and is described in a later

section of this document. For now, the point is that this allows a movie reader to detect stereo-related video tracks and to identify the stereo eyes contributed by that track so it can configure presentation or other processing. Whereas the video track itself uses a video-compression format and signals it has a left and right stereo view, the video extended usage is meant to be more easily parsed and to be applicable to non-multiview video (i.e., not MV-HEVC).

Multiview Video Tracks and MV-HEVC Compression

A QTFF movie (i.e., .mov) or ISOBMFF (e.g., .mp4) movie video track carries video as either uncompressed or compressed video media samples. In the case of compressed video, High Efficiency Video Coding [HEVC] defines extensions to encode more than one view in the compressed bitstream for each coded video frame (or access unit). Defined in Annex G of the HEVC spec [HEVC], *Multiview High Efficiency Video Coding* defines how layers corresponding to views can be encoded and associated. This is sometimes written as *MV-HEVC*, for *Multiview HEVC*.

The QuickTime ImageDescription or ISOBMFF VisualSampleEntry (each referred to as *visual sample-entry*) shall include a *video extended usage* visual sample-entry extension box (described later in this document) indicating which stereo eye views—left, right or both—are carried in the MV-HEVC video track. For MV-HEVC, both left and right-eye views should be available. A hero eye indicating the default stereo eye may optionally be signaled. This construct allows a client to determine the stereoscopic nature of the video track without needing to parse for MV-HEVC bitstream details in the decoder configuration.

The video bitstream requirements of MV-HEVC coding, visual sample-entry and video media samples are described in the document *Apple HEVC Stereo Video Interoperability Profile*.

Video Extended Usage

This specification introduces an optional visual sample-entry extension that can signal additional aspects regarding the use of the video track's decoded frames. The new extension is called the *video extended usage* and uses the box type 'vexu' (optionally pronounced as "vex you"). Details necessary for video frame decoding are still carried in the visual sample-entry header (e.g., the compression type or the dimensions), as well as in other visual sample-entry extensions (e.g., 'colr' and compression type specific decoder configurations). The 'vexu' extension describes aspects beyond fundamental decode. For instance, it may specify that the video frames are stereoscopic or otherwise organized, requiring the video frames to be processed or displayed in a special way before presenting to the viewer.

Traditionally, a video track within a movie file or movie fragments can be decoded and immediately presented with little additional processing other than perhaps scaling, cropping and placement. For video captured in the real world, such as from a camera or computer, this is the norm. Even non-linear editing mostly works with video as stored in the movie files, perhaps applying effects, but otherwise encoding video that is directly presentable.

Increasingly today, a video track may be used as input into a rendering process and may not be suitable to show a viewer immediately. For example, a stereoscopic "3D" movie should present the left-eye view to the viewer's left-eye, and the right-eye view to the viewer's right eye. To do this, it is important to know first that the video track delivers stereoscopic views, and secondly which of those views are available. For MV-HEVC video, the presence of [HEVC] and [ISONALU] constructs such as the 'hvcC' and '1hvc' extension data might seem sufficient, but

unfortunately requires all readers to parse significant HEVC detail. Generally, a guiding principle is not to rely upon codec-level signaling to understand things that are needed at the container level. Here, the video track's 'vexu' visual sample-entry extension serves the role of easy-to-interpret signaling that the video is stereoscopic. The 'vexu' visual sample-entry extension must be consistent with what is signaled in the decoder configuration.

Beyond MV-HEVC, it may be desirable to use other video-compression formats (e.g., non-Multiview HEVC) or uncompressed video to carry stereoscopic video without requiring their fundamental decoded bitstream to signal the stereo use—something the format might not support. A video extended usage extension can be added indicating that a video track carries two stereo eyes or is for only one of the two stereo eyes. A 'vexu' extension can also be added indicating that the decoded video is organized in some other way described in a future version of this specification. This can be combined with stereoscopic detail that there is both a left and a right stereo eye view. Here, playback and processing need to understand the video track uses this alternative organization so it can route the left view and the right view portions of the decoded frame to the respective viewer eye.

Video Extended Usage Box Hierarchy

The video extended usage extension box specifies the usage of the decode of the video samples and details relevant to that usage. This is an optional extension and needed only when special or useful interpretation of the video in playback or processing is required. If the state signaled is not required for playback or processing, the extension may still be present but there is no expectation the reader understands it.

This extension box is a box hierarchy and contains further boxes signaling particular aspects of the video. These contained boxes may be leaf boxes—typically a FullBox—or box hierarchies themselves. There is also a mechanism to indicate that contained boxes must be understood by a reader and if not, that that part of the box hierarchy has failed to be processed. That error can propagate upwards, failing within a local subtree or even in the entire video extended usage box extension. This can in turn indicate the video should not be presented or processed as the reader's implementation lacks sufficient support.

Table 1. Current box types defined in the 'vexu' box hierarchy

FourCC	FourCC	FourCC	FourCC	Box syntax element
vexu				VideoExtendedUsageBox
	must ₁			RequiredBoxTypesBox
	eyes			Video Stereo
		must ₁		RequiredBoxTypesBox
		stri		StereoViewInformationBox
		hero		HeroStereoEyeDescriptionBox
		cams		StereoCameraSystemBox
			blin	StereoCameraSystemBaselineBox
		cmfy		StereoComfortBox
			dadj	StereoComfortDisparityAdjustmentBox
	proj			ProjectionBox
		must ₁		RequiredBoxTypesBox
		prji		ProjectionInformationBox
		. . .		One of different per projection type boxes if needed
hfov				HorizontalFieldOfViewBox

Note: must₁: A must box may occur as a child box in any box hierarchy, to indicate reader responsibility for understanding any child boxes of the must box's parent box. These locations are not enumerated in the above table. Exemplar must boxes are included in the table, but these positions should not be considered exhaustive.

1. Video Extended Usage ('vexu') box

This section describes how the video extended usage extension box is organized and the constituent boxes.

1.1. Definition

Box Type: 'vexu'

Container: Visual Sample Entry (different coding types)

Mandatory: No

Quantity: Zero or one

The video extended usage extension is a QuickTime File Format atom [QTFF] which is the same as a Box in ISO/IEC 14496-12 [ISOBMFF]. As we will use the bitstream syntax from ISO/IEC

14496-12, we will use “box” interchangeably with “atom”. References to ImageDescription for QTFF are also interchangeable with references to the ISO 14496-12 VisualSampleEntry.

The video extended usage box is held in a VideoExtendedUsageBox and has the ISO Box box type of ‘vexu’, for *video extended usage*.

As a box, it can contain zero or more child boxes that together signal the nature of the associated track samples’ extended usage. Having no child boxes is valid but likely not useful. Having only child free-space boxes (i.e., a FreeSpaceBox) is appropriate if the intention is to reserve space in the VisualSampleEntry.

To allow new or otherwise unknown VideoExtendedUsageBox child boxes to be introduced while allowing older readers to know they do not understand enough to process or present the video track, a mechanism is introduced to indicate mandatory and by implication optional child boxes. Additionally, child boxes can indicate whether their own structure can be optional so readers can recognize versions they do not support.

New Boxes should not be introduced into the VideoExtendedUsageBox unless documented in this specification or a successor version of this specification.

1.2. Syntax

```
aligned(8) class VideoExtendedUsageBox extends Box('vexu') {
    RequiredBoxTypesBox();           // optional if no required boxes specified
    StereoViewBox();                 // optional
    ProjectionBox();                 // optional
    Box() any_box;                   // other optional boxes with FreeSpaceBox()
    reserved for its expected use
}
```

1.3. Semantics

The VideoExtendedUsageBox contains zero or more child boxes that signal something about the use of the video. Child boxes will be defined in this specification now or in the future. Child boxes might be defined in external specifications but the box_type used there should be registered so as not to collide with boxes introduced in this or related specifications. The order of child boxes in the VideoExtendedUsageBox and in all contained boxes recursively is not prescribed. A reader should be prepared to find boxes in any order.

Note: As FreeSpaceBox (‘free’) has a very common meaning in ISO BMFF and QTFF, one or more free-space boxes may occur among the child boxes and should be interpreted as having no other meaning than taking up space. There is no guarantee that the payload of a FreeSpaceBox contains exclusively zero (0) bytes, but that is encouraged.

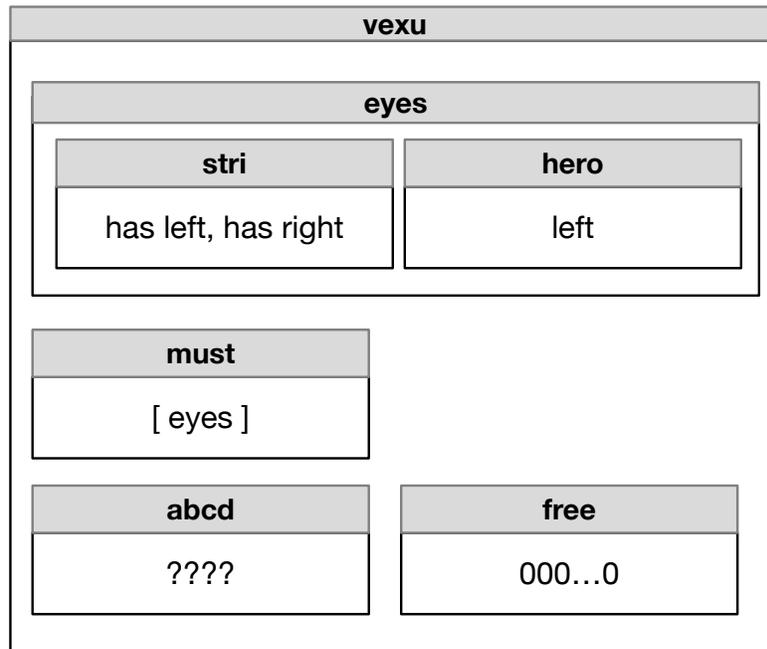
Note: An empty VideoExtendedUsageBox (i.e., containing no child boxes) is allowed but should generally not be included in the VisualSampleEntry. It may however be useful to reserve space by including a VideoExtendedUsageBox in concert with a contained FreeSpaceBox (‘free’).

New child boxes may be introduced in the future that are not described in this spec. There is a mechanism in the structure of `VideoExtendedUsageBox` to signal the set of child boxes that an implementation must understand in order to usefully process the video track. This allows future boxes to be introduced and older implementations to know they should not present the video with newer signaled features.

Besides standard boxes, the `VideoExtendedUsageBox` may contain zero or more boxes that describe specific kinds of signaling. In the following section, each kind of box is described.

Note that a `VideoExtendedUsageBox` should carry only one child box for a specific feature. So, for example, there should not be more than one feature box for stereo view signaling.

A VideoExtendedUsageBox with contained boxes might look like the following:



This 'vexu' box here contains an 'eyes' box and another box. This second box, the 'abcd' box, is optional, but if its type is included in a 'must' box, that indicates that it must be understood by the reader. 'free' boxes allow space to be reserved. The *other* box (i.e., 'abcd') represents an unknown but not required box.

The order of child boxes within a box may vary. Readers should not expect a fixed order of child boxes at any level. A writer should not include a child box of a particular type more than once if it is documented to occur only once.

1.2. Required box types ('must') box

1.2.1. Definition

Box Type: 'must'

Container: A video box within the video extended usage box ('vexu') or within contained Boxes

Mandatory: No

Quantity: Zero or one

If a parent box at any level within the VideoExtendedUsageBox has a 'must' box, that 'must' box contains a list of box types corresponding to boxes that are peers to the 'must' box and that reader must successfully interpret in order for the parent box to be successfully interpreted. In other words, if the reader does not recognize a required box type, or if it fails to parse that box or any *required* child box of that box, the reader must consider that to be a failure to parse the parent box. If the VideoExtendedUsageBox box is considered failed, the track is to be ignored by the reader.

Each box within the `VideoExtendedUsageBox` might serve to signal a feature according to this specification. The set of boxes is interpreted by the reader to understand what the video represents (e.g., it uses stereo views). Some of the signaling is necessary for further processing. Other boxes may be informative but not strictly required for interpretation. It's important to understand all required boxes. `RequiredBoxTypesBox` enables the adding of new boxes in the future that may be required for interpretation and further processing.

Each box within the `VideoExtendedUsageBox` may contain any hierarchy of boxes suitable to signal some aspect about the video. Some of these may be boxes with a hierarchy of other boxes and some may be full boxes. The `RequiredBoxTypesBox` enumerates the box types of its sibling boxes corresponding to required boxes. If not enumerated within the `RequiredBoxTypesBox`, the child box's interpretation is optional.

The `RequiredBoxTypesBox` contains an array of FourCCs corresponding to box types. If an entry is 0, the entry is reserved and is not interpreted as a required box type. Free-space boxes of box type 'free' should not be included in a `RequiredBoxTypesBox`. If a `RequiredBoxTypesBox` includes a box type that doesn't correspond to a child box, the reader can ignore the absence but might want to log this for diagnostic purposes. The use of box types of missing boxes within a `RequiredBoxTypesBox` is however discouraged.

The `FreeSpaceBox` box type of 'free' should not be referenced from a `RequiredBoxTypesBox`.

The `RequiredBoxTypesBox` can also occur within other boxes within the `VideoExtendedUsageBox` box hierarchy that are themselves box hierarchies. These uses of `RequiredBoxTypesBox` serve to indicate local requirements on boxes that must be recognized and understood for local parsing to be valid. A local box can fail and that influences the validity of the parent box if the parent box itself is referenced from another `RequiredBoxTypesBox` that is a sibling of the parent box.

1.2.2.Syntax

```
aligned(8) class RequiredBoxTypesBox extends FullBox('must', 0, 0) {
    unsigned int(32) required_box_types[];
}
```

1.2.3.Semantics

`required_box_types` is an array of zero or more box types corresponding to sibling boxes that must be understood by readers to properly process the video associated with the `VideoExtendedUsageBox`. For each non-zero entry in `required_box_types[]`, the reader should confirm the box type is recognized. A value of zero (0) in a `required_box_types[]` entry can be ignored, allowing for space for entries to be reserved.

1.2.4. Reader Behavior and RequiredBoxTypesBox

A reader of a movie file video track with an associated VideoExtendedUsageBox should be able to detect whether it understands enough about the VideoExtendedUsageBox contents to process the video beyond fundamental decoding. This further processing, interpretation and/or rendering is what *extended* refers to within within the identifier name "VideoExtendedUsageBox".

This specification is intended to be extended in future versions. A particular video track may carry several kinds of signaling that differ from other video tracks within the movie file or in other movie files. The kind of signaling within a box can itself evolve over time. In all these cases, it is important to know if the set of child boxes of a box must be understood. While the most obvious case is child boxes of VideoExtendedUsageBox, the approach can apply to any Box serving as the root of a box hierarchy within the larger hierarchy.

The following describes reader behavior that is aware of RequiredBoxTypesBox:

1. Read (or start processing) the box hierarchy (e.g., VideoExtendedUsageBox).
2. Retrieve the contained RequiredBoxTypesBox child box if any.
 1. If present, confirm all non-zero entries of `required_box_types[]` are recognized box types and if not treat the parent box of the RequiredBoxTypesBox as not processable.
 2. Ignore all zeroed entries of `required_box_types[]`.
3. Enumerate each non-zero box type in the `required_box_types[]` of the child RequiredBoxTypesBox of the VideoExtendedUsageBox using an index from 0 to the length of `required_box_types[]` minus 1 and confirm the referenced box is understood.
 1. For each successive `index`, retrieve the child box with `required_box_types[index]` and confirm understanding of its structure.
 1. For a child FullBox, the reader should consider the version and flags to confirm understanding as well as anything else that may be relevant to its interpretation.
 2. For child boxes that are box hierarchies themselves and allow RequiredBoxTypesBox, the reader should descend into the box, retrieve the optional contained RequiredBoxTypesBox and perform this algorithm recursively.
 3. If the parsing of the FullBox or the child box hierarchy fails, the reader should treat the current level as invalid and propagate that failure upwards. Any semantics discovered at the current level should not be propagated upwards as partial semantics is

misleading (e.g., stereo view and something else both being required should not signal stereo views if the other parsing fails).

The VideoExtendedUsageBox parsing follows this same algorithm. In this case, however, failing to parse required child boxes of VideoExtendedUsageBox means the track has failed to parse. This can best be interpreted as though the entire video track is unavailable.

1.3. Video stereo-view signaling

The StereoViewBox signals if the video track represents stereo 3D content. This can take the form of a track that delivers both a left stereo eye and a right stereo-eye view or a track that carries only the left stereo eye or only the right stereo eye.

If both left and right stereo eyes are carried, the views might be combined in one image and organized in some way or they might be contained in some kind of multiview coding.

If the left stereo eye is in one video track and the right stereo eye is in a second video track, each carries its own VideoExtendedUsageBox with a StereoViewBox. The indication of which eye is carried will be appropriate for each corresponding video track.

For completeness, it is also possible to signal monoscopic video which is to say no stereo view carriage. If this is the case, however, the StereoViewBox can be eliminated from the VideoExtendedUsageBox. If the VideoExtendedUsageBox would be left with no child boxes, the VideoExtendedUsageBox can be eliminated from the VisualSampleEntry as well.

If the recorded stereo video has a designated “hero” eye, the StereoViewBox carries a HeroStereoEyeBox. There are rules that require signaling when the stereo eye video is separated into two video tracks, with each track carrying only one of the stereo eyes.

1.3.1. Definition

Box Type: 'eyes'

Container: Video extended usage box ('vexu')

Mandatory: No

Quantity: Zero or one

1.3.2. Syntax

```
aligned(8) class StereoViewBox extends Box('eyes') {
    RequiredBoxTypesBox();           // as needed
    StereoViewInformationBox();
    HeroStereoEyeDescriptionBox();    // optional
    StereoCameraSystemBox();          // optional
    StereoCameraSystemBaselineBox(); // optional
    Box();                             // other optional boxes
}
```

1.3.3. Semantics

`StereoViewInformationBox` is a required box indicating which stereo eyes are present.

`RequiredBoxTypesBox` indicates the box types for other boxes that must be understood to interpret the current version of the `StereoViewsBox`. The `StereoViewInformationBox` box type of 'stri' is required within `RequiredBoxTypesBox` if a `RequiredBoxTypesBox` is used.

Other boxes indicate additional detail about the stereo view representation and are described in later sections of this document. The set of boxes may evolve.

1.4. Stereo view information ('stri')

1.4.1. Definition

Box Type: 'stri'

Container: Video stereo-view box ('eyes')

Mandatory: Yes

Quantity: One

The `StereoViewInformationBox` can carry the stereography related information indicating the presence of particular stereo eyes (i.e., left stereo eye, right stereo eye) as well as some other flags.

1.4.2. Syntax

```
aligned(8) class StereoViewInformationBox extends FullBox('stri', 0, 0) {
    unsigned int(4) reserved;          // reserved, set to 0
    unsigned int(1) eye_views_reversed;
    unsigned int(1) has_additional_views;
    unsigned int(1) has_right_eye_view; // video contains a right-eye view
    unsigned int(1) has_left_eye_view;  // video contains a left-eye view
}
```

1.4.3. Semantics

`has_left_eye_view`: Indicates the stereo left eye is present in video frames

`has_right_eye_view`: Indicates the stereo right eye is present in video frames

`has_additional_views`: Indicates that one or more additional views may be present beyond stereo left and stereo right eyes (e.g., a "centerline" view)

`eye_views_reversed`: Indicates that the order of the stereo left eye and stereo right eye are reversed from the default order of left being first and right being second

`reserved`: 4 bits reserved for future versions of this specification; for this version of this specification, writers should set it to 0 and readers should treat any non-zero values as if this box is invalid.

Because there is a flag field for the left eye and a field for the right eye, both fields should be set to indicate that both eyes are represented in video frames. Moreover, both

`has_left_eye_view` and `has_right_eye_view` can be set to 0 to indicate that the frame is monoscopic.

Note: If the video is monoscopic, the `StereoViewBox` can also be absent from the `VideoExtendedUsageBox`. If the only signaling is of monoscopic video, the `VideoExtendedUsageBox` can be absent from the `VisualSampleEntry`, too.

If an alternative organization is signaled in the future, the default order of stereo eyes in video will be left eye first, then right eye. Setting the `eye_views_reversed` field reversed the order, so the right view appears to the left of the frame, and the left view appears to the right of the frame. For multiview coding, there is no implied ordering and the `eye_views_reversed` field should be set to 0.

Note: It may be useful to signal in a multiview coding approach the presence of the left stereo eye, the right stereo eye and a third view which is the “centerline” or “down the nose” view which is between these and is neither the left nor the right. It may not be possible or appropriate to use the left or the right eye for this central view. There may be coding efficiencies from being able to include such a view in multiview coding.

The `has_left_eye_view` and `has_right_eye_view` fields specify the presence of the left and right stereo eye views but the fields do not signal how those are stored. That is accomplished with other child boxes of `StereoViewBox`.

Note that the `has_additional_views` field indicates that views beyond those for the left eye and the right eye are present. One example of this might be a “centerline” view. Note that signaling the presence of the “centerline” is not necessary if both the left and right eye flags are zeroed indicating a monoscopic view.

1.5.Hero Stereo Eye Description

1.5.1. Definition

Box Type: ‘hero’

Container: Video stereo-view box (‘eyes’)

Mandatory: No

Quantity: Zero or one

The `HeroStereoEyeDescriptionBox` indicates which stereo eye, if any, has been denoted as the “hero” eye. If so signaled, this indicates that the other stereo eye view is derived from the specified stereo eye and may be useful when choosing which eye to use in a monoscopic viewing environment. If neither eye is the hero eye, the `HeroStereoEyeDescriptionBox` does not need to be included in the `StereoViewBox`. If the hero eye is not known, a `HeroStereoEyeDescriptionBox` might not appear in the `StereoViewBox`.

It is possible to include a `HeroStereoEyeDescriptionBox` but set the flags to indicate that neither the left nor the right stereo eye are set. Though unconventional, this allows an implementation to reserve space for the box, to potentially set later in

processing. Readers should be prepared to recognize such a HeroStereoEyeDescriptionBox that signals no hero eye.

1.5.2. Syntax

```
aligned(8) class HeroStereoEyeDescriptionBox extends FullBox('hero', 0, 0)
{
    unsigned int(8) hero_eye_indicator; // 0 = none, 1 = left, 2 = right, >=
3 reserved
}
```

1.5.3. Semantics

HeroStereoEyeDescriptionBox is used to indicate which of the left or right stereo eye is the “hero” eye, if any.

hero_eye_indicator: is used in the HeroStereoEyeDescriptionBox, to signal which hero eye, if any is specified. Defined values are:

0: The hero eye is not specified

1: Indicates the left eye is the hero eye

2: Indicates the right eye is the hero eye

>= 3: Reserved, and should not be used for implementation of this version of this specification. If a reserved value is read, a reader should treat the signaling as though no hero eye is specified. If the hero eye is not specified, it is recommended that HeroStereoEyeDescriptionBox not be included in the StereoViewBox. The value of 0 is allowed as it can be used to reserve space for the HeroStereoEyeDescriptionBox that might be adjusted later.

The HeroStereoEyeDescriptionBox signals the left or the right stereo eye independently of whether or not the stereo-view box's

StereoViewInformationBox indicates that the order of the stereo eyes is reversed in order. The hero left eye is always the left stereo eye and the hero right eye is always the right stereo eye.

1. Stereo Camera-System Box

1.1. Definition

Box Type: 'cams'

Container: Stereo-view box ('eyes')

Mandatory: No

Quantity: Zero or one

Stereo views are typically produced through capture by a camera system of some kind. Characteristics of that camera system may be useful for rendering or other processing. Moreover, the kinds of information associated with a camera system may vary now or in the future. As such, it seems prudent to allow one or more kinds of information to be carried for consideration by the processing client. It is also the case that this information is optional, as some stereo recording might not use a physical camera system. It may, however, be the case

that the camera system is virtual, so being able to include information that a physical camera system produces also seems useful to allow.

To signal information about the stereo nature of the camera system, the `StereoCameraSystemBox` is introduced. This is an optional `Box` instead of a `FullBox` so that a hierarchy of boxes and full boxes related to the camera system can be carried.

1.2. Syntax

```
aligned(8) class StereoCameraSystemBox extends Box('cams') {
    RequiredBoxTypesBox();           // optional
    StereoCameraSystemBaselineBox(); // optional
    Box[]; // other boxes that signal information about the camera system
}
```

1.3. Semantics

`StereoBaselineBox` is used to describe the baseline dimension between centers of the stereo lenses of the camera system. It is optional. There is no default interpretation when this box is absent.

2. Stereo Camera-System Baseline Box

2.1. Definition

Box Type: 'blin'

Container: Stereo camera-system box ('cams')

Mandatory: No

Quantity: Zero or one

Stereo camera systems may provide a number of characteristics that may prove useful to downstream rendering and processing of stereo frames captured by a camera system. One such characteristic is the distance between the optical centers of the left and right stereo eye camera lenses.

To signal information about the camera-system baseline, the `StereoCameraSystemBaselineBox` is used. This is an optional `FullBox` that holds a field with the distance in micrometers between the lenses. As distances in a camera system are typically expressed in millimeters, this use of micrometers can also be seen as being expressed in thousandths of a millimeter. The value is a fixed-point number expressed as an integer (i.e., 63123 micrometers is 63.123 millimeters).

Note: Although `StereoCameraSystemBaselineBox` is not mandatory in the overall box structure of `VideoExtendedUsageBox`, its presence is required for the containing movie file to be considered "spatial media" by visionOS 2.

2.2. Syntax

```
aligned(8) class StereoCameraSystemBaselineBox extends FullBox('blin') {
```

```
    unsigned int(32) baseline_value;
}
```

2.3. Semantics

`baseline_value` holds the baseline dimension between centers of the stereo lenses of the camera system. It is an unsigned 32-bit integer that is interpreted in micrometers or thousandths of a millimeter (e.g., 63123 micrometers is 63.123 millimeters).

3. Stereo-Comfort Box

3.1. Definition

Box Type: 'cmfy'

Container: Stereo-view box ('eyes')

Mandatory: No

Quantity: Zero or one

Stereo views are presented to the left and the right eye in a display system. How that is performed can influence comfort. To allow different kinds of information related to stereo comfort to be carried with the stereo video frames, it is useful to allow one or more pieces of information to be carried.

To signal information influencing viewer comfort, the `StereoComfortBox` is introduced. This is a `Box` instead of a `FullBox`, so that a hierarchy of boxes and full boxes related to stereo comfort can be carried. This box is optional unless there are contained boxes. Only `StereoComfortDisparityAdjustmentBox` is currently defined.

3.2. Syntax

```
aligned(8) class StereoComfortBox extends Box('cmfy') {
    RequiredBoxTypesBox();           // optional
    StereoComfortDisparityAdjustmentBox(); // optional
    Box[]; // other boxes that signal information about stereo comfort
}
```

3.3. Semantics

`StereoComfortDisparityAdjustmentBox` is used to describe any adjustment in the disparity between the stereo views of the current frame. The absence of this box or a zero value in the `disparity_adjustment` field of the box indicates no change in stereo disparity.

4. Stereo-Comfort Disparity Adjustment Box

4.1. Definition

Box Type: 'dadj'

Container: Stereo-comfort box ('cmfy')

Mandatory: No

Quantity: Zero or one

Stereo view comfort can be improved by allowing an adjustment in disparity to be specified. The source of this disparity is not specified, but the value's carriage is. As this adjustment is optional, this information may not be present.

To signal information about changes to the stereo comfort disparity, the `StereoComfortDisparityAdjustmentBox` is introduced. This is a `FullBox` holding a field that indicates the amount of disparity to apply. If the value is 0, there is no disparity adjustment. The value is a signed 32-bit integer that is interpreted as a uniform number over the range [-1.0...0.0...+1.0]. The valid range of the integer is from -10000 to +10000, which maps from -1.0 to +1.0. The interval of 0.0 to 1.0 and 0.0 to -1.0 are each mapped over the width of a view's image. Half of this value is applied to each stereo eye view.

The value is interpreted this way:

- Half the disparity is added to pixels in the left stereo eye view.
- Half the disparity is subtracted from pixels in the right stereo eye view.

As the disparity value is signed, adding a negative value to the left stereo eye is equivalent to subtracting the absolute value of the disparity value. The right stereo eye's subtraction of a negative disparity value is likewise equivalent to adding the absolute value of the disparity value.

Another interpretation of the sign of the disparity value is that positive denotes increased disparity with respect to the parallel view direction (e.g., horizontal) and negative denotes increased negative disparity with respect to the parallel view direction. Negative disparity is toward the viewer.

If the disparity adjustment value is 0, the `StereoComfortBox` need not contain a `StereoComfortDisparityAdjustmentBox`. If the `StereoComfortBox` would be left with no child boxes, the `StereoComfortBox` can itself be missing. There are rules, however, requiring the `StereoComfortBox` to exist.

Note: Although `StereoComfortDisparityAdjustmentBox` is not mandatory in the overall box structure of `VideoExtendedUsageBox`, its presence is required for the containing movie file to be considered "spatial media" by visionOS 2.

4.2. Syntax

```
aligned(8) class StereoComfortDisparityAdjustmentBox extends
FullBox('dadj') {
    int(32) disparity_adjustment;
}
```

4.3. Semantics

`disparity_adjustment` holds a value describing an optional adjustment in stereo disparity. A value of 0 indicates there is no disparity adjustment and can also be represented by not including a `StereoComfortDisparityAdjustmentBox` in `StereoComfortBox`. The value is a signed 32-bit integer measured in the range of -10000 to 10000, mapping to the uniform range [-1.0...1.0]. The interval of 0.0 to 1.0 or 0 to 10000 maps onto the stereo eye view image width. The negative interval 0.0 to -1.0 or 0 to -10000 similarly maps onto the stereo eye view image width.

Note: This is not measured as a percentage but instead as a uniform value. To express 1.5%, the value 0.015 as a uniform value, or as the signed integer 150 (expressed over 10000), can be used.

Note: If the video is frame packed, for example, as side-by-side, the image width is for one stereo eye rather than the total image width of the side-by-side views. Other kinds of packings or arrangements also use the image width of the view, however stored or represented.

2. Video projection signaling

Projections

A *projection* is a mathematical model mapping 3D world points onto 2D points such as an image or sensor. This mapping in a camera system can be achieved through optical elements such as lenses, through software or through a combination of the two.

2.1. Definition

Box Type: 'proj'

Container: Video stereo-view box ('eyes')

Mandatory: No

Quantity: Zero or one

If the video frame has some form of projection or other mathematical transform, the projection box may be used to indicate the algorithm and the approach taken. Downstream rendering likely needs to perform a mapping to present the frame in a way the viewer understands.

`ProjectionBox` signals that a form of projection is required to present the video track. Its contents signal one kind of projection.

2.2. Syntax

```
aligned(8) class ProjectionBox extends Box('proj') {
    RequiredBoxTypesBox();
    ProjectionInformationBox();
    RectilinearProjectionBox(); // optional: one of the kinds of the
    possible projections (default if absent)
    Box[]; // other kinds of boxes
}
```

2.3. Semantics

`ProjectionInformationBox` indicates the kind of projection described by the `ProjectionBox`. Each possible projection kind may be associated with a sibling box with additional parameters for that projection kind.

`RequiredBoxTypesBox` indicates the box types for other boxes that must be understood to interpret the current version of the `ProjectionBox`.

`RectilinearProjectionBox` indicates that the projection is a rectilinear projection. More detail is found later in this specification.

Other kinds of projections may be introduced in future versions of this specification. Current versions of projections may also be extended in the future.

Note: The projection kinds may not require any additional parameterization, so child box types for a particular projection kind may not exist. Nevertheless, the `box_type` for such projection kinds is reserved.

3. Projection information ('prji')

3.1. Definition

Box Type: 'prji'

Container: Projection box ('proj')

Mandatory: No

Quantity: One

The kinds of projections described by a `ProjectionBox` can vary according to this specification, and in future versions of this specification. The `ProjectionInformationBox` indicates the kind of projection carried.

3.2. Syntax

```
aligned(8) class ProjectionInformationBox extends FullBox('prji', 0, 0) {
    unsigned int(32) projection_kind;    // a FourCC for the kind of
    projection
};
```

3.3. Semantics

`projection_kind` is a FourCC corresponding to the kind of projection. These currently include 'rect', 'equi', 'hequ' and 'prim' with the reserved 'fish' kind.

The `ProjectionInformationBox` can specify one kind of video projection for a video track. With this version of the specification, it should be one of the following `projection_kind` FourCCs with the corresponding projection kind of box:

- 'rect': `RectilinearProjectionBox` indicates that there is no projection, and serves as the default in the absence of the `ProjectionBox` within the `VideoExtendedUsageBox`. Because it is the default, this is likely not seen in movie files.

The `ProjectionBox` containing the `ProjectionInformationBox` may optionally have a child box, using a box type that corresponds to the `projection_kind`. This child box is a sibling of `ProjectionInformationBox`. It signals any parameters specific to that kind of projection, and reserves the corresponding box types. Although no occurrences of such boxes are defined in this specification, such usage is reserved for the future.

4. Rectilinear projection

4.1. Definition

Box Type: 'rect'

Container: Projection box ('proj')

Mandatory: No

Quantity: Zero or one

The rectilinear projection is the default for video. There is no further processing necessary to present video marked with just this projection. As the default, this can be signaled by not including a `ProjectionBox`, or even by not including a `VideoExtendedUsageBox` if there is no other signaling.

If this projection is used, the sibling `ProjectionInformationBox` `projection_kind` should be set to 'rect'.

It is possible to use this box to make it obvious what the intention is.

4.2. Syntax

```
aligned(8) class RectilinearProjectionBox extends FullBox('rect', 0, 0) {
    // fields reserved for future use
};
```

4.3. Semantics

`RectilinearProjectionBox` signals a rectilinear projection. No fields are currently defined within the box, but a child box of type 'rect' is reserved for future addition to this specification.

Use of other signaling extensions

Horizontal field-of-view box

The `VisualSampleEntry` defined an extension box to signal the horizontal field of view.

2. Definition

Box Type: 'hfov'

Container: Visual sample entry

Mandatory: No

Quantity: Zero or one

The horizontal field of view of the decoded video frame image may be important to know. This can be signaled with the optional `HorizontalFieldOfViewBox` extension to `VisualSampleEntry`. The horizontal field of view is an unsigned integer in 1000ths of a degree (e.g., 123.456 is represented as 123456).

The `HorizontalFieldOfViewBox` is optional, but if present, there can be at most one `HorizontalFieldOfViewBox` in a `VisualSampleEntry`.

Note: Although `HorizontalFieldOfViewBox` is not mandatory in the overall `VisualSampleEntry` structure parallel to `VideoExtendedUsageBox`, its presence is required for the containing movie file to be considered "spatial media" by visionOS 2. It may be present for other use cases beyond stereo video.

2.1.1.1.1. Syntax

```
aligned(8) class HorizontalFieldOfViewBox extends Box('hfov') {
    unsigned int(32) field_of_view;
};
```

2.1.1.1.2. Semantics

`field_of_view`: An unsigned 32-bit integer indicating the degrees, in 1000ths of a degree. A 104° field of view would be recorded as 104000.

Auxiliary Video Track Handler Type

To date, video tracks have used the handler type 'vide'. Other track types such as audio and metadata use their own handler types (i.e., 'soun' and 'meta', respectively). This has never been a problem because the decoded video can be presented as-is, though perhaps with scaling or cropping. With stereo video tracks, the decoded video may require additional processing such as view extraction before being presented to the user.

Movies or fragmented movie files for HTTP Live Streaming may now use the *auxiliary video* or 'auxv' handler type to "hide" the video track from naive reader decode and presentation. For example, this can be useful for a video track with an alternative layout of images, signaled using a video extended usage atom.

If the decoded video, however, displays in a backwards compatible way when delivered—such as MV-HEVC showing just a default view from the stereo pair—there is no need to use the 'auxv' handler type. Use 'vide' in this case. Also, HTTP Live Streaming mediates the media shown so that the multivariant playlist can serve to filter display of particular video streams.

In a production workflow where users expect to see and confirm the decoded video even if further processing might be expected when delivered to an end user, it is okay to use 'vide' so tools that already present or process video tracks can find the track.

Spatial Audio

The experience is enhanced if audio can represent the spatial acoustic environment. Just as listening to stereo audio is richer than listening to mono audio, even richer audio representations are possible with appropriate audio coding. A number of advanced audio technologies exist and they may be used in isolation or in combination.

QTFF/ISOBMFF audio tracks use audio codecs to encode and carry audio—uncompressed and compressed—and the codecs can use different audio technologies. Some technologies are

applicable to Spatial Audio, and when used in that way, the audio might be termed *Spatial Audio*. The ISOBMFF format can carry a wide range of uncompressed and compressed format, some supporting spatial playback. By introducing new audio codecs in audio tracks, the movie can carry spatial audio.

ISOBMFF movies may contain any supported audio format. As additional formats are supported, those may prove useful for delivering more richer experiences.

Spatial Audio Technologies

By way of a quick summary, there are three audio technologies typically used in the Spatial Audio realm.

Channel-based audio can include more than one audio channel, each of which is mapped onto the speaker layout. This is called *multi-channel* audio, and is typically used with 5.1 and 7.1 audio. The number and placement of these channels in the soundscape can be more varied and the channel count can be more or less than the six of 5.1 or the eight of 7.1. Indeed, stereo has two channels, so is in fact multichannel, but that term is almost never applied to stereo.

Another technology, termed *ambisonics*, is a modeling of three-dimensional audio in a 360-degree space. It allows for the recording, mixing and playing back of such audio. Just as multi-channel audio can vary in channel count, ambisonic audio can vary in *order* allowing more refined audio with higher degreed ambisonics. Audio is fixed in location but surrounds the user.

A third technology, *object-based audio*, models each sound source as an object with associated metadata describing three-dimensional placement and other relevant characteristics. Individual objects might be fixed in 3D or might move in 3D over time.

This specification does not prescribe which audio encoding formats or which technologies within those formats are used to realize a richer experience.

Timed Metadata and Spatial Media

Spatial media tracks such as video may benefit from having associated timed metadata. Although this might be injected in AVC or HEVC SEI signaling, an alternative is to use a parallel metadata track. This timed metadata track can use the ISOBMFF 'mebx' format's ability to carry a number of metadata items for a time range. Metadata item keys need not be related to other item keys allowing a flexible way to signal a variety of structural or descriptive information.

We consider one kind of timed metadata payload related to describing the parallax of decoded stereoscopic video frames.

Caption parallax timed metadata items

Traditionally, captions are placed in the horizontal and vertical axes over video. With the introduction of stereoscopic video, however, there is a risk of depth collision if captions are placed in Z so they might intersect with stereoscopic elements that have a parallax (i.e., horizontal disparity) that is less than the screen plane. This "depth conflict" can produce viewer discomfort. To account for this, captions can have their parallax adjusted to have a more negative parallax than the video elements so there is no collision.

The document “Video Contour Map Payload Metadata within the QuickTime Movie File Format—Format Additions” [METADATA] specifies the structure of a metadata payload structure that can serve to describe parallax values associated with 2D areas of a stereoscopic video frame. This metadata is specific to the time-aligned video frame.

This payload is carried as metadata items within samples within QuickTime File Format [QTFF] timed metadata or ISOBMFF [ISOBMFF] multiplexed metadata. Note: Both of these use the 'mebx' format type of the 'meta' track handler type. The payloads can also occur in fragmented movie files in both ISOBMFF and QTFF.

Conclusion

This document describes extensions to the QuickTime (.mov) and ISOBMFF movie formats. These extensions are introduced by Apple to allow for delivery of stereoscopic video, spatial audio and timed metadata signaling to influence parallax of any subtitles associated with the video. This is applicable to both ISOBMFF standalone movie and fragmented movie files. It attempts to build on existing structures where that was deemed appropriate and introduces new constructs where there was a perceived deficit or a benefit in introducing a new construct. The evolution of the QTFF/ISOBMFF format extensions described here may be taken through standards processes in time.

Document Revision History

Date	Revision	Notes
2025-06-09	1.0	Version 1.0, Added baseline, disparity adjustment, horizontal field of view. Added view projection with <code>rect</code> projection type.
2023-06-21	0.9	First version