

A URI-based approach for addressing fragments of media resources on the Web

Erik Mannens · Davy Van Deursen · Raphaël Troncy ·
Silvia Pfeiffer · Conrad Parker · Yves Lafon ·
Jack Jansen · Michael Hausenblas · Rik Van de Walle

Published online: 31 December 2010
© Springer Science+Business Media, LLC 2010

Abstract To make media resources a prime citizen on the Web, we have to go beyond simply replicating digital media files. The Web is based on hyperlinks

E. Mannens (✉) · D. Van Deursen · R. Van de Walle
ELIS—Multimedia Lab, Ghent University—IBBT, Ghent, Belgium
e-mail: erik.mannens@ugent.be

D. Van Deursen
e-mail: davy.vandeursen@ugent.be

R. Van de Walle
e-mail: rik.vandewalle@ugent.be

R. Troncy
Multimedia Communications Department, EURECOM, Sophia Antipolis, France
e-mail: raphael.troncy@eurecom.fr

S. Pfeiffer
Vquence, Sydney, Australia
e-mail: silviapfeiffer1@gmail.com

C. Parker
Kyoto University, Kyoto, Japan
e-mail: conrad@metadecks.org

Y. Lafon
W3C/ERCIM, Sophia Antipolis, France
e-mail: ylafon@w3.org

J. Jansen
Distributed Multimedia Languages and Infrastructures, CWI, Amsterdam, Netherlands
e-mail: Jack.Jansen@cwi.nl

M. Hausenblas
Digital Enterprise Research Institute—LiDRC, National University of Ireland,
Galway, Ireland
e-mail: michael.hausenblas@deri.org

between Web resources, and that includes hyperlinking out of resources (e.g., from a word or an image within a Web page) as well as hyperlinking into resources (e.g., fragment URIs into Web pages). To turn video and audio into hypervideo and hyperaudio, we need to enable hyperlinking into and out of them. The W3C Media Fragments Working Group is taking on the challenge to further embrace W3C's mission to lead the World Wide Web to its full potential by developing a Media Fragment protocol and guidelines that ensure the long-term growth of the Web. The major contribution of this paper is the introduction of Media Fragments as a media-format independent, standard means of addressing media resources using URIs. Moreover, we explain how the HTTP protocol can be used and extended to serve Media Fragments and what the impact is for current Web-enabled media formats.

Keywords Media fragments · W3C standardisation · HTML5

1 Introduction

Video clips on the World Wide Web (WWW) used to be treated as “foreign” objects as they could only be embedded using a plugin that is capable of decoding and interacting with these clips. The HTML5 specification is a game changer and all of the major browser vendors have committed to support the newly introduced `<video>` and `<audio>` elements [7].¹ However, in order to make video clips accessible in a transparent way, it needs to be as easily linkable as a simple HTML page. In order to share or bookmark only the interesting parts of a video, we should be able to link into or link out of this time-linear media resource. If we want to further meet the prevailing accessibility needs of a video, we should be able to dynamically choose our preferred tracks that are encapsulated within this video resource, and we should be able to easily show only specific regions-of-interest within this video resource. And last but not least, if we want to browse or scan several video resources based on (encapsulated) semantics, we should be able to master the full complexity of rich media by also enabling standardised media annotation [21, 29]. Note that we can generalise the above observations to other media, such as audio resources. This way, media resources truly become “first-class citizens” on the Web.

The mission of the W3C Media Fragments Working Group (MFWG) [31], which is part of W3C's Video in the Web activity,² is to provide a mechanism to address media fragments on the Web using Uniform Resource Identifiers (URIs) [6, 9]. The objective of the proposed specification is to improve the support for the addressing and retrieval of sub-parts of so-called media resources (e.g., audio, video and image), as well as the automated processing of such sub-parts for reuse within the current and future Web infrastructure [19]. Example use cases are the bookmarking or sharing of excerpts of video clips with friends in social networks, the automated creation of fragment URIs in search engine interfaces by having selective previews, or the annotation of media fragments when tagging audio and video spatially and/or

¹At the time of writing, the following browsers support the HTML5 media elements: IE 9, Firefox 3.5, Chrome 4, Safari 4, Opera 10.

²<http://www.w3.org/2008/WebVideo/Activity.html>

temporally [18]. The examples given throughout this paper to explain the Media Fragments URI specification are based on the following two scenarios. In scenario (a), Steve –a long-time basketball enthusiast– posts a message on his team’s blog containing a Media Fragment URI, that highlights 10 s of an NBA video clip showing the same nifty move that he himself performed in last Saturday’s game. In scenario (b), Sarah –a video artist by profession– quickly previews the video footage in search of a particular high quality 10 s sub-clip to finalise editing her new video clip ‘The Editors’.

In this paper, we present the rationale for a Media Fragments specification in Section 2. In Section 3, we outline the boundaries and semantics of a Media Fragment URI and show how the syntax should look like, whereas Section 4 elaborates on how a media fragment specified as an URI fragment can be resolved stepwise using the HTTP protocol [8]. We then identify the influence of the current media formats on fragment extraction in Section 5. Finally, we outline open issues and future work in Section 6 and give our conclusions in Section 7.

2 Related work

Before video can become a “first-class citizen” on the Web, one urgently needs to provide standardised ways to localise the temporal, spatial, and track sub-parts of audio-visual media content [6, 29]. Previous efforts to do so include both URI-based and non-URI based mechanisms.

In the non-URI based class of solutions, the SMIL specification over HTTP allows to play only a temporal fragment of the video by using the *clipBegin* and *clipEnd* attributes [1]. However, current implementations have to first get the complete media resource and then cut it up locally, which entails a terrible waste of bandwidth in case of large video resources. Using MPEG-7, a video is divided into *VideoSegments* that can be described by a *MediaTimePoint* and *MediaDuration* corresponding to the starting time and media duration respectively [10]. Using TV-Anytime, temporal intervals can also be defined, accessed, and manipulated through *segments* within an audio/video stream using MPEG-7 types for specifying the temporal boundaries [3]. For images, one can use either MPEG-7 or an SVG code snippet to define the bounding box coordinates of specific regions [4]. However, this approach implies an extra *indirection* since a semantic description of this region will actually be *about* a piece of an XML document just defining a multimedia fragment and not the fragment itself. As such, the identification and the description of the temporal fragment or region is intertwined (this use of indirection) and one needs to first parse and understand the metadata model in order to get access to the media fragment afterwards, which is not desirable in all circumstances. Finally, HTML ImageMaps can also define spatial regions (rectangles, circles, or polygons) via the `<area>` and `<a>` elements. However, the complete media resource first has to be downloaded to the user agent too [23].

As for the URI-based mechanisms, SVG has a spatial URI mechanism using the `#` that specifies the region of an SVG image to be viewed, but having the same limitations as SMIL. The temporalURI draft specification defines a temporal fragment of multimedia resources using the query parameter `?`, thus creating a new resource which is not desirable as fragments should still have a direct link to their

“parent” resource [22]. An in-depth analysis of # versus ? is discussed in Section 3.3. MPEG-21, on the other hand, specifies a normative syntax to be used in URIs for addressing parts of any resource using the #, but the supported media types are restricted to the MPEG formats only [15]. Furthermore, this specification has a very complex syntax which can be ambiguous. Four schemes –ffp(), offset(), mp(), and mask()– are defined and e.g., both mp() and ffp() can be used to identify tracks. Since our rationale is to find a scheme that is easy enough to get implemented and have a real impact on the web, our impression is that MPEG-21 Part 17 was over-designed even for the most simple use cases, thus preventing its adoption, as there are no real-world applications implementing this specification since it was launched in 2006. YouTube released a tool³ to link to particular time points in videos and to annotate parts of those videos spatio-temporally. It uses the URI fragment marker #, but the entire resource still must be downloaded by the client in order to enable seeking in the media file afterwards. In contrast, the solution advocated by the MFWG is to only send the bytes corresponding to the media fragments requested and still be able to cache them. Finally, for fragment-aware protocols such as RTSP, we observe that this behaviour is possible. While RTSP doesn’t provide an URI scheme to address media fragments, it does allow temporal range requests in the PLAY method, track requests in the SETUP method, and mappings of named dimensions through the SDP service. With the media fragment URI specification, we want to be compatible with solutions already widely deployed as much as possible and enable for HTTP what is already possible with RTSP.

3 Media fragments URIs

3.1 Media resource model

We assume that media fragments are defined for “time-linear” media resources, which are characterised by a single timeline (see also Fig. 1). Such media resources usually include multiple tracks of data all parallel along this uniform timeline. These tracks can contain video, audio, text, images, or any other time-aligned data. Each individual media resource also contains control information in data headers, which may be located at certain positions within the resource, either at the beginning or at the end, or spread throughout the data tracks as headers for those data packets. There is also typically a general header for the complete media resource. To comply with progressive decoding, these different data tracks may be encoded in an interleaved fashion. Normally, all of this is contained within one single container file.

3.2 Requirements

We formally define a number of requirements for the identification and access of media fragments. Based on these requirements, we motivate a number of design choices for processing media fragment URIs.

³http://www.youtube.com/t/annotations_about

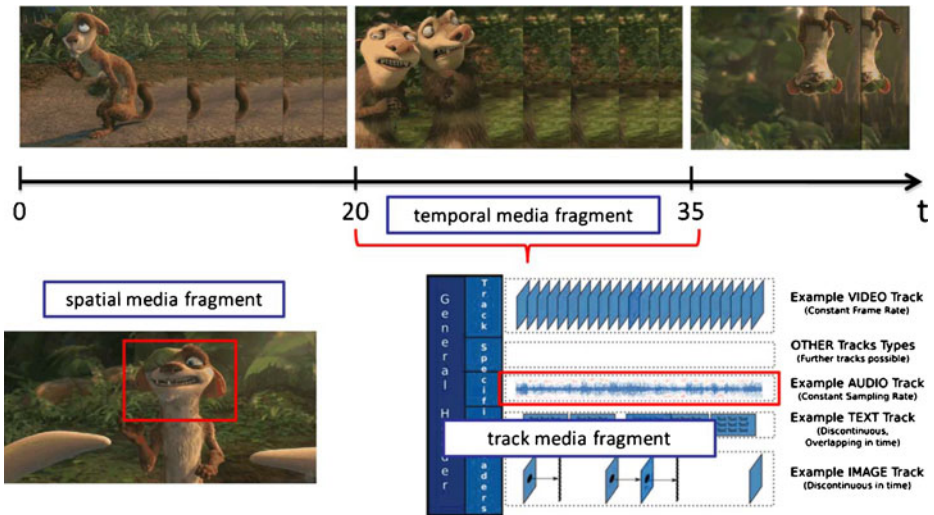


Fig. 1 Example of media fragments and media resource model

- *Independent of media formats.* The media fragments URI specification needs to be independent of underlying media formats such as MP4 [11], Ogg [20] or WebM [28].
- *Fragment axes.* Media fragments need to be accessed along three different axes: temporal, spatial, and track. Additionally, media fragments could be identified through names.
- *Context awareness.* A media fragment must be a secondary resource of its parent resource. This way, the relationship between the media fragment and its parent resource is kept. Moreover, when accessing media fragments, user agents need to be aware of the context of the media fragment (i.e., where the fragment is located within the original parent resource).
- *Low complexity.* The Media Fragment URI specification should be kept as simple as possible. For instance, defining spatial regions is limited to the definition of rectangular regions which should be sufficient for most applications.
- *Minimise Impact on Existing Infrastructure.* Necessary changes to all software components –be it user agents, proxies, or media servers– in the complete media delivery chain should be kept to a bare minimum. Furthermore, the access to media fragments should work as much as possible within the boundaries of existing protocols such as FILE, HTTP(S) and RTSP [24].
- *Fragment by extraction.* Preferably, it should be possible to express media fragments in terms of byte ranges pointing to the parent media resource. This makes the fragments real sub-resources of the “de facto” media resource. Therefore, we consider media segments obtained through a re-encoding process not as media fragments.

3.3 URI fragments vs. URI queries

According to the URI specification [9], URI fragment markers # point at so-called “secondary” resources. Per definition, such a secondary resource may be

“some portion or subset of the primary resource, some view on representations of the primary resource, or some other resource defined or described by those representations”. As such, it makes a viable syntax for Media Fragment URIs. A further consequence of the use of URI fragments is that the media type of the retrieved fragment should be the same as the media type of the primary resource, which means that an URI fragment pointing to a single video frame within a video results in a one-frame video, and not in a still image. To make these URI fragment addressing methods work, only byte-identical segments of a media resource can be considered, as we assume a simple mapping between the media fragment and its elementary byte-aligned building blocks. Where byte-identity cannot be maintained and thus a form of transcoding of the resource is needed, the user agent is not able to resolve the fragmentation by itself and an extra server interaction is required. In this case, URI queries have to be used as they result in a server interaction to deliver a newly created transcoded resource.

The main difference between an URI query and an URI fragment is indeed that an URI query creates a completely new resource having no relationship whatsoever with the resource it is created from, while an URI fragment delivers a secondary resource that relates to the primary resource. As a consequence, URI query created resources cannot be mapped byte-identical to their parent resource (this notion does not even exist), and are thus considered a re-encoded segment. As of the aforementioned requirements *Context awareness* and *Fragment by extraction* described in Section 3.2, the use of URI fragments is preferable over the use of URI queries since byte range requests are an inherent part of HTTP including the caching proxy infrastructure, while providing a query mechanism requires extending the server software. We discuss how to resolve media fragments using the # in Section 4.

In the case of playlists composed of media fragment resources, the use of URI queries (receiving a completely new resource instead of just byte segments from existing resources) could be desirable since it does not have to deal with the inconvenience of the original primary resources - its larger file headers, its longer duration, and its automatic access to the original primary resources. On top of that, URI queries have to take care of an extra caveat of creating a fully valid new resource. This implies typically a reconstruction of the media header to accurately describe the new resource, possibly applying a non-zero start-time or using a different encoding parameter set. Such a resource will then be cached in Web proxies as a different resource to the original primary resource.

3.4 Fragment dimensions

3.4.1 Temporal axis

The most obvious *temporal dimension* denotes a specific time range in the original media, such as “starting at second 10, continuing until second 20”. Temporal clipping is represented by the identifier τ , and specified as an interval with a begin and an end time (or an in-point and an out-point, in video editing terms). If either or both are omitted, the begin time defaults to 0 second and the end time defaults to the end of the entire media resource. The interval is considered half-open: the begin time is part of the interval whereas the end time on the other hand is the first time point that is not part of the interval. The time units that can be used are Normal Play Time (npt),

real-world clock time (clock), and SMPTE timecodes [24, 26]. The time format is specified by name, followed by a colon, with `npt` : being the default. Some examples are:

```
t=npt:10,20      # => results in the time interval [10,20[
t=,20           # => results in the time interval [0,20[
t=smppte:0:02:00, # => results in the time interval [120,end[
```

3.4.2 Spatial axis

The *spatial dimension* denotes a specific spatial rectangle of pixels from the original media resource. The rectangle can either be specified as pixel coordinates or percentages. A rectangular selection is represented by the identifier `xywh`, and the values are specified by an optional format `pixel` : or `percent` : (defaulting to `pixel`) and four comma-separated integers. These integers denote the top left corner coordinate (x,y) of the rectangle, its width and its height. If `percent` is used, x and width should be interpreted as a percentage of the width of the original media, and y and height should be interpreted as a percentage of the original height. Some examples are:

```
xywh=160,120,320,240      # => results in a 320x240 box at
                           x=160 and y=120
xywh=pixel:160,120,320,240 # => results in a 320x240 box at
                           x=160 and y=120
xywh=percent:25,25,50,50  # => results in a 50%x50% box at
                           x=25% and y=25%
```

3.4.3 Track dimension

The *track dimension* denotes one or multiple tracks, such as “the English audio track” from a media container that supports multiple tracks (audio, video, subtitles, etc). Track selection is represented by the identifier `track`, which has a string as a value. Multiple tracks are identified by multiple name/value pairs. Note that the interpretation of such track names depends on the container format of the original media resource as some formats only allow numbers, whereas others allow full names. Some examples are:

```
track=1&track=2          # => results in only extracting
                           track '1' and '2'
track=video               # => results in only extracting
                           track 'video'
track=Kids%20Video       # => results in only extracting
                           track 'Kids Video'
```

3.4.4 Named dimension

The *named dimension* denotes a named section of the original media, such as “chapter 2”. It is in fact a semantic replacement for addressing any range along the aforementioned three axes (temporal, spatial, and track). Name-based selection is represented by the identifier `id`, with again the value being a string. Percent-encoding can be used in the string to include unsafe characters (such as a single

quote). Interpretation of such strings depends on the container format of the original media resource as some formats support named chapters or numbered chapters (leading to temporal clipping), whereas others may support naming of groups of tracks or other objects. As with track selection, determining which names are valid requires knowledge of the original media resource and its media container format.

```
id=1          # => results in only extracting the section
               called '1'
id=chapter-1 # => results in only extracting the section
               called 'chapter-1'
id=My%20Kids # => results in only extracting the section
               called 'My Kids'
```

3.4.5 Combined dimensions

As the temporal, spatial, and track dimensions are logically independent, they can be combined where the outcome is also independent of the order of the dimensions. As such, following fragments should be byte-identical:

```
http://example.com/video.ogv#t=10,20&track=vid&xywh=
  pixel:0,0,320,240
http://example.com/video.ogv#track=vid&xywh=
  0,0,320,240&t=npt:10,20
http://example.com/video.ogv#xywh=0,0,320,240&t=
  smpte:0:00:10,0:00:20&track=vid
```

4 Resolving media fragments with HTTP

We described in the previous section the Media Fragment URI syntax. We present in this section how a Media Fragment URI should be processed using the HTTP protocol. We foresee that the logic of the processing of media fragments URI will be implemented within smart user agents, smart servers and sometimes in a proxy cacheable way. We observe that spatial media fragments are typically interpreted on the user agent side only (i.e., no spatial fragment extraction is performed on server-side) for the following reasons:

- Spatial fragments are typically not expressible in terms of byte ranges. Spatial fragment extraction would thus require transcoding operations resulting in new resources rather than fragments of the original media resource according to the semantics of the URI fragments defined in the corresponding RFC [9].
- The contextual information of extracted spatial fragments is not really usable for visualisation on client-side.

In the remainder of this section, we describe how to resolve media fragments URIs using the HTTP protocol focusing on the temporal and track dimensions.

4.1 User agent mapped byte ranges

As stated in Section 1 (scenario (a)), Steve can now show off his awesome play using his smart phone displaying the specific scene posted on his blog by using the following media fragment URI:

<http://example.com/video.ogv#t=10,20>

Since Steve does not want to blow his entire monthly operator fee with the unnecessary bandwidth cost of downloading the full movie, he uses a smart user agent that is able to interpret the URI, determine that it only relates to a sub-part of a complete media resource, and thus requests only the appropriate data for download and playback.

We assume that Steve's smart phone runs a smart user agent (e.g., an HTML5 compliant browser) that can resolve the temporal fragment identifier of the media fragment URI through an HTTP byte range request. A click on this URI triggers the following chain of events:

1. The user agent checks if a local copy of the requested fragment is available in its buffer, which is not the case.
2. We assume the use of a user agent that knows how time is mapped to byte offsets for this particular Ogg media format.⁴ It just parses the media fragment identifier and maps the fragment to the corresponding byte range(s).
3. The user agent sets up the decoding pipeline for the media resource at <http://example.com/video.ogv> by just downloading the first couple of bytes of the file that correspond to the resource's header information.
4. The MIME-type of the resource requested is confirmed. The user agent can use the header information to resolve the fragment byte ranges. Based on the calculated time-to-byte mapping and the extracted information from the resource's header of where to find which byte, the user agent sends one or more HTTP requests using the *HTTP Range request* header (see Fig. 2) for the relevant bytes of the fragment to the server. In this particular case, an HTTP 1.1 compliant Web server will be able to serve the media fragment.
5. The server extracts the bytes corresponding to the requested range and responds with an *HTTP 206 Partial Content response* containing the requested bytes.
6. Finally, the user agent receives these byte ranges and is able to decode and start playing back the initially requested media fragment.

In this scenario, media fragments can be served by traditional HTTP Web servers that support byte range requests. However, a smart user agent is necessary to parse the media fragment URI. Furthermore, it requires knowledge about the syntax and semantics of various media codecs formats. More profound client side examples (fragment requests which are already buffered and fragment requests of a changed resource) can be found in the Media Fragments specification [19]. It is expected that this recipe will be supported for this kind of media fragment URI and the temporal dimension with the HTML5 media elements.

⁴This is possible by using Ogg Index, http://wiki.xiph.org/Ogg_Index.

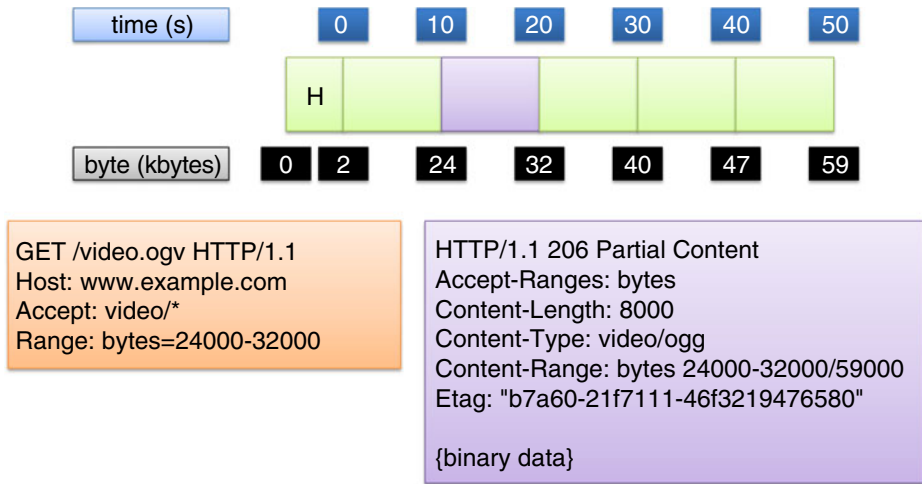


Fig. 2 User agent mapped byte ranges

4.2 Server mapped byte ranges

We assume now that the user agent is able to parse and understand the media fragment URI syntax, but is unable to perform by itself the byte range mapping for a given request. In this case, the user agent needs some help from the media server to perform this mapping and deliver the appropriate bytes.

4.2.1 Server mapped byte ranges including header information

In a first case, the server has to help the client to setup the initial decoding pipeline (i.e., there is no header information from the resource on client side yet). When Steve starts to play this “infamous” 10 s of video, the following events occur:

1. The user agent parses the media fragment identifier and creates an HTTP Range request expressed in a different unit than bytes, e.g., a time unit expressed in seconds. Furthermore, it extends this header with the keyword `include-setup` (see Fig. 3) in order to let the server know that it also requires the initial header of the resource to initiate the decoding pipeline.
2. The server extracts the header information of the media resource as requested with `include-setup`.
3. The server, which also understands this time unit, resolves the HTTP Range request and performs the mapping between the media fragment time unit and byte ranges, and extracts the corresponding bytes.
4. Once the mapping is done, the server then wraps both the header information and the bytes requested in a multi-part HTTP 206 Partial Content response. As depicted in Fig. 3, the first part contains the header data needed for setting up the decoding pipeline, whereas subsequent parts contain the requested bytes of the needed fragment. Note that a new response header, named `Content-Range-Mapping`, is introduced to provide the *exact* mapping of the retrieved byte range corresponding to the original `Content-Range` request

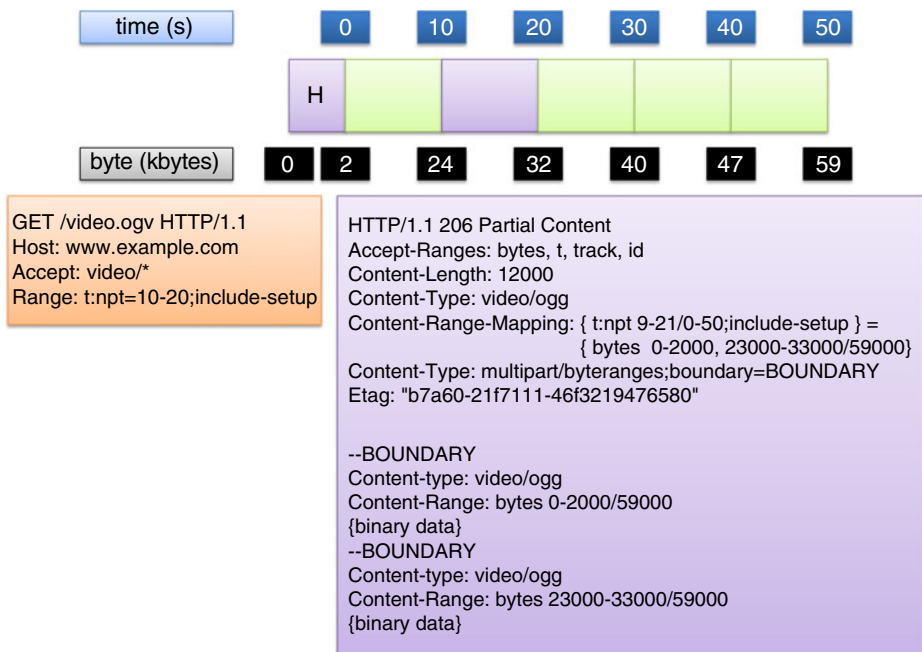


Fig. 3 Server mapped byte ranges including header information

expressed with a time-unit (and not in bytes). The decoder might need extra data, before the beginning (hence npt 9), and/or after the end (hence npt 21) of the requested sequence (e.g., npt 10-20), since this initial requested period might not correlate to a random access unit of the clip to start/end with. In analogy with the Content-Range header, the Content-Range-Mapping header also adds the instance-length after the slash-character /, thus providing context information regarding the parent resource in case the HTTP Range request contained a temporal dimension. More specifically, the header contains the start and end time of the parent resource, so the user agent is able to understand and visualise the temporal context of the media fragment.

5. Finally, the user agent receives the multi-part byte ranges and is able to setup the decoding pipeline (using the header information), decodes, and starts playing back the media fragment requested.

A plugin named *Media Fragments 1.0* developed for the Firefox browser has successfully implemented this recipe and is now available in the Mozilla add-ons library.

4.2.2 Server mapped byte ranges with initialized client

On the other hand, if the server can assume that the client already initiated the decoding pipeline, i.e., it knows about the header of the requested media resource, then the following events occur:

1. The user agent parses the media fragment identifier and creates an HTTP Range request expressed in a different unit than bytes, e.g., a time unit expressed in Normal Playing Time (npt) seconds (Fig. 4).
2. The server, which understands this time unit, resolves the HTTP Range request and performs the mapping between the media fragment time unit and byte ranges.
3. Once the mapping is done, the server extracts the proper bytes and wraps them within a HTTP 206 Partial Content response. As displayed in Fig. 4, the response contains the additional Content-Range-Mapping header describing the content range in terms of time (ending with the total duration of the complete resource). Where multiple byte ranges are required to satisfy the HTTP Range request, these are transmitted as a multipart message body (with media type ‘multipart/byteranges’), as can be seen in Fig. 2.
4. Finally, the user agent receives these byte ranges and is able to decode and start playing back the media fragment requested.

If the server does not understand the HTTP Range request –which means it does not support media fragments–, it must ignore that header field –per the current HTTP specification– and deliver the complete resource. This also means we can combine both byte range and fragment range headers in one request, since the server will only react to the HTTP Range header it understands.

4.3 Proxy cacheable server mapped byte ranges

To prevent wasting unnecessary bandwidth and to maximise throughput speeds, the current internet architecture heavily relies on caching Web proxies. In the case of media fragment URIs that are sometimes resolved by servers as described in the previous section, existing Web proxies have no means of caching these requests as they only understand byte ranges.

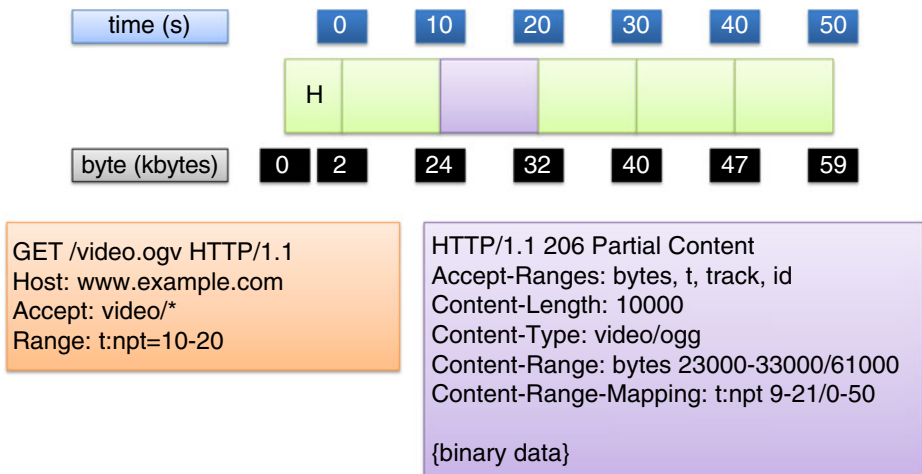


Fig. 4 Server mapped byte ranges with initialised client

In order to make use of the Web proxies, the user agent should therefore only ask for byte ranges. To be able to do so, we foresee an extra communication between the server and the user agent. In a first step, the server will just redirect the original request giving extra hints of the byte ranges to request corresponding to the media fragment instead of replying with the actual data. In a second step, the user agent will re-issue another range request, this time expressed in bytes which will therefore be cacheable for current Web proxies.

In the scenario (b) described in the Section 1, Sarah wants to finalise the new music clip of “The Editors”. Therefore, she scrolls through the video footage she has gathered in order to find the appropriate final scene. Since she is only interested in finding the right frames to start editing the high quality footage, she would like to watch only 10 seconds of the video track in low-resolution of those gathered clips, which can be translated into the following media fragment URI:

<http://example.com/video.ogv#t=10,20&track=video>

The chain of events for getting 10 seconds of the video track of this footage is as follows:

1. The user agent parses the media fragment identifier and creates an HTTP Range request expressing the need for 10 s of the video track and further requests to retrieve just the mapping to byte ranges, hence the extra `Accept-Range-Redirect` HTTP header (Fig. 5), which signals to the server that only a redirect to the correct byte ranges is necessary and the result should be delivered in the HTTP Range-Redirect header.

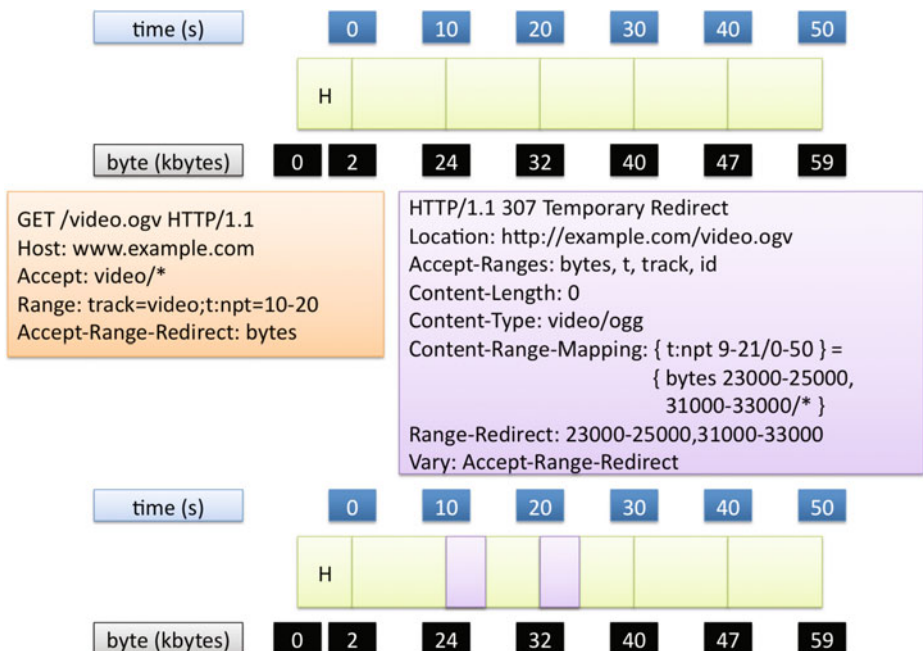


Fig. 5 Proxy cacheable server mapped byte ranges

2. The server resolves the HTTP Range request, knows it only has to look for the correct byte ranges of the video track, and performs the mapping between the media fragment time unit and byte ranges.
3. Afterwards the server sends back an empty HTTP 307 Temporary Redirect that refers the user agent to the right byte range(s) for the previous requested correct time range. Note that for the sake of simplicity only 2 byte ranges are sent back, whereas interleaving normally means a lot more data chunks for one track.
4. After this initial indirection, the steps to follow by the user agent are the same as those we discussed earlier in Section 4.1.

More advanced server side examples, e.g., server triggered redirects, can again be found in the Media Fragments specification [19].

5 Influence of media formats on fragment extraction

The extraction of media fragments is dependent on the underlying media format. Moreover, media formats introduce restrictions regarding possibilities to extract media fragments along different axes. For instance, very few coding formats support the extraction of spatial fragments without the need to re-encode the media resource. As such, depending on the fragment axis, there are different requirements regarding the extraction of media fragments (i.e., obtaining media fragments without transcoding operations).

- *Temporal*: random access points need to occur in a regular way. Random access refers to the ability of a decoder to start decoding at a point in a media resource other than at the beginning and to recover an exact or approximate representation of the decoded resource [5]. In case of video, random access points typically correspond to intra-coded pictures.
- *Spatial*: independently coded spatial regions are required. More specifically, (mostly rectangular) regions of the picture need to be coded independently from each other. Typically, Regions Of Interest (ROIs) and a background are coded independently, which results in the possibility to extract these ROIs. Another, more advanced approach is “interactive” ROI. Interactive ROI extraction can be used in applications in which the ROI cannot be defined during the encoding phase of the video sequence [2]. In order to support interactive (rectangular) ROI scalability, a video frame has to be divided into different tiles that can be selected on an individual basis. Each tile has to be coded such that the tiles can be decoded independently of other tiles. This way, the tiles belonging to a certain ROI can be selected on-the-fly during the extraction process. Scalable Video Coding [25] (SVC) is an example of a video coding format supporting interactive ROI coding.
- *Track*: the way tracks are extracted from a media resource is dependent on the container format. In contrast to temporal and spatial fragment extraction, tracks are not “encoded” within a container format and can thus always be extracted without low-level transcoding operations. Based on the headers of the container format, it is possible to locate the proper chunks corresponding to the desired track.

- *Named*: whether named fragments are supported or not depends again on the container format. Typically, named fragments are supported through container formats by using other metadata formats describing these named fragments. An instance of such a metadata format can, for example, correspond to a separate track or can be included in the headers of the container format. Examples of such metadata formats are MPEG-4 TimedText [14] and Rich Open multitrack media Exposition⁵ (ROE).

Hereafter, we provide more information regarding the extraction of media fragments applied to three popular media format combinations:

- H.264/AVC-encoded video [16] and AAC-encoded audio [13] packed in an MP4 container;
- Theora-encoded video [32] and Vorbis-encoded audio [33] packed in an Ogg container.
- VP8-encoded video [27] and Vorbis-encoded audio packed in a WebM container.

For each combination, we discuss how media fragments along different axes can be extracted, based on the above defined requirements.

5.1 H.264/AVC and AAC in MP4

MPEG-4 Part 14 or MP4 file format, formally known as ISO/IEC 14496-14:2003 [11], is a multimedia container format standard specified as a part of MPEG-4. It is able to store digital video and digital audio streams as well as other data such as subtitles and still images. The MP4 file format is built on top of the ISO Base Media File Format [12] (known as MPEG-4 part 12), which is in turn based on Apple's QuickTime file format.⁶ H.264/AVC [16] is the latest block-oriented motion-compensation-based codec standard developed by the Joint Video Team⁷ (JVT) and is used in applications such as Blu-ray Disc, YouTube and the iTunes Store. Advanced Audio Coding (AAC) is a standardised, lossy compression and encoding scheme for digital audio and was designed to be the successor of the popular MP3 format.⁸ AAC has been standardised by ISO and IEC, as part of the MPEG-2⁹ and MPEG-4¹⁰ specifications and is the standard audio format for Apple's iPhone, Sony's PlayStation 3, Android based phones, etc.

Temporal fragment extraction is both supported by H.264/AVC and AAC. For H.264/AVC, random access points correspond to Instantaneous Decoding Refresh (IDR) frames, which are intra-coded frames that are not referenced by any frames outside the current GOP. Thus, a temporal H.264/AVC fragment always starts with an IDR frame. An AAC audio stream consists of a stream of audio frames, each containing typically 1,024 audio samples. Each audio frame corresponds to

⁵<http://wiki.xiph.org/ROE>

⁶<http://www.quicktime.com>

⁷<http://www.itu.int/ITU-T/studygroups/com16/jvt/>

⁸<http://www.iis.fraunhofer.de/EN/bf/amm/products/mp3/index.jsp>

⁹<http://mpeg.chiariglione.org/standards/mpeg-2/mpeg-2.htm>

¹⁰<http://mpeg.chiariglione.org/standards/mpeg-4/mpeg-4.htm>

a random access point, resulting in a fine temporal granularity. The MP4 format provides support for headers containing information regarding random access points of the contained media streams. Hence, by interpreting the MP4 header, temporal fragments can be extracted on condition that random video access points occur in a regular way.

Spatial fragments extraction is only supported by H.264/AVC in a limited way. More specifically, ROI coding can be realised in H.264/AVC using the Flexible Macroblock Ordering (FMO) tool. By using H.264/AVC FMO type 2, rectangular regions can be coded independently from each other (i.e., each region corresponds to a slice group). Extracting a particular ROI corresponds to the detection of coded P and B slices located in non-ROI slice groups and the substitution of these slices with place-holder slices [17].

Tracks are described in the headers of the MP4 format. More specifically, the track box contains general information about the track (e.g., name, creation time, width, height, etc.) as well as references to the location of the bytes representing the track (through the sample table box). Hence, by interpreting these boxes, track fragments can be extracted, based on the names provided by the MP4 container.

Named fragments for MP4 can, for instance, be obtained using MPEG-4 Timed Text, which is the text-based subtitle format for MPEG-4. It was developed in response to the need for a generic method for coding of text as one of the multimedia components within audiovisual presentations. As such, it can be used to give names to temporal fragments.

5.2 Theora and Vorbis in Ogg

The Ogg container format is a multimedia container format and the native file and stream format for the Xiph.org multimedia codecs [20]. It is an open format, free for anyone to use, and is able to encapsulate compressed audio and video streams. Ogg is a stream-oriented container, meaning it can be written and read in one pass, making it a natural fit for internet streaming and use in processing pipelines. Note that this stream orientation is the major design difference over other file-based container formats such as MP4. Theora is a video codec whose bitstream format was frozen in July 2004 by Xiph.org. The standard codec is frequently found on the Web and used by large sites complying with the Wikimedia Commons like Wikipedia.¹¹ Vorbis, also developed by the Xiph.Org Foundation, is an audio format specification for lossy audio compression and has proven popular among supporters of free software. Both Vorbis and Theora contain solely royalty-free and open technology.

Similar to H.264/AVC, random access points in Theora streams correspond to I-frames. Random access in Vorbis streams is obtained through audio frames, similar to AAC (note that the number of samples is variable in case of Vorbis). Thus, temporal fragment extraction is both supported by Theora and Vorbis.

Spatial fragment extraction is not supported by Theora. More specifically, it does not provide support to encode spatial regions independent of each other.

An Ogg container is a contiguous stream of sequential Ogg pages. Ogg allows for separate tracks to be mixed at page granularity in an Ogg container. Tracks are

¹¹<http://www.wikipedia.org>

identified by a unique serial number, which is located in each Ogg page header. Further, to link Ogg tracks to track names, again the ROE format can be used. ROE is a metadata format for describing the relationships between tracks of media in a stream. ROE descriptions are represented in an Ogg container by means of an Ogg Skeleton track. This way, track fragment extraction is supported by Ogg through fields containing the track names stored in Ogg Skeleton.

Named fragment extraction for temporal and track fragments in Ogg can be obtained by combining ROE, Skeleton, and Continuous Media Markup Language¹² (CMML). CMML allows a time-continuously sampled data file to be structured by dividing it into temporal sections (so-called clips) and provides these clips with some additional information. Thus, named fragment extraction can be done based on the CMML and ROE description of an Ogg container.

5.3 VP8 and Vorbis in WebM

The WebM container format, which is introduced by Google, is an open and royalty-free multimedia container format designed for the Web. WebM is based on the Matroska media container format. It uses Vorbis for the representation of audio streams while VP8 is used as video codec. VP8 [27], released by Google and originally created by On2 Technologies, is the latest open source video codec. Compression-wise, VP8 generally performs better than Theora.

The constraints implied by WebM for extracting temporal media fragments are very similar to the ones we discussed for Theora and H.264/AVC. More specifically, random access points in VP8 streams correspond to I-frames. Temporal fragment extraction is thus supported by VP8. Also, similar to Theora, spatial fragment extraction is not supported by VP8.

Tracks within WebM are identified by a track number and/or a track identifier. The track identifier can be used within a Media Fragment URI using the track axis, while the track number is used internally to indicate for each frame (audio or video) to which track it belongs. This way, we can address a track in a WebM container and also extract the track. Further, named fragments within WebM containers can be obtained by defining chapters. A chapter in WebM is a combination of a time segment and one or more tracks. The name of the chapter could then be used as an identifier for the named fragment. Note that WebM also provides support for adding tags to chapters. However, these tags cannot always be used as identifier for a named fragment because they could result in multiple time segments (e.g., two chapters having the same tag). Addressing multiple time segments within one Media Fragment URI is not supported, as discussed in Section 3.4.1.

6 Open issues and future work

Currently, it is unclear for all the specified media fragment axes, how their defined media fragments should be rendered to the end-user in a meaningful way. Temporal fragments could be highlighted on a timing bar whereas spatial fragments could be emphasised by means of bounding boxes or they could be played back in colour

¹²<http://wiki.xiph.org/CMML>

while the background is played back in grey-scale. Others have suggested that spatial media fragments should result in cropping the images. Finally, track selection could be done via dropdown boxes or buttons. Whether the media fragment URIs should be hidden from the end-user or not is an application implementation issue.

There is currently no standardised way for user agents to discover the available named and track fragments. One could use ROE, which makes it possible to express the track composition of a media resource, in addition to naming these tracks and extra metadata info on language, role and content-type which could further help selecting the right tracks. Another candidate to use is the *Media Multitrack API*¹³ from the HTML5 Working Group. This is a JavaScript API for HTML5 media elements that allows content authors to determine which tracks are available from a media resource. Not only does it expose the tracks that a media resource contains, but it also specifies the type of data that is encapsulated within the resource (e.g., audio/vorbis, text/srt, video/theora), the role this data plays (e.g., audio description, caption, sign language), and the actual language (e.g., RFC3066¹⁴ language code).

With such media fragments addressing schemes available, there is still a need to hook up the addressing with the resource. For the temporal and the spatial dimension, resolving the addressing into actual byte ranges is relatively obvious across any media type. However, track addressing and named addressing need to be resolved too. Track addressing will become easier when we solve the above stated requirement of exposing the track structure of a media resource. The name definition, on the other hand, will require association of an id or name with temporal offsets, spatial areas, or tracks.

Finally, hyperlinking out of media resources is something that is not generally supported at this stage. Certainly, some types of media resources such as QuickTime, Flash,¹⁵ MPEG-4, and Ogg support the definition of tracks that can contain HTML marked-up text and thus can also contain hyperlinks. It seems to be clear that hyperlinks out of media files will come from some type of textual track. On 6 May 2010, the WHATWG version of the HTML5 draft specification introduced the WebSRT format (Web Subtitle Resource Tracks), a format intended for marking up timed track resources. WebSRT will have a means of addressing segments of a media resource by name and also a means to include text with hyperlinks, though the exact scope of WebSRT and its inclusion as part of HTML5 is still under debate.

7 Conclusion

In this paper, we presented the rationale for a Media Fragment URIs specification to make video a “first-class citizen” on the web and pinpointed the drawbacks/inconsistencies of existing solutions. We outlined the boundaries, requirements and semantics of a Media Fragment URI. We clearly defined the syntax of

¹³http://www.w3.org/WAI/PF/HTML/wiki/Media_MultitrackAPI

¹⁴<http://www.ietf.org/rfc/rfc3066.txt>

¹⁵<http://www.flash.com>

Media Fragments over the different possible fragment dimensions (i.e., temporal, spatial, track, and named) and showed how the defined Media Fragments can be resolved stepwise using the HTTP protocol. We then identified the influence of current media formats on such Media Fragments extraction.

We already have a HTTP implementation for the W3C Media Fragments 1.0 specification [30] in order to verify all the test cases defined by the working group. Furthermore, several browser vendors are on their way to implementing an HTTP implementation themselves. In the near future, we also foresee reference implementations for the RTSP, RTMP,¹⁶ and File protocols.

In the meantime this specification really opens up time-related media to the internet crowd and makes time-related annotation [18, 19] feasible for hyperlinking into the media, thus providing the necessary support for this already omnipresent “third dimension” into the internet.

Acknowledgements The authors would like to thank the other W3C Media Fragments’ participants: Eric Carlson (Apple, Inc.), Thierry Michel (W3C/ERCIM), Guillaume Olivrin (Meraka Institute, South Africa), Soohong Daniel Park (Samsung Electronics Co., Ltd.), David Singer (Apple, Inc.), Philip Jägenstedt (Opera Software) for their willingness to discuss the definition of media fragments and more generally their adequacy within the Web architecture. The research activities that have been described in this paper were partially funded by W3C/ERCIM, Ghent University, Interdisciplinary Institute for Broadband Technology (IBBT), the Institute for the Promotion of Innovation by Science and Technology in Flanders (IWT), the Fund for Scientific Research-Flanders (FWO-Flanders), the French Ministry of Industry (Innovative Web call) under contract 09.2.93.0966, EURECOM, CWI, Kyoto University, National University of Ireland, Mozilla Corporation, and the European Union within the 7th framework program (FP7/2007-2013) under agreement ICT-2007-214793.

References

1. Bulterman D, Grassel G, Jansen J, Koivisto A, Layaïda N, Michel T, Mullender S, Zucker D (eds) (2005) Synchronized multimedia integration language (SMIL 2.1). W3C recommendation. World Wide Web Consortium. Available at <http://www.w3.org/TR/SMIL2/>
2. De Schrijver D, De Neve W, Van Deursen D, De Bruyne S, Van de Walle R (2006) Exploitation of interactive region of interest scalability in scalable video coding by using an XML-driven adaptation framework. In: Proceedings of the 2nd international conference on automated production of cross media content for multi-channel distribution, pp 223–231. Leeds, UK
3. EBU/ETSI (2008) ETSI TS 102 822-3-1: TV-anytime—part 3: metadata. Available at <http://www.tv-anytime.org/>
4. Ferraiolo J, Jackson D (eds) (2009) Scalable vector graphics (SVG) 1.1 specification. W3C recommendation. World Wide Web Consortium. Available at <http://www.w3.org/TR/SVG11/>
5. Hannuksela MM, Wang YK, Gabbouj M (2004) Isolated regions in video coding. *IEEE Trans Multimedia* 6:259–267
6. Hausenblas M, Troncy R, Burger T, Raimond Y (2009) Interlinking multimedia: how to apply linked data principles to multimedia fragments. In: Proceedings of the 2nd workshop on linked data on the Web (LDOW’09). Madrid, Spain
7. Hickson I, Hyatt D (eds) (2010) HTML5—a vocabulary and associated APIs for HTML and XHTML. W3C working draft. World Wide Web Consortium. Available at <http://www.w3.org/TR/html5/>

¹⁶RTMP: Adobe’s Real Time Messaging Protocol available at <http://www.adobe.com/devnet/rtmp/>.

8. Internet Engineering Task Force (199) RFC 2616: hypertext transfer protocol—HTTP/1.1. Available at <http://www.ietf.org/rfc/rfc2616.txt>
9. Internet Engineering Task Force (2005) RFC 3986: uniform resource identifier (URI)—generic syntax. Available at <http://tools.ietf.org/html/rfc3986>
10. ISO/IEC (2002) Information technology—multimedia content description interface—part 1: systems. ISO/IEC 15938-1:2002
11. ISO/IEC (2003) Information technology—coding of audio, picture, multimedia and hypermedia information—part 14: MP4 file format. ISO/IEC 14496-14:2003
12. ISO/IEC (2005) Information technology—coding of audio-visual objects—part 12: ISO base media file format. ISO/IEC 14496-12:2005
13. ISO/IEC (2005) Information technology—coding of audio-visual objects—part 3: audio. ISO/IEC 14496-3:2005
14. ISO/IEC (2006) Information technology—coding of audio-visual objects—part 17: streaming text format. ISO/IEC 14496-17:2006
15. ISO/IEC (2006) Information technology—MPEG-21, part 17: fragment identification of MPEG resources. ISO/IEC 21000-17:2006
16. ITU-T and ISO/IEC (2003) Advanced video coding for generic audiovisual services. ITU-T Rec. H.264 and ISO/IEC 14496-10 AVC
17. Lambert P, De Schrijver D, Van Deursen D, De Neve W, Dhondt Y, Van de Walle R (2006) A real-time content adaptation framework for exploiting ROI scalability in H.264/AVC. In: Blanc-Talon J, Philips W, Popescu D, Scheunders P (eds) Advanced concepts for intelligent vision systems. Lecture notes in computer science, vol 4179. Springer, Berlin, pp 442–453
18. Mannens E, Troncy R (eds) (2009) Use cases and requirements for media fragments. W3C working draft. World Wide Web Consortium. Available at <http://www.w3.org/TR/media-frags-reqs/>
19. Mannens E, Troncy R, Pfeiffer S, Van Deursen D (eds) (2010) Media fragments URI 1.0. W3C working draft. World Wide Web Consortium. Available at <http://www.w3.org/TR/media-frags>
20. Pfeiffer S (2003) RFC 3533: The Ogg encapsulation format version 0. Available at <http://www.ietf.org/rfc/rfc3533.txt>
21. Pfeiffer S (2007) Architecture of a video web—experience with annodex. In: Proceedings of the W3C video on the Web workshop. Brussels, Belgium
22. Pfeiffer S, Parker C, Pang A (2005) Internet draft: specifying time intervals in URI queries and fragments of time-based Web resources. Available at <http://annodex.net/TR/draft-pfeiffer-temporal-fragments-03.html>
23. Raggett D, Le Hors A, Jacobs I (eds) (1999) HyperText markup language (HTML) 4.01—specification. W3C recommendation. World Wide Web Consortium. Available at <http://www.w3.org/TR/html4/>
24. Schulzrinne H, Rao A, Lanphier R (1998) RFC 2326: real time streaming protocol. Available at <http://www.ietf.org/rfc/rfc2326.txt>
25. Schwarz H, Marpe D, Wiegand T (2007) Overview of the scalable video coding extension of the H.264/AVC standard. IEEE Trans Circuits Syst Video Technol 17(9):1103–1120
26. Society of Motion Picture and Television Engineers (2004) SMPTE RP 136: time and control codes for 24, 25 or 30 frame-per-second motion—picture systems. Available at <http://www.smpte.org/standards>
27. The WebM Project (2010) VP8 data format and decoding guide. Available at <http://www.webmproject.org/code/specs/>
28. The WebM Project (2010) WebM multimedia container guidelines. Available at <http://www.webmproject.org/code/specs/>
29. Troncy R, Hardman L, Van Ossenbruggen J, Hausenblas M (2007) Identifying spatial and temporal media fragments on the Web. In: Proceedings of the W3C video on the Web workshop. Brussels, Belgium
30. Van Deursen D, Troncy R, Mannens E, Pfeiffer S, Lafon Y, Van de Walle R (2010) Implementing the media fragments URI specification. In: Proceedings of the 19th international World Wide Web conference (WWW), pp 1361–1363. Raleigh, USA
31. W3C Media Fragments Working Group (2010) Media fragments working group. Available at <http://www.w3.org/2008/WebVideo/Fragments/>
32. Xiph.org Foundation (2009) Theora specification. Available at <http://theora.org/doc/Theora.pdf>
33. Xiph.org Foundation (2010) Vorbis I specification. Available at http://xiph.org/vorbis/doc/Vorbis_I_spec.pdf



Erik Mannens received his MSc degree in Electro-Mechanical Engineering (1992) at KAHO Ghent and his MSc degree in Computer Science (1995) at KU Leuven University. Before joining IBBT-MMLab in 2005 as a project manager, he was a software engineering consultant and Java architect for over a decade. His major expertise is centered around metadata modeling, semantic Web technologies, broadcasting workflows, iDTV and Web development in general. He is involved in several projects as senior researcher and is finishing up his PhD on Semantic News Production. He is co-chair of W3C's Media Fragments Working Group and actively participating in other W3C's semantic Web standardization activities. On all of these subjects he has published several papers and book chapters. He is also an active member of the technical committees of ACM MultiMedia, IEEE CCNC, IEEE MMTC, SAMT and MARESO. His full bio can be obtained from <http://www.mmlab.be/emannens>.



Davy Van Deursen received his MSc and PhD degrees in Computer Science from Ghent University, Belgium in 2005 and 2009, respectively. He is currently a post-doctoral researcher at Ghent University—IBBT (Ghent, Belgium), where he is working for the Multimedia Lab research group (<http://multimedialab.elis.ugent.be>). Between 2007 and 2009, he was involved in Siruna, a spin-off of IBBT and Ghent University which offered a platform serving as a thin client mobile application gateway. Since 2008, he is also actively participating within the W3C Media Fragments Working Group, whose mission is to specify a URI scheme for the addressing of media fragments. His research interests and areas of publication include video coding, media content adaptation and delivery, multichannel publishing, media content presentation, Semantic Web technologies, advanced reasoning techniques, e-learning, and standardization.



Raphaël Troncy is an assistant professor at EURECOM since 2009. He obtained with honors his Master's thesis in Computer Science at the University Joseph Fourier of Grenoble (France), after one year spent in the University of Montreal (Canada). He benefited from a PhD fellowship at the National Audio-Visual Institute (INA) of Paris where he received with honors his PhD from the University of Grenoble (INRIA/INA) in 2004. He selected as an ERCIM Post-Doctorate Research Associate (2004–2006) where he visited the National Research Council (CNR) in Pisa (Italy) and the National Research Institute for Mathematics and Computer Science (CWI) in Amsterdam (The Netherlands). He was a senior researcher for CWI from 2006 till 2009 where he actively participated in the EU K-Space Network of Excellence.

He is now involved in a number of National (ACAV, DataLift) and European (FP7 Petamedia, AAL ALIAS) projects where semantic web and multimedia analysis technologies are used together. Raphaël Troncy is co-chair of the W3C Incubator Group on Multimedia Semantics and the W3C Media Fragments Working Group and contributes to the W3C Media Annotations Working Group. He is an expert in audio-visual metadata and in combining existing metadata standards (such as MPEG-7) with current Semantic Web technologies. He works closely with the IPTC standardization body on the relationship between the NewsML language family and Semantic Web technologies.



Silvia Pfeiffer has worked on novel media technology for more than 15 years and is an internationally renowned expert in new Web video standards.

After finishing her double degree in Computer Science and Business Management at Mannheim University, Germany, she worked as a consultant at Accenture (then: Andersen Consulting) for six months. She then returned to Mannheim University to completed her PhD on audio-visual content analysis. Subsequently, she spent seven years at the CSIRO, amongst other activities developing new video technology for the Web in the “Annodex” project.

She continues to promote and pursue open media technology, being the main organiser for the annual Foundations of Open Media Software workshop, being the president of the Annodex Association, and a member of the Xiph.org community.

She is now an invited expert on four W3C video-related working groups. She is making contributions to the new audio and video elements in HTML5, to media annotation standards, to media fragment addressing via URIs, and to video accessibility technology for hearing and seeing-impaired people (captions, audio annotations etc).

In 2007, she co-founded Australian video company Vquence. Vquence offers consulting and technology services for Web 2.0 video. This includes strategy planning, social video channel setup, audience measurement, reporting, and recommendation for strategy improvement. It also includes services to create a strategic integration of video into a working online marketing and communication strategy.



Conrad Parker is a PhD candidate at Kyoto University, worked on the Annodex project at CSIRO Australia, and studied Mathematics and Computer Engineering at the University of New South Wales. He is a keen contributor to open source multimedia software including the Sweep sound editor and various Ogg utilities. He has worked on Ogg video support in Firefox; created the oggz-chop media server used by the Internet Archive; and developed a video streaming server and Linux hardware MPEG codec interfaces for a Japanese semiconductor manufacturer. Hobbies include computer music and Japanese language.



Yves Lafon studied Mathematics and Computer Science at ENSEEIHT in Toulouse, France, and at Ecole Polytechnique de Montreal in Montreal, Canada. His field of study was signal recognition and processing. He discovered Internet Relay Chat and the Web in Montreal in 1993 and has been making robots and games for both. He joined the W3C in October 1995 to work on W3C's

experimental browser, Arena. Then he worked on Jigsaw, W3C's Java-based server, on HTTP/1.1 and started the work on SOAP 1.2. Yves is now the Web Services Activity leader, HTTPbis editor and the TAG Team Contact.



Jack Jansen has finally decided to embark on a PhD track in the field of multimedia, after 30 years of working as a research programmer in distributed operating systems (Amoeba), programming languages (Python), augmented reality (Personal Space Station) and multimedia (CMIF, SMIL). His current work is in multimedia protocols and languages, and in how multimedia applications can be better integrated in their environment. Social aspects of multimedia, including privacy issues and accessibility also have his interest.



Michael Hausenblas is a postdoctoral researcher at DERI, National University of Ireland, Galway and the co-ordinator of DERI's Linked Data Research Centre. At DERI, Michael is managing the European project iMP (realising the Virtual Film Factory) and is acting as the coordinator of the European Support Action LATC (in the Linked Data domain). He has published over 40 papers at conferences, workshops and in journals and contributed to the Web of Data in the past couple of years in various domains, including the RDFa standardisation in W3C, in the W3C Media Fragments Working Group and as a co-chair the W3C RDB2DF Working Group.



Rik Van de Walle received his MSc and PhD degrees in Engineering from Ghent University, Belgium in 1994 and 1998, respectively. After a visiting scholarship at the University of Arizona (Tucson, USA), he returned to Ghent University. In 2001 he became a professor at the Department of Electronics and Information Systems (Ghent University—IMEC, Belgium) and founded the Multimedia Lab. Rik Van de Walle has been/is editor of the following MPEG specifications: MPEG-21 Digital Item Declaration Language; MPEG-21 Digital Item Processing; MPEG-21 Digital Item Processing—Technologies under Consideration; and MPEG-21 Reference Software. Rik Van de Walle has been involved in the organization of and/or review of papers for several international conferences (e.g., IEEE ICME, WIAMIS, ISAS-SCI, ACIVS, Mirage, EUROMEDIA Mediatec). His current research interests include multimedia content delivery, presentation and archiving, coding and description of multimedia data, content adaptation, interactive (mobile) multimedia applications, interactive digital TV. Multimedia Lab is one of the partners of the Interdisciplinary Institute for Broadband Technology (IBBT), which was founded by the Flemish Government in 2004.