Deploying Multimedia Metadata in Cultural Heritage on the Semantic Web

Michael Hausenblas, Werner Bailer, and Harald Mayer *

Institute of Information Systems & Information Management,
JOANNEUM RESEARCH Forschungsges. mbH,
Steyrergasse 17, 8010 Graz, Austria
{firstname.lastname}@joanneum.at

Abstract. More and more archives and libraries start publishing parts of their collection on the Web. While many problems have been solved, one critical shortcoming has not yet been addressed: The effective and efficient deployment of multimedia metadata (M3) along with the content. There exists a vast array of existing multimedia metadata formats that have been used for years in diverse applications. However, when one is after using these formats in the context of the Semantic Web, the options are limited. What is basically missing is a framework that allows hooking existing multimedia metadata into the Semantic Web.

1 Introduction

Managing and using multimedia metadata (M3) to facilitate access to cultural objects has always been of particular importance for memory institutions. With the advances of the available technologies, such as digitisation, it has become possible to provide access to catalogues, and often the digitised objects themselves are made available on the Web. These Web-based applications provide for several kinds of browsing and searching facilities.

In recent projects we observe growing interest by cultural heritage institutions to make part of their collections available on the Web. This may be realised through an on-line exhibition of a single museum (e.g., the "Museum im alten Zeughaus"[1]), or a joint effort of several museums of a region (as implemented in the DOMINICO project[2]), or audio-visual archive content, such as the BIRTH Television Archive[3]. We further gathered experiences with the publication platform developed as part of the PrestoSpace project[4], an EU 6th Framework Programme IST Integrated Project.

* The authors wish to express their gratitude to Steiermärkisches Landesarchiv for providing the scanned newspaper image.

1 http://www.stadtbadradkersburg.at/deutsch/stadt/diestadt/Museum/Museumsaustellungen.php

2 http://dominico.joanneum.at

3 http://www.birth-of-tv.org

4 http://www.prestospace.org

The frequently used approach is to provide full text search, which works well with text based objects. It enables conventional Web search engines to offer retrieval results on a global scale. However, with multimedia assets, such as still images, audio recordings and video material this mechanism falls short—they only become searchable when published together with metadata [16].

The actual deployment of existing M3 formats to a Semantic Web agent is a critical success factor for a Semantic Web application dealing with multimedia assets. This is pretty much the same as communication providers nowadays struggle dealing with the so called *last mile*—the final leg of delivering connectivity from a communications provider to a customer.[5]

Two factors are obstacles, when linking catalogues of cultural heritage institutions to the Semantic Web: Firstly, their interfaces are not standardised and primarily focus on serving human visitors. Secondly, many different M3 formats and flavours have been developed and established for different domains, not to mention the plethora of in-house and proprietary formats developed by broadcasters, national institutions, and digital libraries.

While many issues—such as storage, encoding, rendering, transmission, etc. of the media assets—have already been addressed, one critical shortcoming has still not yet been extensively researched: The effective and efficient deployment of the multimedia metadata (M3) along with the content in a Semantic Web environment.

The paper is organised as follows. Section 2 illustrates the core problem with use cases from the cultural heritage domain. To overcome this situation we propose to use an M3 deployment framework based on RDFa [1], called *RDFa-deployed Multimedia Metadata* (ramm.x) [11]. Further, Section 3 describes the rationale behind ramm.x, the vocabulary and the usage, and how it can be employed to our use cases. In Section 4 we give an overview on related efforts to ours, and finally, Section 5 concludes the discussion and outlines future extensions.

2 Real-World Use Case

Imagine the case of an archive collecting historical newspapers (cf. Fig. 1), which are scanned per page. Optical character recognition (OCR) can be applied to retrieve the text in the articles and to make it searchable by full text. Other elements of the pages, such as illustrations, photographs, advertisements, etc. can be located and extracted during the digitisation process, but are not self-descriptive. In addition, a number of metadata about the asset exists, for example descriptive and administrative metadata. This metadata is commonly represented using the METS standard [9]; in our example the suggested historical newspaper profile [5] would be appropriate.

Another type of metadata is information about the digitisation process—for example device, resolution, date and time—usually stored as Exif data [8] embedded in the digital image.

[5] `http://en.wikipedia.org/wiki/Last_mile`

Fig. 1. A cultural heritage newspaper scan in an XHTML container document.

The archive in our example decides to make its collection available on the Web. It publishes the original scanned images, the text transcript and the extracted non-text elements. The most relevant of the available metadata elements are put into the asset description on the HTML page. This is very useful for a human viewer of the page.

Let us now assume that a TV journalist wants to edit a documentary on the Hungarian Revolution of 1956. He uses a Semantic Web agent to gather video and image material on this event. Clearly the image on the frontpage of the newsletter depicted in Fig. 1 would be relevant in this context, but how could it be linked to other information on the Semantic Web? Some simple descriptive metadata could be represented using the Dublin Core vocabulary [6]. As the Semantic Web agent understands RDF, we could use the RDF representation of Dublin Core [4], either in a separate document or preferably embedded in the HTML page using RDFa [1]. *But what about the information contained in the Exif and METS descriptions?*

For a second case, consider a broadcast archive offering video footage from its historical news collection for sale, such as the BBC Motion Gallery[6]. Broadcast archives have often comprehensive metadata for their content (especially in the case of news), represented in broadcast industry data formats/models such as EBU P_Meta [19], MXF DMS-1 [7] or SMEF [3]. Especially for time-based media such as video the use of these specific and more comprehensive metadata formats provides a big advantage, as they are for example able to describe the temporal structure of the video and thus allow to reference only the relevant clip of a longer video. But the problem is similar as above: *How could the journalist's Semantic Web agent access these metadata if they are in their native format or rendered into HTML for a human user?*

Taking the two above outlined use cases as a starting point, we now try to derive more generic characteristics of scenarios. We propose a solution for deploying M3 formats in the cultural heritage domain, for which the following conditions hold:

- Multimedia assets are published on the Web, i.e., for example in an HTML page (container for both the media asset, as well as the metadata);
- The media asset is published together with some multimedia metadata in an existing format, be it standardised such as Exif, MPEG-7, etc., or proprietary;
- The metadata itself is neither free text nor natively represented using an RDF-based vocabulary (as hAudio/RDFa), but uses some existing multimedia metadata format (cf. [12]);
- There is an added value from making the multimedia metadata available to a Semantic Web agent. This may be due to data integration requirements, interoperability issues or advanced query use cases.

[6] `http://www.bbcmotiongallery.com`

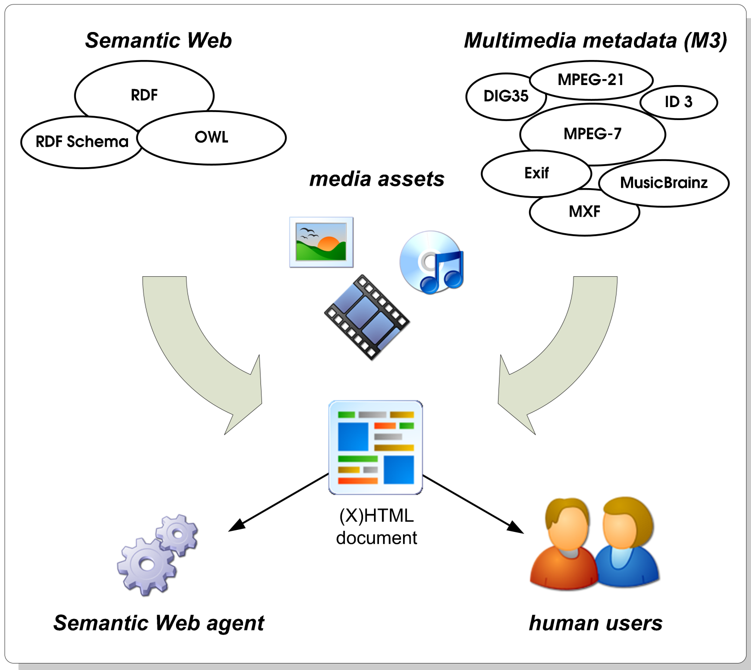

Fig. 2. Requirements for deploying M3 formats on the Semantic Web.

3 RDFa-deployed Multimedia Metadata

A vast array of multimedia metadata (M3) formats [12], as Exif, MPEG-7, ID3, etc. is available to describe what a multimedia asset is about. To enable an effective and efficient deployment of this M3 along with the content in a Web environment, the RDFa-deployed Multimedia Metadata (ramm.x) specification has been proposed. The ramm.x specification builds upon RDFa [1, 2]—a concrete serialisation syntax of the core Semantic Web data model RDF [14] as the main deployment format. Further, ramm.x utilises formalisations of multimedia metadata formats[12, Sec. 4], this is to say RDF-based schemas of the formats for the actual descriptions. The ramm.x specification allows a Semantic Web agent to determine the formalisation steps in order to perform a validation, or carry out inference.

From the use cases outlined in Section 2, we can derive a set of requirements for an M3 deployment format usable on the Semantic Web as depicted in Fig. 2:

- Embed references to existing M3 formats in Web container documents, such as (X)HTML. It should be possible to describe also certain parts of a page rather, and the resulting description must be interpretable by a Semantic Web agent operating on the RDF model;

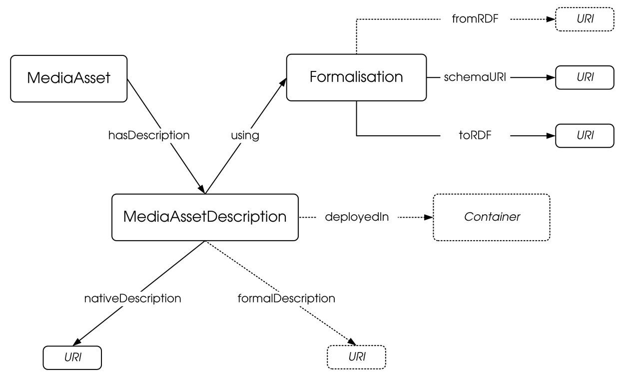

Fig. 3. The ramm.x core vocabulary at a glance.

- Provide references to Web services capable of mapping between a specific M3 format and its RDF representation (RDFising, or formalisation);
- Allow for different descriptions that may be available for a media asset. There may be alternative M3 formats covering different aspects, or there may be several ways to formalise the descriptions available in a certain format;

The ramm.x specification has the following core objectives. Firstly we aim at deploying any existing M3 format [12, Sec. 3] in a Web document (as XHTML, SMIL, SVG, etc.), enabling it to be part of the Semantic Web[7]. Secondly, ramm.x focuses on offering self-descriptive media asset descriptions, allowing to apply the "follow-your-nose" principle[8].

Last but not least, performance and scalability issues [13] have a high priority. Both human users and Semantic Web agents consuming the content, and the metadata respectively, demand an effective and efficient service actually rendered.

3.1 The ramm.x Vocabulary

In the ramm.x core vocabulary four modelling primitives (classes) are defined: `MediaAsset`, `MediaAssetDescription`, `Formalisation`, and `Container`. In the following these terms, and their relations are described. An overview of the ramm.x core vocabulary is given in Fig. 3, for a more detailed explanation, the reader is referred to the ramm.x specification [11], available at the ramm.x home page: `http://sw.joanneum.at/rammx/spec/`.

A **media asset** (`MediaAsset`) is an information resource as of [10], hence we use[9] the URL of a media asset to identify the media asset itself.

[7] `http://www.w3.org/2001/sw/`
[8] `http://esw.w3.org/topic/FollowLinksForMoreInformation`
[9] URI declaration vs. usage—cf. `http://dbooth.org/2007/uri-decl/`

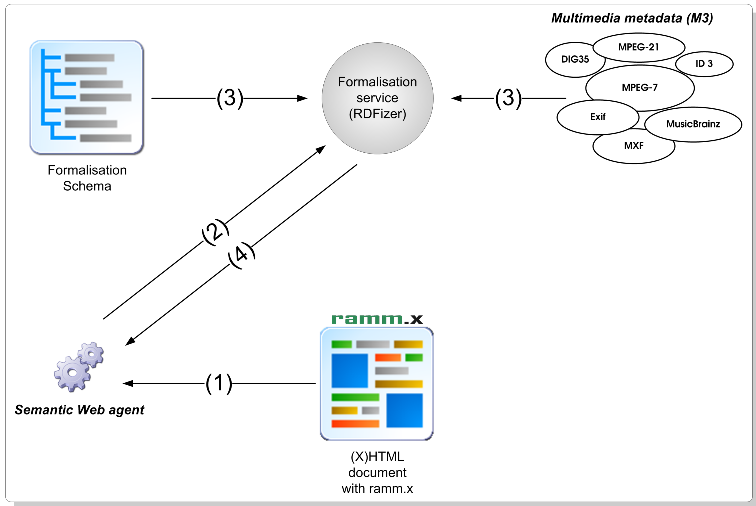

Fig. 4. Processing steps for using ramm.x descriptions.

Media asset descriptions (`MediaAssetDescription`—MAD) are the pivot element for bridging the gap between a multimedia metadata format in its native representation and the formalised version usable on the Semantic Web. Each MAD points exactly to one native description (URL of the native metadata document), and uses exactly one formalisation. Optionally, there might be a formalised version of the description available as well.

A **formalisation** (`Formalisation`) is an explicit representation of the transformation (or RDFising) process regarding a schema (in RDF-S, OWL, etc.). It provides for URLs pointing to converter services from and to a RDF-based representation, along with the respective schema.

A **container** (`Container`), is an—optional—logical unit used to bundle MAD. It acts as host for MAD to enable assigning certain properties to an array of MAD. Take a copyright note for example: Using a container, a `cc:license` is attached and all contained MAD inherit this property. Using an explicit `Container` also allows to have parts of a page (e.g., a `<div>` element) as the host, which is especially valuable in limited control environments as Wikis, Blogs, etc. The fallback rule states that in case no explicit container is present, the page in which the MAD is embedded is the container.

3.2 Processing ramm.x Descriptions

To actually use the ramm.x deployed metadata, a Semantic Web agent needs to process the container document. The following steps (cf. in Fig. 4) are necessary to access the formalised M3:

1. The Semantic Web agent fetches the ramm.x container (typically an XHTML document) with the embedded ramm.x and extracts the RDF using an RDFa extractor (on-board or as an online service—see the RDFa Implementation page[10] for further details);
2. The formalisation service is detected and invoked—this might be a lookup of a Web service, a preparation of a SOAP message, etc.
3. Using the `toRDF` property of the `Formalisation` along with the defined schema (from `schemaURI`), the original M3 format—denoted by the value of `nativeDescription`—is converted to its RDF representation by the formalisation service;
4. The resulting RDF graph (i.e. the formal representation of the M3 document) is transferred to the Semantic Web agent—ready to be used in a query, or to perform an inference on top of it.

3.3 A Real-World Example with ramm.x

We now revisit the use case introduced in Section 2 regarding a cultural heritage application that deals with newspapers. In the following we discuss, how to use ramm.x to deploy a scan of a newspaper's frontpage along with its embedded Exif metadata. When using an XHTML document as the container for both the media asset (the still image of the scanned paper) and the ramm.x MAD, it may yield a result as depicted in Fig. 1 on page 3.

The interesting part of the XHTML document is the RDFa-embedded ramm.x MAD along with the formalisation; an excerpt of an exemplary document might look as depicted in Fig. 5, below[11]. Put in simple words, the code snippet in Fig. 5 tells us that there is a media asset (`ch_example.jpg`), which we identify by its URL, hence is both a information and a non-information resource. The media asset has a description `#sample_mad` referring to the native Exif metadata embedded in the media asset. Using the formalisation `#exif_formal`, the full formal description can be obtained. The formal description conforms to the schema `http://www.kanzaki.com/ns/exif`, and can be generated using `http://www.kanzaki.com/test/exif2rdf`.

A Semantic Web agent may access the Web page with the ramm.x deployed media asset (`http://sw.joanneum.at/rammx/usecases/example/ch_example.html`), and use an RDFa extractor[12] to retrieve the embedded RDF in this page, such as depicted in Fig. 6 (in RDF N3 notation[13]).

Further, the Semantic Web agent may want to perform a SPARQL [20] query as shown in Fig. 7 to answer the question: *When was the scanned newspaper image digitised?*

---

[10] `http://esw.w3.org/topic/RDFa#Implementations`

[11] The full document, incl. a live query demo is available at the ramm.x homepage, cf. `http://sw.joanneum.at/rammx/usecases/example/ch_example.html`

[12] such as `http://triplr.org/`

[13] The shown RDF is shortened and to be understood to be interpreted with the base URI `http://sw.joanneum.at/rammx/usecases/example/ch_example.html`

```
<div about="#exif_formal" href="[ramm:Formalisation]" rel="rdf:type">
  <span rel="ramm:schemaURI" href="http://www.kanzaki.com/ns/exif" />
  <span rel="ramm:toRDF" href="http://www.kanzaki.com/test/exif2rdf" />
</div>

<div about="#sample_mad" href="[ramm:MediaAssetDescription]"
     rel="rdf:type" style="text-align: center;">
  <div about=
       "http://sw.joanneum.at/rammx/usecases/example/ch_example.jpg"
       rel="rdf:type" href="[ramm:MediaAsset]" >
    <img src="ch_example.jpg"
         alt="An example cultural heritage newspaper image from 1956."
         style="border: solid black 1px;" />
      <span rel="ramm:hasDescription" href="#sample_mad"></span>
      <p property="dc:title" datatype="xsd:string"
         style="border-top: dotted black 1px; width:500px">
        An example cultural heritage newspaper image from 1956.
      </p>
  </div>
  <span property="dcterms:created" content="2007-08-24"
        datatype="xsd:date" />
  <span rel="ramm:nativeDescription" href=
        "http://sw.joanneum.at/rammx/usecases/example/ch_example.jpg" />
  <span rel="ramm:using" href="#exif_formal" />
</div>
```

Fig. 5. An excerpt of the XHTML source code of the deployed exemplary media asset.

When the Semantic Web agent executes the SPARQL query on the embedded Exif metadata of the scanned newspaper image, it yields the following answer: `http://sw.joanneum.at/rammx/usecases/example/ch_example.jpg` was digitised on `2006-11-13T14:47:0`.

4 Related Work

To the best of our knowledge no comparable proposal to ramm.x exists. However, we build on existing standards and works. The two main works worth mentioning are the *Multimedia Vocabularies on the Semantic Web* XGR [12] and *Embedding RDF in XHTML* RDFa [1, 2]—ramm.x's deployment and main serialisation syntax. Another related work that motivated partly the ramm.x use cases is the *Image Annotation on the Semantic Web* XGR [21].

The W3C's Protocol for Web Description Resources (POWDER) Working Group[14] works on a very powerful, but rather generic standard to facilitate the publication of descriptions of multiple resources such as all those available from

[14] `http://www.w3.org/TR/powder-grouping/`

```
@prefix xsd: <http://www.w3.org/2001/XMLSchema#>
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix dc: <http://purl.org/dc/elements/1.1/> .
@prefix dcterms: <http://purl.org/dc/terms/> .
@prefix h: <http://www.w3.org/1999/xhtml> .
@prefix ramm: <http://sw.joanneum.at/rammx/ns#> .

<#exif_formal>
  a ramm:Formalisation;
  ramm:schemaURI <http://www.kanzaki.com/ns/exif>;
  ramm:toRDF <http://www.kanzaki.com/test/exif2rdf> .

<#sample_mad>
  a ramm:MediaAssetDescription;
  dcterms:created "2007-08-24"^^<xsd:date>;
  ramm:nativeDescription↩
   <http://sw.joanneum.at/rammx/usecases/example/ch_example.jpg>;
  ramm:using <#exif_formal> .

<http://sw.joanneum.at/rammx/usecases/example/ch_example.jpg>
  a ramm:MediaAsset;
  ramm:hasDescription <#sample_mad>;
  dc:title "An example newspaper image from 1956."^^<xsd:#string> .
```

Fig. 6. Extracted RDF from the example historical newspaper page.

```
prefix rdf:  <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
prefix foaf: <http://xmlns.com/foaf/0.1/>
prefix k: <http://www.kanzaki.com/ns/exif#>

SELECT ?img ?digiValue

FROM <http://www.kanzaki.com/test/exif2rdf?u=↩
      http%3A%2F%2Fsw.joanneum.at%2Frammx%2Fusecases%2Fexample↩
      %2Fch_example.jpg>

WHERE {
 ?img rdf:type foaf:Image ;
      k:exifdata ?exifData .
 ?exifData k:exif_IFD_Pointer ?ePointer .
 ?ePointer k:dateTimeDigitized ?digiValue.
}
```

Fig. 7. Querying the embedded RDF metadata of the newspaper scan.

a Web site. In the POWDER *Grouping of Resources* Working Draft [17] they explain how sets of resources may be defined, either for use in so called Description Resources or in an other context. An OWL Class is to be interpreted as the Resource Set with its predicates and objects either defining the characteristics that elements of the set share, or directly listing its elements. Resources that are directly identified or that can be interpreted as being elements of the set can then be used as the subject of further RDF triples.

Adobe's Extensible Metadata Platform (XMP)[15], which is primarily used for PDF documents (but also usable with other formats, such as JPEG, PNG, etc.), shares some of our objectives. The Open Archives Initiative (OAI) has published the Compound Information Objects draft [15], a document that deals with the publishing of aggregations of distinct information units.

Finally, the Upper Mapping and Binding Exchange Layer (UMBEL) specification[16] is a high-level subject layer for mapping various ontologies with simple binding mechanisms for any structured formalism. Simply stated, UMBEL is both a high-level reference "bag of subjects" and light-weight mechanisms for binding to Web ontologies via proxies for those subjects.

A somehow related approach to ramm.x was proposed by Pfeiffer et al. [18] back in 2005: The Continuous Media Markup Language (CMML). CMML specifies XML based markup for time-continuous data to allow it to become an integral part of the WWW. The specification allows to attach free-text annotations, metadata, captions and other textual information to clips of time-continuous data, thus enabling a timed textual representation of the data, which can be indexed by Web search engines. Further, CMML also allows to attach hyperlinks to clips of time-continuous data, enabling Web search engines to crawl the content. This also enables users to surf seamlessly between time-continuous data and other Web resources, integrating clips of media into the browsing history of a Web browser. Although the proposal seems very promising, a wide adoption may be not have been taken place due to the complexity and the disruptive model CMML is based on. This is in contrast to ramm.x, which builds on existing standards and only defines a very light-weight vocabulary.

5 Conclusion and Future Work

Archives and libraries increasingly publish parts of their collection on the Web, including not only textual information but also multimedia assets. In contrast to text documents multimedia assets can only be indexed and interpreted if published together with metadata describing them. We have identified a shortcoming in deploying multimedia metadata (M3) on the Semantic Web and have proposed a solution which allows to deploy M3 vocabularies so that they can be used by Semantic Web agents understanding RDF. We have described two related use cases motivated by the search for image and video motivation, one based on a newspaper archive and the other on a broadcast archive. Using the newspaper

[15] http://www.adobe.com/devnet/xmp/pdfs/xmp_specification.pdf
[16] http://www.umbel.org/proposal.xhtml

example, we have outlined how ramm.x can be employed for deploying metadata so that it is accessible for a Semantic Web agent.

The RDFa-deployed Multimedia Metadata (ramm.x) framework allows to deploy references to media asset descriptions represented in any multimedia metadata vocabulary in a Web document. It defines and utilises a light-weight vocabulary consisting of four classes which is embedded into the host document using RDFa. It has to be emphasized that ramm.x is not an RDF-based multimedia description vocabulary or a multimedia ontology, nor does it replace existing technologies. It is complementary to multimedia metadata vocabularies, their formalisation and Web technologies in that it allows the deployment of non-RDF media asset descriptions on the Semantic Web. This is achieved by linking the host document of the media asset, the media asset description and a service capable of providing a formalisation (i.e., an RDF representation) of the media asset description.

ramm.x is useful if the media asset description is not natively available in RDF, otherwise a plain RDFa description would be sufficient. The same is true for global-only metadata, for example a `dc:title` or the like. A rather rough analysis already performed in popular social media sites, such as `flickr.com`, or `last.fm` yielded that most of the current metadata is of global type. Using specific M3 vocabularies is necessary the more detailed the media description is, e.g., for describing the (spatio-)temporal structure of a video, with annotations on all of the segments. Also in many archives and libraries, the description in a native M3 format is readily available and thus ramm.x simplifies deployment. Of course the successful use of ramm.x depends on the availability of a service capable of formalising the native description format.

In contrast to other RDF access methods—such as SPARQL-end points[17]—ramm.x is resource-oriented. This is to say instead of a global and coarse-grained access, the focus of ramm.x is on enabling fine-grained deployment of media asset descriptions. Especially regarding contextual issues this is of great importance. When a media asset is taken out of its original context, say, the hosting HTML page, one would certainly expect that the associated metadata is lost. One of the future research directions in ramm.x is how to address this issue through the utilisation of watermarking techniques that carry a pointer to the original container URI. Once such a back-link is provided from within the media asset, the standard ramm.x description processing as described in Section 3.2 can be applied.

From the exemplary use of ramm.x in the cultural heritage domain two issues have been identified that should be addressed by future extensions. One concerns collecting formalisations of different M3 vocabularies as well as the services capable of doing the formalisation, which facilitates the practical use of ramm.x. The other issue stems from the broadcast archive use case described in Section 2. It concerns the deployment of large media asset descriptions, such as a detailed annotation of a one hour documentary. Clearly ramm.x is more beneficial in such a case than when just describing a few global metadata elements.

[17] `http://esw.w3.org/topic/SparqlEndpointDescription`

But as the media asset description can be very large in that case, a kind of "streaming mode" allowing to access the part of the description that is relevant to the current segment of the media asset would be useful.

References

1. B. Adida and M. Birbeck. RDFa Primer 1.0. W3C Working Draft, W3C Semantic Web Deployment Working Group, 2007.
2. B. Adida and M. Hausenblas. RDFa Use Cases: Scenarios for Embedding RDF in HTML. W3C Working Draft, W3C Semantic Web Deployment Working Group, 2007.
3. British Broadcasting Corporation (BBC). Standard Media Exchange Framework (SMEF) Data Model 1.5. `http://www.bbc.co.uk/guidelines/smef`, 2000.
4. D. Beckett, E. Miller, and D. Brickley. Expressing Simple Dublin Core in RDF/XML (DCMI Recommendation). `http://dublincore.org/documents/dcmes-xml`, 2002.
5. M. V. Cundiff. METS Profile for Historical Newspapers (Draft). `http://www.loc.gov/standards/mets/test/ndnp/00000010.html`, 2006.
6. DCMI. Information and documentation—The Dublin Core metadata element set. ISO 15836, 2003.
7. DMS-1. Material Exchange Format (MXF) – Descriptive Metadata Scheme-1. SMPTE 380M, 2004.
8. Exif. Digital Still Camera Image File Format Standard (Exchangeable image file format for Digital Still Camera: Exif). JEIDA-49, version 2.1, 1998.
9. Digital Library Federation. METS – Metadata Encoding and Transmission Standard. `http://www.loc.gov/standards/mets`, 2006.
10. W3C Technical Architecture Group. httpRange-14: What is the range of the HTTP dereference function? `http://www.w3.org/2001/tag/issues.html?type=1#httpRange-14`, 2005.
11. M. Hausenblas and W. Bailer. RDFa-deployed Multimedia Metadata (ramm.x). Specification, ramm.x Working Group, 2007.
12. M. Hausenblas, S. Boll, T. Bürger, O. Celma, C. Halaschek-Wiener, E. Mannens, and R. Troncy. Multimedia Vocabularies on the Semantic Web. W3C Incubator Group Report, W3C Multimedia Semantics Incubator Group, 2007.
13. M. Hausenblas, W. Slany, and D. Ayers. A Performance and Scalability Metric for Virtual RDF Graphs. In *3^{rd} Workshop on Scripting for the Semantic Web (SFSW07)*, Innsbruck, Austria, 2007.
14. G. Klyne, J. J. Carroll, and B. McBride. Resource Description Framework (RDF): Concepts and Abstract Syntax. W3C Recommendation, World Wide Web Consortium, 2004.
15. C. Lagoze and H. Van de Sompel. Compound Information Objects: The OAI-ORE Perspective. Discussion Document, Open Archives Initiative, 2007.
16. S. Ortiz. Searching the Visual Web. *Computer*, 40(6):12–14, 2007.
17. A. Perego and P. Archer. Protocol for Web Description Resources (POWDER): Grouping of Resources. W3C Working Draft 9 July 2007, W3C Protocol for Web Description Resources Working Group, 2007.
18. S. Pfeiffer, C. Parker, and A. Pang. The Continuous Media Markup Language (CMML), Version 2.0. Internet-Draft, The Internet Engineering Task Force, 2005.
19. P_Meta. The EBU Metadata Exchange Scheme. EBU Tech 3295, version 1.2, 2005.

20. E. Prud'hommeaux and A. Seaborne. SPARQL Query Language for RDF. W3C Candidate Recommendation 14 June 2007, W3C RDF Data Access Working Group, 2007.
21. R. Troncy, J. van Ossenbruggen, J. Z. Pan, and G. Stamou. Image Annotation on the Semantic Web. W3C Incubator Group Report, W3C Multimedia Semantics Incubator Group, 2007.