The LinkedTV Platform –Towards a Reactive Linked Media Management System

Jan Thomsen, Ali Sarioglu, and Rolf Fricke

Condat AG, Alt-Moabit 91D, 10559 Berlin, Germany
{jan.thomsen,ali.sarioglu,rolf.fricke}@condat.de
http://www.condat.de

Abstract. The LinkedTV Platform was developed in the EU FP7 project LinkedTV[1] with the objective of supporting semi-automatically linking TV content with additional information and content. The realized approach can be used for many different applications and purposes, in which the LinkedTV Platform serves as the backbone connecting and managing the different services for analyzing, annotating, linking and enriching media resources and storing the aggregated metadata. This paper describes the subsequent evolution of the LinkedTV Platform towards a Linked Media Management System supporting flexible and scalable linked media applications.

Keywords: linked media, media fragments, semantic media, media analysis, media annotation, media enrichment, media interlinking, reactive systems

1 Introduction

In recent years research efforts have sought to deal with the structured analysis, annotation, enrichment[2] and interlinking of multimedia resources, subsumed under the label of Semantic Multimedia [1],[2] or more recently (with the use of Linked Data for concept identification and linkage) Linked Media [3], in order to improve computer-based processing and re-use of digital media. The creation of Linked Media covers a wide range of approaches, including media metadata extraction, text, audio and video analyses, transcription / speech-to-text, Named Entity Recognition (NER) and semantic similarity measures / graph transversal to link to conceptually related resources. Once this Linked Media has been produced it could support a wide range of specific applications and user experiences including Second Screen applications to complement and enhance media consumption, media content recommendation and personalization, on-screen information overlays and so on. However current media management does not take

[1] http://www.linkedtv.eu

[2] Here, we mean the association of some content with a part of the multimedia resource based on its annotation with some concept(s), to which the content is related (e.g. explanatory).

the support of such functionalities into account nor provide a holistic workflow for their combination, leading to the necessity for a new type of media management system which we call Linked Media Management Systems. The LinkedTV Platform[4] is an example of this new type of system and was an outcome of the EU FP7 project LinkedTV [5] which ran from October 2012 until March 2015 and has been further developed since. This paper describes the current state of the evolution of the LinkedTV Platform to become a general solution to Linked Media Management for future media services. In Section 2, we introduce the concept of Linked Media Management and the functionalities required and provide a short overview of the state of the art with related approaches. Section 3 describes the workflow as realized by the LinkedTV Platform as a result of the LinkedTV project. In Section 4, we present the evolution of the LinkedTV Platform towards becoming a Linked Media Management System. Section 5 concludes the paper by discussing the approach and giving a short outlook.

2 Linked Media Management: State of the Art

Today's media service providers rely on legacy MAM systems without the functionalities required for Linked Media, and faced with the need for costly ad-hoc integration of diverse services to realize Linked Media it is clear that this is currently an unassailable barrier to industry uptake.

Based on the experience with LinkedTV we propose that Linked Media Management Systems should support the following main functionalities:

- Not just ingestion of media resources but also files related to them, such as subtitles/transcriptions (SRT, VTT) or metadata descriptions such as TVAnytime [6]
- Analysis of video and audio tracks of a media resource with respect to temporal or spatial aspects, such as fragmentation (to shots and scenes), object (re-)detection, face recognition, speaker identification, visual concept classification, etc. resulting in an audio-visual description of the media resource, including temporal and spatial segments. This can be captured in e.g. an MPEG-7 AVDP profile [7]
- Generation of Media Fragment URIs [8] from the profile segments, which supports simplified and Web-friendly references to temporal and spatial parts of a media resource
- Generation of structured annotations at the media fragment level by classification (i.e. shot, visual concept type, detected face, etc.) and provenance (e.g. which classifiers were used, when), preferably by using commonly understood and used models such as the W3C Annotation Data Model [9] and W3C PROV Data Model [10]
- Named Entity Recognition on textual resources (subtitles, speech transcriptions) for associating named entities such as persons, events, locations, objects to the media fragments (semantic annotation)
- Linking named entities to Linked Data resources on the Web, such as dbpedia (a structured metadata representation of the information in Wikipedia),

which provide both a global disambiguation mechanism and access to additional metadata and links
- A repository for storing this metadata information. Given the proposed specifications, a graph-based storage and retrieval approach makes most sense, e.g. RDF.
- Connecting to enrichment services which can suggest related content such as conceptual descriptions, Web pages or similar images/videos for media fragments according to their annotated entities
- Exposing this data to client applications as a Web service through a consistent REST based interface, which is in line with service-oriented architectures and Web-based communication, enabling any connected application (e.g. on a SmartTV, Set Top Box or mobile device acting as a Second Screen) to easily and asynchronously retrieve data for an enhanced media service (e.g. LinkedTV providing related content synchronized to the TV broadcast on a second screen)
- Providing for manual oversight, correction and confirmation of the generated annotations and enrichments for a media resource
- Managing the workflow of all these different steps and phases
- Offering a framework for integrating these different systems and services, as well as further ones, e.g. integration with TV Program Planning or Production Systems, streaming servers, Content Delivery Networks, Content Management Systems, or Rights Management Systems.

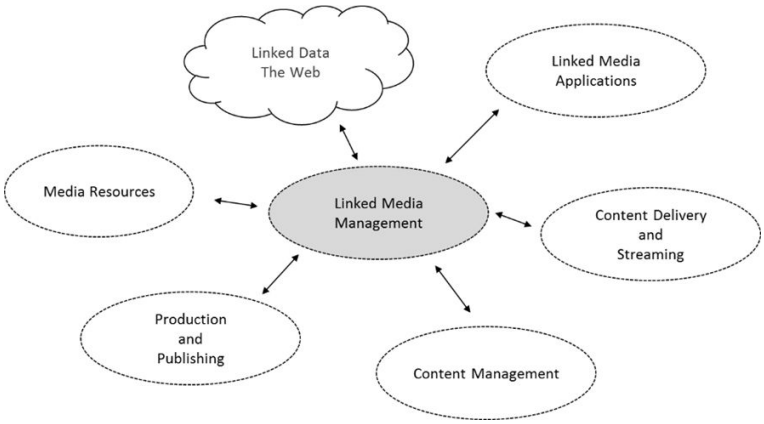

Fig. 1. Linked Media Management connecting different processes, services and resources.

Linked Media Management is thus a complex and innovative new media management process which connects and manages different kinds of services from production to consumption of media content Fig. 1). However, a Linked Media

Management System does not necessarily have to host all mentioned functionalities locally. In line with current trends towards service-oriented architectures and cloud-based computing, the architecture should allow the connection of external 3rd party services, open or commercial. The crucial core components are the workflow management, the connections to the different services (Linked Media Services) and the metadata store (Linked Data Repository).

As of today Linked Media Management Systems do not exist as a category named as such. However, there are similar and related approaches such as the former Linked Media Framework (LMF) [11], now being continued within the Apache Marmotta project [12], or the MICO Platform [13], [14] (currently being developed on top of Apache Marmotta) which is very similar to the LinkedTV Platform approach, even down to technological choices, as we will return to in the conclusion.

3 The LinkedTV Workflow

The LinkedTV Platform would be an example of a Linked Media Management System. The generation, aggregation and usage of the annotated media fragments is done through a general workflow of different tasks which are grouped into three main successive workflows (Fig. 2):

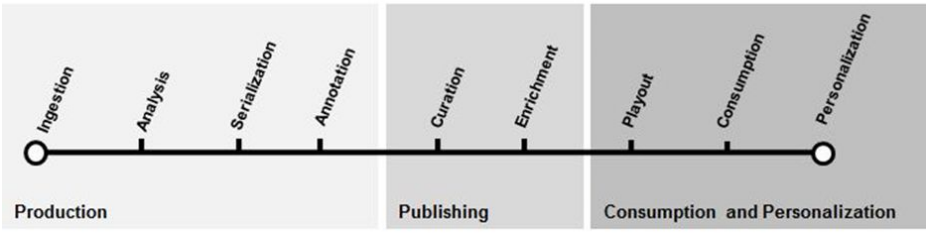

Fig. 2. The basic LinkedTV workf

The **production workflow** has the objective to make the media resources Linked Media-ready. It consists of the steps: (1) ingestion of the video itself and related metadata (TVAnytime metadata and subtitle files), including also an encoding to various resolutions and formats and a transfer to a streaming server, (2) deep analysis of the video and audio tracks with various techniques developed within the project [15], (3) serialization of the analysis results and metadata files into a common LinkedTV data model [16] (being an RDF-based description format making use of all kinds of existing ontologies such as the W3C Media Ontology and which provides annotated media fragments identified using Media Fragment URIs), (4) an automatic annotation of the media fragments with provenance information, named entity recognition results and other media resource information [17].

The **publishing workflow** is mainly a manually curation and enrichment process. Its objective is to take the raw LinkedTV production data, evaluate it, correct it (e.g. incorrect chapter segmentations) filter out unwanted data, and most notably, enrich it by adding all kinds of related material to the various chapters or entities by making use of a rich set of enrichment services developed within the LinkedTV project [18]. For this a specific LinkedTV Editor Tool has been developed [19].

The **playout, consumption and personalization workflow** is the process of: (1) playing the video itself to a viewer, (2) displaying the related content either on the same screen or a second screen depending on the respective scenario, (3) adapting the related information to the viewers profile, (4) reacting to viewer events like pause, fast forward or switch channel, (5) building the user profile out of her implicit or explicit preferences. Steps (3) to (5) are the personalization part of the consumption and, of course, optional; see [20] for an overview of personalized content delivery in LinkedTV.

4 The LinkedTV Platform Evolution

The first version of the LinkedTV Platform was designed to linearly process the different steps as depicted in Fig.2. However, as this approach does not scale very well and is not resilient as well as not flexible enough for adding new services, the platform has been redesigned according to the principles of reactive design (responsiveness, resilience, elasticity, message driven) [21]. In fact, a lot of the above mentioned tasks do not need to be executed subsequently but can run in parallel and independently from each other. After a redesign, the new distributed processing path can be illustrated like a "tube map" (Fig. 3).

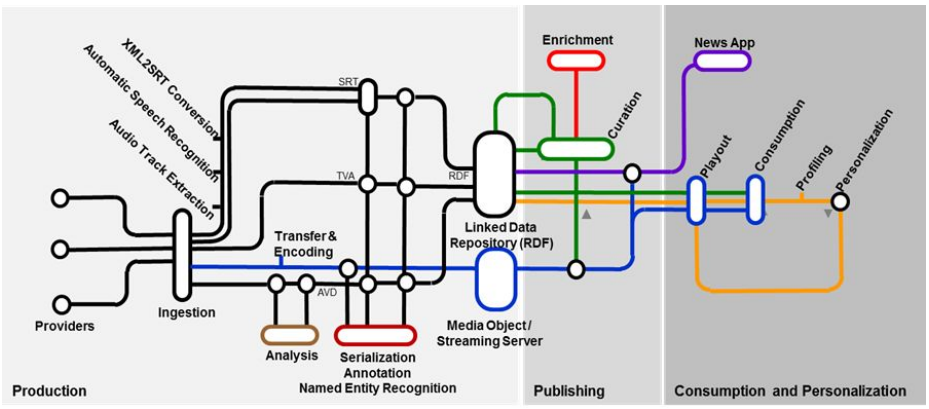

Fig. 3. The LinkedTV Workflow Map Evolution

In this map, the black lines depict the services running locally in the LinkedTV platform, whereas the colored ones depict the connected services provided by third parties via REST APIs.[3]

The LinkedTV Platform (i.e. the system running the black lines) has been designed and implemented now in a way which maps system components almost 1:1 to parts as depicted in Fig. 2. Thus, this figure can be read also as an architectural diagram. Using a message driven architecture, the nodes represent individual micro services whereas the lines represent queues. Normally, the nodes act as consumers, reading messages from the dedicated queue, do some processing (which may or may not include invoking external REST services) and then produce again messages and send into the next queue. Sending messages to several queues (like the Ingestion service does) evokes parallel processing. But there are also pure producers (like the Providers, which initiate the processing by sending media resources and associated metadata files to the ingestion service) or pure consumers (like the micro service storing RDF graphs in the Linked Data Repository).

With the vision of creating an open, reactive Linked Media Management System, this redesign has, among others, the following main properties which are relevant in the LinkedTV context: (1) it is highly scalable as each node and the respective queues can be multiplied so an implicit work and load balance is ensured. Through this, a high level of elasticity and resilience is achieved. Of course, however, the degree of these properties in the overall system is generally limited by how reactive the connected services themselves are, so these have to scale in the same way. (2) The system is very open as easily new queues and new micro services can be added without affecting current workflows. (3) It is also quite easy to configure new set-ups (routes) for individual clients or use case profiles, e.g. depending on whether subtitle files are already provided by the source or have to be generated from automatic speech recognition. (4) The whole process management, delivery and acknowledgement control is done by the underlying message system.

4.1 Technologies used and links

The current version of the LinkedTV Platform has been implemented on base of RabbitMQ [22] which employs the AMPQ protocol [23]. Client micro services are implemented in either Java or Python, but a lot of others are supported by RabbitMQ as well, such as JavaScript or Scala. For other technologies like Node.js there exist specific AMPQ libraries. For the storage of RDF graphs, Openlink Virtuoso is used. A web based interface for uploading videos and starting the LinkedTV process is available under `http://api.linkedtv.eu`. A REST inter-face for accessing the media generated media fragment annotations exists under `http://data.linkedtv.eu` (realized with Elda [24], an open source implementation of the Linked Data API [25] by Epimorphics) and also a SPARQL

[3] in our case LinkedTV partners, a full list of LinkedTV platform compatible tools and services and the respective partner has been published at `http://www.linkedtv.eu/demos-materials/tools-and-services/`.

endpoint under `http://data.linkedtv.eu/sparql`. Fig. 4 shows a screenshot of the LinkedTV dash-board displaying metadata and to a media resource processed by the LinkedTV system.

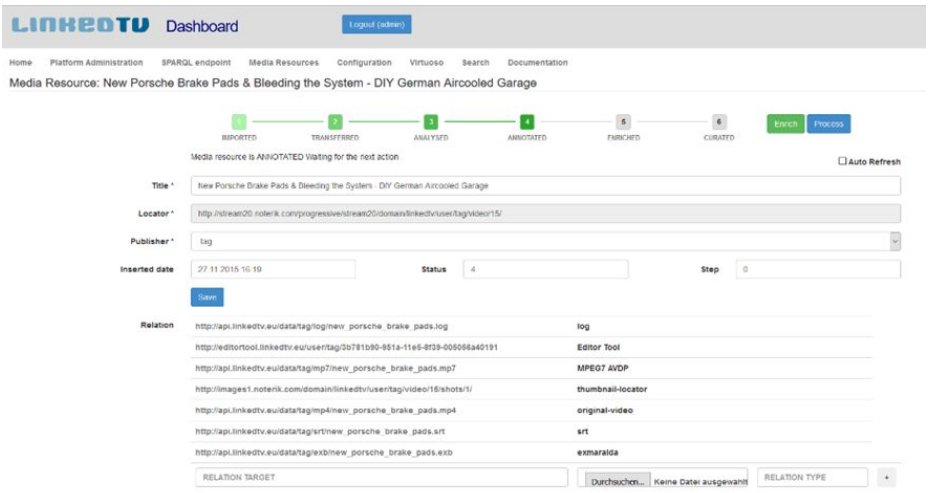

Fig. 4. Screenshot of the LinkedTV Dashboard

5 Summary, discussion and outlook

The core functionalities required for the purposes of Linked Media Management as followed within LinkedTV and described in Section 2 project have already been realized within the first version of the LinkedTV Platform. The current version of the LinkedTV Platform did not add any new functionality, but focused on a redesign and refactoring process in order to prepare the LinkedTV Platform for industrial use. We did this by applying principles of reactive design. In comparison with the above mentioned related approaches, the LinkedTV Platform architecture is in fact quite similar to the approach taken by the MICO Platform: Both ones are based on a message driven backend with even RabbitMQ as the same technology; within LinkedTV we also consider using Apache Camel for Service Orchestration. Both employ Linked Data Repositories (LinkedTV: Virtuoso/Elda, MICO: Apache Marmotta).

The main difference between these two architecture lies in the integration of the connected Linked Media Services: while LinkedTV is a very distributed architecture connecting the different services over the Web via REST services, MICO seems to use mainly native components. While the LinkedTV approach seems to be more open and more directed to creating a web-scale envisioned Linked Media Layer, the MICO approach will quite surely be more efficient.

As next steps we will be addressing the evaluation of the new platform architecture by testing different configurations and profiles. From anecdotal experience we can say that within the first platform version the processing ratio was about 1.5, i.e. the whole automated processing of a 20-min German news show took about 30 min, whereas within the new platform architecture this ratio went down to about 0.75. By far the most time takes the visual analysis, but this can be done now almost completely in parallel to the other steps.

Acknowledgments. This work has been partially supported by the European Commission via the FP7 project LinkedTV (GA 287911).[4] We wish to thank particularly Lampis Apostolidis (CERTH), José Luis Redondo Garcia (EURECOM), Daniel Oeckeloen and Pieter van Leeuwen (Noterik), Jaap Blom (Sound and Vision) for greatly supporting the further development of the LinkedTV Platform after the end of the project LinkedTV and LinkedTVs former scientific coordinator Lyndon Nixon for his input on this paper as well as all of the developers and advocates for the different Web services which are used by the Platform.

References

1. Evaggelos Spyrou et al., Semantic Multimedia Analysis and Processing, CRC Press, June 2014. ISBN 978-1-4665-7549-3.
2. Lyndon Nixon and Raphaël Troncy, Survey of Semantic Multimedia Annotation Tools for the Web, in ESWC Satellite Events 2014, LNCS 8798, pp. 100-114, 2014.
3. Lyndon Nixon, The importance of Linked Media to the Future Web, in WWW 2013 Companion, pp. 455-456, 2013. http://www2013.wwwconference.org/companion/p455.pdf
4. Jan Thomsen, Final LinkedTV End-to-end Platform, LinkedTV project deliverable D5.6, October 2014. http://www.linkedtv.eu/deliverables
5. Lyndon J. B. Nixon, Delivering related web content synchronized to online television: The LinkedTV solution, Proc. 2015 IEEE 5th International Conference on Consumer Electronics (ICCE 2015), September 2015
6. TVAnytime, http://www.tv-anytime.org/
7. MPEG7 Audiovisual Description Profile, http://mpeg.chiariglione.org/standards/mpeg-7/profiles
8. Media Fragments URI 1.0, https://www.w3.org/TR/media-frags/
9. W3C Web Annotation Working Group, https://www.w3.org/annotation/
10. PROV-Overview, https://www.w3.org/TR/prov-overview/
11. Linked Media Framework, http://semanticweb.org/wiki/Linked_Media_Framework.html
12. Apache Marmotta, http://marmotta.apache.org/
13. MICO Platform, http://www.mico-project.eu/platform/
14. Patrick Aichroth et al., MICO Technical Report Vol. 3 Enabling Technology Modules: Version I, Salzburg Research Forschungsgesellschaft mbH, Salzburg, Austria, 2016

[4] The full list of LinkedTV partners can be found at http://www.linkedtv.eu/about-the-project/consortium/.

15. Evlampios Apostolidis et al., Visual, text and audio information analysis for hypervideo, final release, LinkedTV project deliverable 1.4, September 2014. http://www.linkedtv.eu/deliverables
16. LinkedTV Ontology, http://www.linkedtv.eu/ontology
17. J.L. Redondo Garcia et al., Annotation and retrieval module of media fragments, LinkedTV project deliverable 2.4, October 2013. http://www.linkedtv.eu/deliverables
18. Tomaš Kliegr et al., LinkedTV Framework for generating video enrichments with annotations, LinkedTV project deliverable 2.6, October 2014. http://www.linkedtv.eu/deliverables
19. LinkedTV EditorTool, http://editortoolv2.linkedtv.eu/
20. Dorothea Tsatsou et al., When TV meets the Web: towards personalised digital media, book chapter in [1].
21. The Reactive Manifesto, http://reactivemanifesto.org
22. RabbitMQ, http://www.rabbitmq.com/
23. AMPQ, https://www.amqp.org/
24. Elda, https://github.com/epimorphics/elda
25. Linked Data API , https://code.google.com/archive/p/linked-data-api/wikis/Specification.wiki