Title	Podcast Pinpointer: A Multimedia Semantic Web Application
Author(s)	Hogan, Aidan; Harth, Andreas; Breslin, John
Publication Date	2005
Publication Information	Aidan Hogan, Andreas Harth, John Breslin "Podcast Pinpointer: A Multimedia Semantic Web Application", Proceedings of the 2nd European Workshop on the Integration of Knowledge, Semantic and Digital Media Technologies, 2005.
Publisher	IEEE
Item record	http://hdl.handle.net/10379/487

PODCAST PINPOINTER: A MULTIMEDIA SEMANTIC WEB APPLICATION

Aidan Hogan, Andreas Harth, John G. Breslin

Digital Enterprise Research Institute, National University of Ireland Galway,
Galway, Ireland

firstname.lastname@deri.org

Keywords: Semantic Web, Podcast, Multimedia.

Abstract

In late 2004, a new method of publishing multimedia broadcasts on the Internet became popular called 'Podcasting'. Podcasting incorporates existing feed description formats, namely RSS 2.0, to deliver various enclosed files. This allows users to subscribe to feeds, receiving updates periodically. Originally intended for self-publishing and syndication of audio files, usage of Podcasting for video files has become quite popular. Indeed, thousands of Podcast feeds are now available, reaching a wide range of listeners and viewers. The rapid development of such a technology proves the demand for structured formats of describing multimedia data, facilitating location and retrieval of desirable media for consumers. This paper proposes such a format, using Semantic Web technologies, that has been used in a prototype application for the intelligent location and retrieval of Podcasts.

1 Introduction

As bandwidth increases and new internet-friendly compressed multimedia formats become popular, demand for, and consequently supply of multimedia files on the Web has experienced a surge in volume. However, such a proliferation of multimedia data corresponds in difficulties with regards data retrieval. This problem is not specific to this domain however, and is a more general symptom of the expansion of the Web. Many researchers are focusing their efforts on solving this problem, one such venture being in the area of the Semantic Web [1].

Podcasts have illustrated the demand for structured formats to describe multimedia. They too have experienced a phenomenal escalation in popularity (through radio station Podcasts or through individuals "audio blogging" and "video blogging" / "vlogging"). As is, Podcasts are described in RSS 2.0 format files, with supplementary specifications cropping up to allow multimedia-specific descriptions to be stipulated. Unfortunately, however, the RSS 2.0 format of feed description does not exploit the capabilities of RDF and is not a participant of the Semantic Web venture.

RDF, as a W3C recommendation, is a framework for providing the Web with structured data. There is a considerable amount of research being pursued in developing the technology and applications pertaining to RDF and the Semantic Web. Currently, many applications exist to handle and analyse such data and ultimately provide services of utility to users, and to bring order to a rather chaotic World Wide Web.

RSS 1.0[1] is a format for creating feeds based on RDF/XML. As such, it is within the realm of the Semantic Web. Indeed, by using an RSS 1.0 format template to describe multimedia data, the framework exists for bringing structured metadata to multimedia files, thus introducing these files into the Semantic Web. By converting Podcast feeds from RSS 2.0 to RSS 1.0, they could begin to realise the potential of all these new applications, technology and research.

In addition to this, but perhaps more challenging, it would be beneficial to create such RSS 1.0 descriptions for more generic multimedia files which do not possess a corresponding feed. In doing so, there would be a wealth of multimedia data available to various Semantic Web applications, enriching the services such applications provide to users and making such data more easily accessible and so benefiting both.

2 Podcasting and Metadata

Podcasting is very much a booming technology[2]. From humble beginnings, it has become a prevalent force in multimedia syndication and distribution. Part of its strength lies in its relative simplicity, allowing casual users to create and publish *online radio shows* and get them to a wide audience. All a user needs to create a Podcast is some simple recording equipment, a basic understanding of RSS 2.0 and some web space. It is also convenient for the consumer, who can use traditional feed-catching methods to *subscribe* to a Podcast feed and receive automatic intermittent updates.

However, it is not only casual users that are publishing Podcasts, as larger organisations have seen the positive aspects of such technologies. Many companies have begun publishing media via Podcasts, ranging from the BBC and ABC News to NASA and Disney. Indeed, there has been much interest shown by many American radio stations, who have begun making Podcasts of their programmes available online (e.g. NPR's Science Friday).

---

[1] http://purl.org/rss/1.0/

[2] http://www.tdgresearch.com/press044.htm

Other companies have taken interest in promoting the technology. Apple have become heavily involved in the area of Podcasting. Being implicitly involved from the start (Podcasting being a portmanteau involving a reference to their portable MP3 player, iPod, and broadcasting), they took an active role in June 05. Some of their products such as iPod and iTunes were already Podcast friendly, featuring music file synchronisation features exploited by Podcasting origins. They began providing Podcatching software, which allows users to browse and subscribe to feeds, and also a Podcast directory, a categorised listing of broadcasts. They upgraded the firmware on iPods to display Podcasts in the top-level music menu. They also built a specification for describing aspects of Podcasts which would appear in a predefined format on iTunes and iPod displays[3]. This *itunes* namespace[4] is intended for use within an RSS 2.0 environment.

Besides this specification, Yahoo! have also created a namespace for syndicating media items[5]. This specification is intended as a replacement for the RSS enclosure element, offering a more expressive vocabulary for describing media items.

To introduce Podcasts to the Semantic Web, these feeds would have to be converted to RDF. Whilst RSS 2.0 can quite neatly convert over to RSS 1.0, with the exception of a few features, the above specifications are incompatible with a direct conversion, so it was necessary that a unique RDF/XML specification[6] be created.

Existing Podcast metadata, following the RSS 2.0 template, can then be converted to follow this specification. By formulating such a method of conversion, we enable other RDF tools to interpret and handle Podcast files, and create a solid structure in which more general multimedia entities can be described. Indeed, a search engine was constructed by using existing RDF tools on the converted data. In doing so, we ensured that the created RDF/XML specification would be interoperable and easily manageable by other Semantic Web software agents.

The main feature of the specification is the *podcast:media* container and *podcast:content* element. The *podcast:media* device is a container for all *podcast:content* elements. Such elements which refer to the same content can be grouped in the one *podcast:media* container. The *podcast:content* element itself allows various attributes to be specified, each described in the specification document.

Other features include *podcast:owner*, a container for a *foaf:Person* element, from the FOAF specification. Within the FOAF[7] area exists an expressive vocabulary of describing personal details, which are all valid within the *foaf:Person* element.

Also present is a category element, which allows conveyance of the topic of a channel or item. It allows specification of the category description and the domain to which the description refers.

```
<?xml version="1.0" encoding="UTF-8" ?>
<rdf:RDF
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:podcast="http://sw.deri.org/2005/07/podcast#"
  xmlns:foaf="http://xmlns.com/foaf/0.1/"
  xmlns:dc="http://purl.org/dc/elements/1.1/"
  xmlns="http://purl.org/rss/1.0/">

<channel rdf:about="http://example.com/podcasts/index.html">
  <title>New Specification</title>
  <description>An example Podcast feed describing this new
specification.</description>
  <podcast:owner>
    <foaf:Person>
      <foaf:name>Aidan Hogan</foaf:name>
      <foaf:mbox>aidan.hogan@deri.org</foaf:mbox>
    </foaf:Person>
  </podcast:owner>
  <podcast:topic>
    <podcast:category podcast:domain="http://example.com/Taxonomy/>
      <podcast:name>Semantic Web</podcast:name>
    </podcast:category>
  </podcast:topic>
  <items>
    <rdf:Seq>
      <rdf:li rdf:resource="http://example.com/podcasts/newShow1.mp3" />
      <rdf:li rdf:resource="http://example.com/podcasts/newShow2.mp3" />
    </rdf:Seq>
  </items>
</channel>

<item rdf:about="http://example.com/podcasts/newShow1.mp3">
  <dc:creator>Andreas Harth</dc:creator>
  <title>Item #1</title>
  <description>This week Andreas talks all about the new RSS 1.0 Podcast
Specification</description>
  <podcast:media>
    <podcast:content
      podcast:url="http://example.com/podcasts/newShow1.mp3"
      podcast:fileSize="16490812"
      podcast:type="audio/mpeg"
      podcast:expression="full"
      podcast:isDefault="true"
    />
    <podcast:content
      podcast:url="http://example.com/podcasts/newShow1.ra"
      podcast:fileSize="482342"
      podcast:type="audio/x-realaudio"
      podcast:expression="full"
      podcast:isDefault="sample"
    />
</item>

<item rdf:about="http://example.com/podcasts/newShow2.mp3">
  <dc:creator>John Breslin</dc:creator>
  <title>Item #2</title>
  <description>This week John gives his two cents on the new RSS 1.0
Podcast Specification</description>
  <podcast:media>
    <podcast:content
      podcast:url="http://example.com/podcasts/archive/mediafile.mp3"
      podcast:fileSize="8727310"
      podcast:type="audio/x-mpeg"
      podcast:bitrate="128"
      podcast:duration="1058"
    />
  </podcast:media>
</item>
</rdf:RDF>
```

Figure 1: An example usage of the new specification.

[3] http://phobos.apple.com/static/podcast_specifications.pdf
[4] http://www.itunes.com/DTDs/Podcast-1.0.dtd
[5] http://search.yahoo.com/mrss
[6] http://sw.deri.org/2005/07/podcast
[7] http://xmlns.com/foaf/0.1/

3 Syndicating Other Multimedia

The expression of Podcasts in RSS 1.0 format is essentially a process of reorganising pre-structured data (see Figure 2). In fact, it may also be possible to explicitly define metadata in an RSS 1.0 format for multimedia data where such metadata does not already exist.

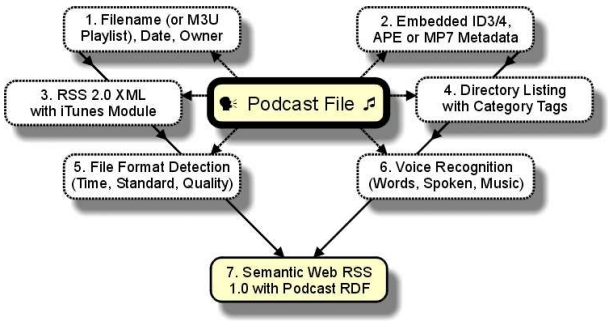

Figure 2: Some sources of metadata for a Semantic Web representation of a Podcast file.

One possibility would be to parse and convert metadata embedded in multimedia files. An example of such would be ID3/ID4/APE tags embedded in MP3 files. Such tags provide information relating to the file name, song or piece name, creator or artist, album, genre and year. Other multimedia metadata standards include the MPEG series of standards. Of particular interest would be MPEG-7, a means of expressing audio-visual metadata in XML. Upon parsing out such information, a pre-templated RSS 1.0 file could be filled with the available information. This would then be interpretable by the same tools as the modified Podcasts.

Another interesting application of the RDF Podcast specification is in relation to the Asterisk project[8], an open source Linux-based PBX application. A potential side use for Asterisk on online bulletin boards has already been touted as the next evolutionary step in web-based discussions by Drupal/CivicSpace, whereby phone conversations made through the Asterisk PBX would be recorded and stored as streamed or downloadable audio conversations for other readers of the bulletin board discussion.

Many sites have begun using voice recognition technology in the indexing of multimedia files, one such site being 'blinkx'[9]. Voice recognition software has seen many advances in recent years, and is becoming more and more accurate. Such sites use the ever-more advanced technology to create a transcript of spoken words in the audio of files. Indeed this would be quite useful in keyword searches.

On top of these transcripts then, HLT (Human Language Technology) could be implemented to derive a structure from the prose [3]. This structure could take the form of various elements within an RSS 1.0 document, accompanying all other metadata already located.

[8] http://www.asterisk.org/

[9] http://www.blinkx.tv/downloads/blinkx_TV_White_Paper_v1.0.pdf

4 Locating & Crawling Multimedia on the Web

To create a central repository of multimedia data, it would be obligatory to locate and crawl such data. However, such data, including Podcasts, is not interlinked, and so achieving a complete set of data is demanding. The most complete method of data acquisition would be a crawl of the entire Web, however the resources necessary to complete said task are quite substantial.

To begin with we needed a raw dataset of crawled Podcasts. To attain an initial set of links we visited various public Podcast Directories (e.g. iPodder) and achieved a list of about 6,000 Podcast pages. This was numerically greater than the amount claimed to be referenced by any of the main Podcast Directories.

Some sites such as Odeo[10] are making the Podcast crawling process easier by providing multi-format metadata about both Podcasts and site members. Odeo has made the Podcast categorisation process easier by allowing users to both create and subscribe to Podcasts that are tagged as belonging to a certain category. The site also provide links to RSS 2.0 metadata / M3U files corresponding to both a particular Podcast series and to the various series that a particular member has subscribed to.

A more practical means of achieving new feeds would be to offer user submission options, whereby creators of new broadcasts can issue their feed address to the system and a crawler could pick it up. Also, some of directories, including the iPodder directory offer a 'New Today' listing. By monitoring such listings, new Podcasts could easily be located.

Once a set of url's were achieved, they were crawled and cached. Podcast feeds should not be neglected for more than a few hours however, as new shows are continually being added and local versions of the data would become obsolete. Currently, the crawler is run intermittently however, it is soon planned to have the crawler run continually. Also, a feature supported by sites such as Odeo includes a ping listener, where users can ping the site as notification that their feed has been updated.

5 Multimedia Search Engine

Having acquired or created metadata in RDF relating to multimedia entities, existing Semantic Web applications can be employed to provide various services relating to multimedia retrieval. Indeed, we have created an example of such an application called "Podcast Pinpointer", a Podcast search engine employing an existing Semantic Web tool: 'YARS' [2], or Yet Another RDF Store, which is an optimised database of RDF data. It stores data in N3 format, and features an N3QL[11] query interface.

[10] http://www.odeo.com/

[11] http://www.w3.org/DesignIssues/N3QL.html

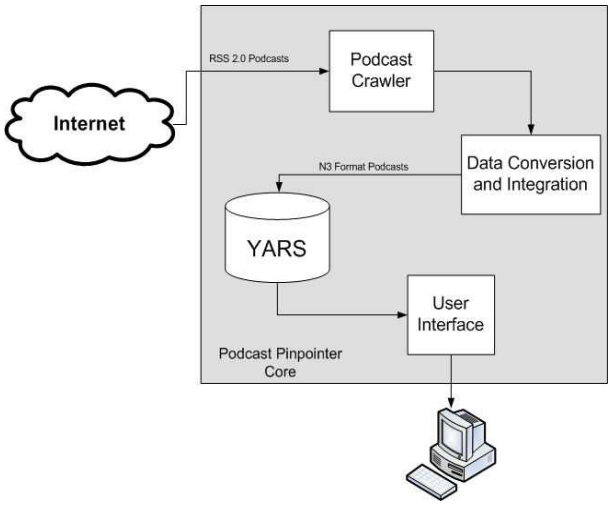

Figure 3: The architecture of the current Podcast Pinpointer application.

By converting Podcast feeds to RDF/XML, and in the future, by creating RDF/XML descriptions of more general multimedia, such data can be stored with ease within YARS. This provides a central repository for data which can then be queried. YARS features advanced querying features, in line with the structured dataset it stores. It also features keyword search capabilities. On top of the N3QL search capabilities, a user interface is built with advanced searches including keywords, multimedia types, acceptable range of duration, creator, acceptable file size range etc. Presently, a simple interface exists which work on keyword searches, but in the future, such advanced searches would be trivial to implement within the defined architecture.

6. Conclusions

This paper has described methods for representing Podcast information on the Semantic Web, beginning with the development of an ontology to represent Podcast metadata (and more generally, information on other multimedia audio and video files) using RDF. A Podcast Pinpointer application has been created to aid in the intelligent search and retrieval of relevant Podcasts by combining Podcast metadata with other data that already exists in the Semantic Web framework. As such, this prototype demonstrates interesting possibilities for the use of audio and video data in future Semantic Web multimedia applications.

Acknowledgements

We thank Matteo Magni for comments on draft versions of the paper.

References

[1] T. Berners Lee, “Semantic Web Road Map”, http://www.w3.org/DesignIssues/Semantic.html, September 1998.

[2] A. Harth, S. Decker. “Optimized Index Structures for Querying RDF from the Web”, Accepted for the 3rd Latin American Web Congress, Argentina, October 2005.

[3] W. Minker. “Semantic Analysis for Automatic Spoken Language Translation and Information Retrieval”, European-Japanese Conference on Information Modelling and Knowledge Bases (EJC), Iwate, Japan, May 1999.