

THE MUSIC ONTOLOGY

Yves Raimond, Samer Abdallah, Mark Sandler
Centre for Digital Music
Queen Mary, University of London

Frederick Giasson
Zitgist LLC

ABSTRACT

In this paper, we overview some Semantic Web technologies and describe the Music Ontology: a formal framework for dealing with music-related information on the Semantic Web, including *editorial*, *cultural* and *acoustic* information. We detail how this ontology can act as a grounding for more domain-specific knowledge representation. In addition, we describe current projects involving the Music Ontology and interlinked repositories of music-related knowledge.

1 INTRODUCTION

Information management is becoming an important part of multimedia related technologies, ranging from the management of personal collections to the construction of large, distributed, ‘semantic’ databases. The latter can be addressed, to some extent, using ‘Semantic Web’ technologies, allowing to create a machine understandable *web of data*. Such a web of data embeds interlinked *ontologies*, giving formal specification of a set of concepts and relationships relevant for describing a set of objects in a given domain (music, people, places, etc.)—the data is linked to its meaning. We can therefore gather in a unique distributed knowledge environment heterogeneous sources of music-related information. These include manual annotations, editorial information, social information, and automatically generated representations.

In this paper, we first give in §2 a quick overview of Semantic Web technologies, and particularly some of them, able to create a web of data. We then focus in §3 on the ‘Music Ontology’, trying to create a formal modular framework for describing music-related information on the Semantic Web. We first focus on describing the workflow beginning with the creation of a musical work to its release on a particular record. Then, we describe how the Music Ontology connects with other ontologies available on the Semantic Web. We also describe in §4 several projects using this ontology in order to publish and interlink music-related information. The aim of these projects is to create a large, distributed, machine-understandable cultural environment dealing with all sorts of concepts related to music production, music consumption, music recommendation, etc.

2 TOWARDS A SEMANTIC WEB OF DATA

In this section, we first describe how Semantic Web technologies help to create a distributed, machine-understandable, web of data. We then overview different user agents able to access this knowledge.

2.1 RDF, dereferencable URIs, HTTP and linked data

RDF (Resource Description Framework [1]) allows the description of resources, by expressing statements about them in the form of *triples*: subject, predicate and object. Resources are identified by a URI (Unique Resource Identifier). An RDF file made available on the web will express a set of triples, which can be interpreted as a graph through the use of these identifiers.

Now, if we use HTTP identifiers, we can add a look up mechanism to the resources mentioned in such a graph. *Dereferencing* such identifiers can then provide useful information about the corresponding resources, as well as links (through the use of some predicate) to other relevant identifiers. This allows an agent interacting with this environment to discover more information, as they browse through *linked data* [2].

For example, a URI *A* might represent a music track. When dereferencing this URI, we should access a RDF description of this track. In this description, we might have a statement such as ‘*A* is on album *B*’, where *B* is a URI representing an album. We can then dereference *B* in order to have more information about this album (release date, cover, etc.).

2.2 Using the web of data

A user agent on the web of data makes use of available resources, links and dereferencing mechanisms in order to dynamically access available data: it can be a crawler, a browser (such as the Tabulator [3]), a ‘client’ (looking programmatically for a particular information¹), etc.

A typical agent holds a cache of RDF triples, and when looking for more information about a particular resource, it dereferences its URI, therefore appending more triples to its RDF cache (an algorithm to implement such a mechanism is described in [3]).

Moreover, some triples in this cache can represent an ontology—a formal description of concepts and relation-

¹such as <http://sites.wiwiwiss.fu-berlin.de/suhl/bizer/ng4j/semwebclient/>

ships in a domain. Using such statements, we can add a reasoning process, such as Description Logic reasoning [4] to this cache. We can therefore check whether our cache has been made inconsistent by dereferencing a particular URI, which may have lead to asserting statements that can't be true according to what we already know. We can also merge resources together, by using functional or inverse-functional properties, or equivalence relationships (such as *owl:sameAs*, meaning 'this resource is the same as that resource'), and we can infer new statements from existing ones, using for example transitive properties, defined class, symmetric properties, etc.

Trust can also be handled within such an agent, by using a Named Graph [5] approach. Each RDF graph within the agent's cache can indeed be associated with an URI, corresponding to its actual provenance. We can then choose to trust, or not, a particular data source.

3 THE MUSIC ONTOLOGY

In this section, we describe the Music Ontology [6], and the main concepts and relationships it defines. We first overview the main ontologies we are using: the Timeline and the Event ontologies, Then, we describe the core of the Music Ontology, and the work flow (from a musical work to its release on a record) it is able to express. We also explain how the Music Ontology can address a wide range of use cases, by dividing itself in several levels of expressiveness. Finally, we detail the different anchor points this ontology provides — allowing to plug onto it more domain-specific ontologies. This leads to a flexible knowledge representation framework for dealing with all sorts of music-related information on the Semantic Web.

3.1 The Timeline ontology

Temporal information is the first thing we want to express when dealing with music-related knowledge. We indeed need to cover a large range of temporal description, from 'this performance happened the 9th of March, 1984' to 'this beat is occurring around sample 32480', through 'the first chorus of this song is here, and is just before the second verse'.

The Music Ontology is therefore built on top of an ontology able to express such temporal information: the Timeline Ontology [7]. This ontology is itself built on top of two concepts defined in OWL-Time [8]: **Interval** and **Instant**, respectively representing time intervals and instants. The Timeline ontology defines another concept: **TimeLine**, representing a coherent backbone for addressing temporal information. For example, we may express the release date of a particular record using the physical time line, and information such as 'from 0 to 3 seconds on that particular track' using the time line backing an audio signal.

The original Timeline Ontology is written in the OWL-DL sublanguage [4] of the Ontology Web Language (OWL [9]). In this ontology, a single time line may have

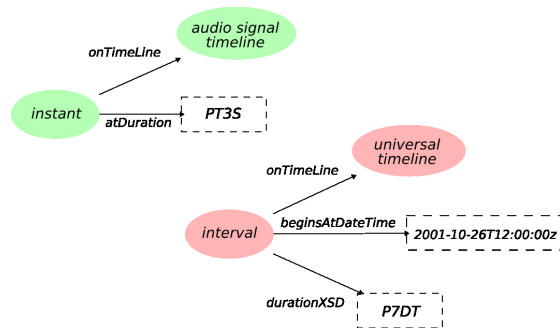


Figure 1. Describing an instant on an audio signal time line (at 3 seconds) and an interval on the universal time line (7 days starting on 26 October 2001, 12:00 UTC)

several coordinate systems allowing to address time points and intervals on it. The ontology also defines a way to relate two time lines together, through the use of time line maps (**TimeLineMap**) — for example, we can use it to express the relationship between the continuous time line supporting an analog signal and the discrete time line supporting the sampled signal. However, we introduce several simplification of this ontology when wrapping it into the Music Ontology. We define one canonical coordinate system per type of time line. We also consider the existence of two main types of time lines: *physical* and *relative* (the time line of a track, for example). The first one is addressed through the XML schema² *dateTime* datatype, and the second one through the *duration* datatype (duration from the beginning of the time line to the point we want to address) in the *continuous* case and the *int* datatype in the *discrete* case.

In fig. 1, we show how an instant on an audio signal time line and an interval on the universal time line can be represented using the Timeline ontology.

3.2 The Event ontology

The music production process involves physical events that occur at a certain place and time and that can involve the participation of a number of physical objects both animate and inanimate. Such events include performances, involving some musicians and their instruments.

The event representation we have adopted is based on the *token-reification* [10] approach. We consider an event occurrence as a first class object or 'token', acting like a hook for additional information pertaining to the event. Regarding the ontological status of event tokens, we consider them as being the way by which cognitive agents classify arbitrary regions of space-time. Our definition of an event is broad enough to include sounds (an acoustic field defined over some space-time region), performances, compositions, and even transduction and recording to produce a digital signal. It is also broad enough to include 'acts of classification' by artificial cognitive agents—this definition will therefore be used in the Feature Ontology described in § 3.5.5.

² <http://www.w3.org/TR/xmlschema-2/>

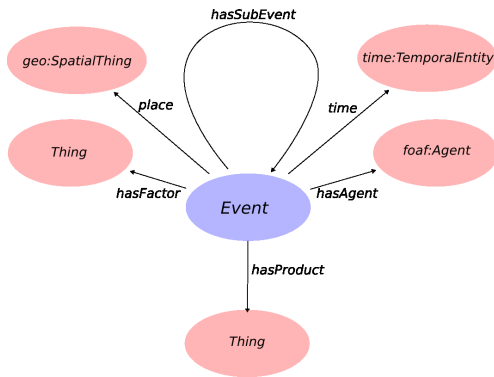


Figure 2. Overview of the Event ontology

We define an **Event** concept, having a number of *factors* (such as a musical instrument, for example), *agents* (such as a particular performer) and *products* (such as the physical sound that a performance produces). This concept can be linked to a particular place through the predicate *event:place* (linking the Event ontology to the Geonames ontology) and to a particular time through *event:time*, linking the Event ontology to the Timeline ontology specified in the previous section.

We also consider the existence of *sub-events* to represent information about complex events in a structured and non-ambiguous way. A complex event, perhaps involving many agents and factors, can be broken into simpler sub-events, each of which can carry part of the information pertaining to the complex whole. For example, a group performance can be described in more detail by considering a number of parallel sub-events, each of which representing the participation of one performer using one musical instrument.

This leads to the ontology defined in [11]. An event concept and its relations to other concepts can be depicted as in fig. 2.

3.3 Music production specific concepts

The Music Ontology is built on top of the Timeline ontology and the Event ontology, as well as the Functional Requirements for Bibliographic Records ontology (FRBR [12]), mainly used for its concept of **Work** (an abstract, distinct, artistic creation), **Manifestation** (physical embodiment, like a record, for example), and **Item** (a single exemplar of such a manifestation, like a particular vinyl). However, we do not directly reuse **Expression** (realization of a **Work**), as a complex workflow can happen during the production process (involving for example an arrangement, a performance, a recording and a mastering), and we want to be able to cover it. We also use the Friend-of-a-friend ontology (FOAF [13]), and its concepts of **Person** and **Group**.

We define a number of music-specific concepts, on top of these three ontologies. First, on top of FRBR, we define **MusicalWork**—an abstract musical creation (such as Franz Schubert’s Trout quintet), **MusicalManifestation**

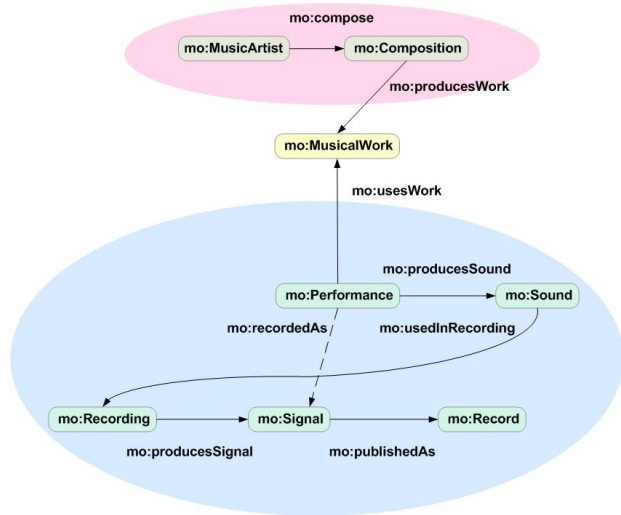


Figure 3. Describing a music production work flow using the Music Ontology

(which can be a **Record** or a **Track** among others), and **MusicalItem**, which can be a **Stream**, a particular CD or a particular vinyl, etc.

On top of the FOAF ontology, we define **MusicArtist** and **MusicGroup** (note that these particular concepts can be considered as defined classes—any person contributing to a musical event can be inferred as being a **MusicArtist**).

On top of the Event ontology, we also define a number of concepts, relative to the music creation work flow. **Composition** deals with the creation of a **MusicalWork**. **Arrangement** deals with an arrangement of a **MusicalWork** and can have as a factor a **MusicalWork**, as an agent an **Arranger** (which can also be considered as a defined class, as any *role* in this ontology) and as a product a **Score**. **Performance** denotes a particular Performance, and can have as factors a **MusicalWork** and a **Score**, a number of musical instruments, equipments, and as agents a number of musicians, sound engineers, conductors, listeners, etc. We can use the event decomposition process defined in § 3.2 and apply it to **Performance** events in order to express things such as ‘this musician was playing this instrument at that particular time’. A **Performance** can have as a product another event: **Sound** — a physical sound. This sound may itself be a factor of a **Recording**, which may produce a **Signal**. This **Signal** can then be *published* as a **MusicalManifestation**. This leads to a work flow such as the one depicted in fig. 3.

We also define a number of other concepts, also linked to the Event ontology. For example, we define **Festival**, which can act as a wrapper for a set of performances, using event decomposition.

3.4 Levels of expressiveness

The main goal of the Music Ontology is to be really flexible. Therefore, it should allow a “blogger” to put online a recording of a concert he attended the day before, or a

musicologist to express complex tonality modulations of a Jazz piece.

In order to allow such a large range of granularity when using the ontology, we divided it in several levels of expressiveness:

- The first level only deals with purely editorial information. We can use it to express ‘*this* track was on *this* particular album and *this* compilation and was created by *that* artist’.
- The second level introduces the **Event** concept. We can use it to describe a work flow involving the composition of a musical work, an arrangement of this work, a performance of this arrangement and a recording of this performance. For example, this level of expressiveness can be used to state: ‘I attended a concert last night, which I recorded using my cell phone, and *here* is the corresponding audio stream’.
- The third level introduces event decomposition, as described in § 3.2. It allows to express things such as ‘in *this* performance, *this* key was played at *this* particular time by *this* person, who was playing the piano’.

Concrete examples of these three levels of expressiveness can be found in [6], along with best practice guidelines.

3.5 Extensions of the ontology

In the previous section, we in fact described a number of anchor points — these different sub classes of **Event** act as a hook for more domain specific information. We now describe available and possible extensions of this ontology.

3.5.1 The Key ontology

One of the first extension of the ontology is the Key ontology [14]. This ontology allows to describe (as a factor of a **Performance**, for example) a particular musical **Key** used. Using event decomposition, we can also express key changes.

3.5.2 The Instrument taxonomy

Another extension of the music ontology is the Instrument Taxonomy [15], extracted from the Musicbrainz instrument taxonomy and expressed using the Simple Knowledge Organisation System (SKOS [16]).

3.5.3 The Genre taxonomy

We also consider using the dbpedia³ RDF repository, exposing information automatically extracted from

³ see <http://dbpedia.org/>

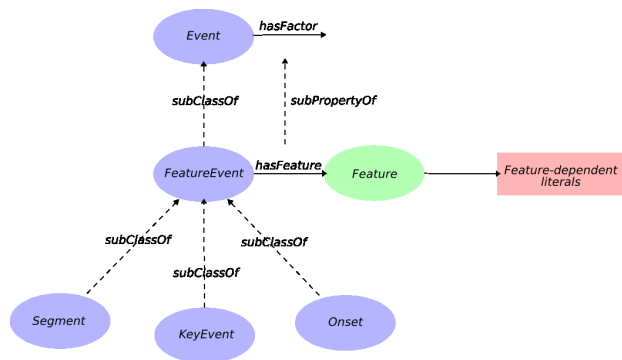


Figure 4. The Feature Ontology, and some of the feature events it defines

Wikipedia for describing musical genres and the relationships between these genres. Such genres can then be used, for example, as factors of some performances.

3.5.4 Social networking information

As the Music Ontology is also based on FOAF, we consider using this ontology in order to express tastes of music consumers, as described in [17], and relationships between such consumers (which may be based on these tastes). This provides a grounding for collaborative filtering applications.

3.5.5 The Features ontology

The feature ontology [18] aims at creating a generic framework for expressing features of audio signals (Mel Frequency Cepstral Coefficients, chromagram, onsets, etc.). It uses the broad definition of the **Event** concept in order to express an artificial classification of a time region, corresponding to a particular feature. Therefore, it defines a sub class of **Event**: **FeatureEvent**, allowing to classify time regions corresponding to features. **FeatureEvent** may have a number of **Feature** factors, representing a particular feature, such as a chromagram or a key. This ontology can be represented as in fig. 4.

3.5.6 Further possible developments

We can also imagine further improvements to the ontology. For example, we might plug an ontology dealing with sound cognition under the **Sound** class, or one dealing with studio recording equipment under the **Recording** one. We might even plug an ontology dealing with mixing several signals together at different volume and pan, by creating a **Mixing** event dealing with **Signal** objects, or one dealing with **Sampling** (taking a part of a signal to mix it with other signals). Another improvement could be to create a score ontology under the **Score** concept, allowing to deal with symbolic music notation or abstract composition rules. We are also developing a chord ontology, using the symbolic representation defined in [19].

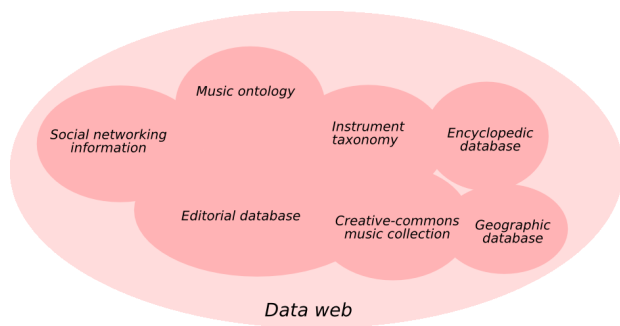


Figure 5. Interlinking music-related information on the web of data

3.6 A flexible knowledge representation framework

Now, using the Music Ontology along with other ontologies and taxonomies available on the Semantic Web, we are able to express either musical ‘metadata’ (editorial information) in a really detailed way, ‘cultural’ metadata (in the form of musical genres and social networking information) and content-based information. We cover the three types of music-related information described in [20].

4 TOWARDS A MUSIC-RELATED WEB OF DATA

In this section, we detail some projects aiming at publishing and interlinking several repositories of music-related data (as depicted in fig. 5), using the Music Ontology and the technologies described in § 2.

4.1 Linking Open Data on the Semantic Web

The *open data* movement aims at making data freely available to everyone. Such data sources include Musicbrainz, Wikipedia, Wikibooks, Geonames, Wordnet, DBLP Bibliography, Magnatune, Jamendo, and many others. The ‘Linking Open Data on the Semantic Web’ community project⁴ of the W3C Semantic Web Education and Outreach group aims at making available such data sources on the web as RDF. Moreover, it also aims to create RDF links between them, in order to make user agents such as the ones described in § 2.2 able to jump from one of these data sources to another.

As an example of such a linking, we may provide information about a festival happening in Montreal, Canada on 28 June 2007. We know that a geographical location database called Geonames exists with 6,4 millions descriptions of locations and that is expressed using RDF and their own geographic ontology. We would like to extend the information we provide about that festival with the full description of its geographical location. Therefore, we can link our **Festival** instance using the *event:place* property to its geographical location resource in Geonames. A user agent crawling the web of data can then jump from our knowledge base to the Geonames one,

⁴ <http://esw.w3.org/topic/SweoIG/TaskForces/CommunityProjects/LinkingOpenData>

by following this *link*, and get detailed information about the place where the festival is happening.

4.2 Current music-related projects

The Zitgist project⁵ converted the Musicbrainz metadata repository in RDF, using the Music Ontology. Every resource in this repository is dereferencable, and it is therefore really easy to link a particular track to relevant metadata (by just stating ‘*this* track is the same as *that* track in the Musicbrainz RDF dump’). The DBTune⁶ project aims at publishing and interlinking several Creative Commons music repositories, in order to provide URIs and allow everyone to share annotations, features, etc. on the available tracks. So far, the Magnatune and the Jamendo repositories have been published. Moreover, DBTune provides links to the dbpedia data source, as well as the Geonames and the Musicbrainz ones.

Other projects are using the Music Ontology. Foafing-the-music [17] is describing user profiles using FOAF and relevant tracks using the Music Ontology. In order to do so, it implements a small web-service to read the ID3 tags from a MP3 file. Then, it uses that information to query other web services such as Musicbrainz and Amazon to aggregate information about the MP3. Finally the web service sends back all that information, described in RDF using the Music Ontology, to the user. The EASA-IER⁷ European project is using the Music Ontology as a knowledge representation foundation, and provides an interface to produce instance data from an audio archive. This project is also developing several new ontology modules, mainly to handle speech feature extraction and links from relevant resources to other medias (video of a performance, scan of a musical score, etc.). The OMRAS2⁸ project is also using the Music Ontology as a knowledge representation framework.

5 CONCLUSIONS AND FUTURE WORK

In this paper, we described some technologies allowing to create a semantic web of data: a distributed, shared, trusted, and machine-understandable knowledge environment. Then, we described the Music Ontology, providing a modular knowledge representation framework for dealing with music-related information in such an environment. Using this ontology along with other ontologies and taxonomies available on the Semantic Web, we are able to express either editorial information in a really detailed way, ‘cultural’ metadata (in the form of musical genres and social networking information) and content-based information. We then described several projects aiming at publishing and interlinking large amounts of information on the web of data, using this ontology. This leads to the creation of a large music-related knowledge environment.

⁵ <http://zitgist.com/>

⁶ <http://purl.org/dbtune/>

⁷ see <http://easaier.org/>

⁸ see <http://omras2.com/>

This knowledge environment can then be interpreted by a large range of applications: from *semantic workspaces* for music analysis researchers [21] to information retrieval tools (such as music recommenders), making use of either ‘metadata’, collaborative or content-based filtering.

Further work includes more publishing and interlinking of data sources, in order to bootstrap the Semantic Web as a whole by demonstrating that it provides the necessary technologies to create a large machine understandable cultural environment. We also aim to develop and manage the quickly growing Music Ontology community, as well as related projects: tools, published data and extensions of the ontology.

6 ACKNOWLEDGMENTS

The authors acknowledge the support of both the Centre For Digital Music and the Department of Computer Science at Queen Mary University of London for the studentship for Yves Raimond. This work has been partially supported by the EPSRC-funded ICT project OMRAS-2 (EP/E017614/1), and by the European Community 6th Framework Programme project EASAIER (IST-033902).

7 REFERENCES

- [1] O. Lassila and R. Swick, “Resource description framework (RDF) model and syntax specification,” 1998. [Online]. Available: citeseer.ist.psu.edu/article/lassila98resource.html
- [2] T. Berners-Lee, “Linked data,” World wide web design issues, July 2006. [Online]. Available: <http://www.w3.org/DesignIssues/LinkedData.html>
- [3] T. Berners-Lee, Y. Chen, L. Chilton, D. Connolly, R. Dhanaraj, J. Hollenbach, A. Lerer, and D. Sheets, “Tabulator: Exploring and analyzing linked data on the semantic web,” in *Proceedings of the 3rd International Semantic Web User Interaction Workshop*, 2006. [Online]. Available: <http://swui.semanticweb.org/swui06/papers/Berners-Lee/Berners-Lee.pdf>
- [4] F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, and P. F. Patel-Schneider, Eds., *The Description Logic Handbook: Theory, Implementation and Application*. Cambridge University Press, 2002.
- [5] J. J. Carroll, C. Bizer, P. Hayes, and P. Stickler, “Named graphs,” *Journal of Web Semantics*, 2005.
- [6] F. Giasson and Y. Raimond, “Music ontology specification,” Working draft, February 2007. [Online]. Available: <http://www.musicontology.com/>
- [7] Y. Raimond and S. A. Abdallah, “The timeline ontology,” OWL-DL ontology, 2006. [Online]. Available: <http://purl.org/NET/c4dm/timeline.owl>
- [8] J. R. Hobbs and F. Pan, “Time ontology in owl,” Working draft, September 2006. [Online]. Available: <http://www.w3.org/TR/owl-time/>
- [9] D. L. McGuinness and F. van Harmelen, “Owl web ontology language: Overview,” World Wide Web Consortium,” Working Draft, March 2003. [Online]. Available: <http://www.w3.org/TR/2003/WD-owl-features-20030331/>
- [10] M. P. Shanahan, “The event calculus explained,” in *Artificial Intelligence Today, Lecture Notes in AI no. 1600*, M. J. Woolridge and M. Veloso, Eds. Springer, 1999, pp. 409–430.
- [11] Y. Raimond and S. A. Abdallah, “The event ontology,” OWL-DL ontology, 2006. [Online]. Available: <http://purl.org/NET/c4dm/event.owl>
- [12] I. Davis and R. Newman, “Expression of core frbr concepts in rdf,” Working draft, 2005. [Online]. Available: <http://vocab.org/frbr/core>
- [13] D. Brickley and L. Miller, “Foaf vocabulary specification,” Working draft, 2005. [Online]. Available: <http://xmlns.com/foaf/0.1/>
- [14] Y. Raimond, “The key ontology,” Working draft, 2006. [Online]. Available: <http://purl.org/net/c4dm/keys.owl>
- [15] I. Herman, “Musicbrainz instrument taxonomy in skos,” Working draft, 2007. [Online]. Available: <http://purl.org/ontology/mo/mit/>
- [16] A. Miles and D. Brickley, “Skos core vocabulary specification,” Working draft, November 2005. [Online]. Available: <http://www.w3.org/TR/2005/WD-swbp-skos-core-spec-20051102/>
- [17] O. Celma, M. Ramirez, and P. Herrera, “Foafing the music: A music recommendation system based on rss feeds and user preferences,” in *International Conference on Music Information Retrieval*, 2005.
- [18] Y. Raimond, “The musical features ontology,” Working draft, February 2007. [Online]. Available: <http://purl.org/ontology/mo/mf/>
- [19] C. Harte, M. Sandler, S. Abdallah, and E. Gomez, “Symbolic representation of musical chords: A proposed syntax for text annotations,” in *Proceedings of the International Conference on Music Information Retrieval*, 2005.
- [20] F. Pachet, *Knowledge Management and Musical Metadata*. Idea Group, 2005.
- [21] Y. Raimond, S. Abdallah, M. Sandler, and M. Lalmas, “A scalable framework for multimedia knowledge management,” in *Proceedings of the 1st conference on Semantic and Digital Media Technologies*, ser. Lecture Notes in Computer Science, Springer, Ed., December 2006.