Semantically Connected Web Resources with MPEG-21

H. Castro[1], M. T. Andrade[1], F. Almeida[1], G. Tropea[2], N. Blefari Melazzi[2], A. S. Mousas[3], D. I. Kaklamani[3], L. Chiariglione[4], and A. Difino[4]

[1]Faculty of Engineering of Oporto University, INESC TEC, Porto, Portugal
[2]UdR Roma Tor Vergata, CNIT Rome, Italy
[3]School of Electrical and Computer Engineering, NTUA, Athens, Greece
[4]CEDEO.net, Via Borgionera, 103 - I-10040 Villar Dora (TO), Italy

Abstract. The Web is rapidly becoming the prime medium for human socialization. The resources that enable that process (social web sites, blogs, media objects, etc.) present growing complexity and, collectively, weave an ever more intricate web of relationships. Current technology for declaring those relationships is predominantly implicit, ambiguous and semantically poor. As a consequence, their automatic assessment is complex and error prone, preventing the satisfaction of users' needs such as effective semantic searches. To address these limitations, whilst enabling the explicit declaration of semantically unambiguous relationships between digital resources, a solution employing structured semantic descriptors and ontologies was conceived, based on MPEG-21. This paper explains the functioning of the devised mechanism, and goes beyond that, into the definition of two novel employment venues for it, at the service of two real-world usage scenarios. These demonstrate the mechanism's added value as a powerful alternative for the semantically aware interconnection of web resources, and highlight the increased QoE that said mechanism enables.

Keywords: Semantic Web, relationships, MPEG-21, DID, DII, Cloud, QoE, Content Search, Metadata, RDF/OWL.

1   Introduction

The Internet was developed based on a host-centric perspective. Still, it is now evolving into a "content-centric" operating structure [2]. Within this new operational paradigm, the key elements are no longer the "hosts", but rather the exchanged data and services [1]. Availability of information concerning the relationships between such elements is fundamental to foster the successful establishment of the emerging new paradigm and of the added QoE that it brings about to Internet usage.

However, current technology enables the expression of, only, very rigid and semantically ambiguous relationships, between Internet's informational resources. For instance, the standard way of declaring a relationship between two digital documents, on

the web, is to insert, in one of them, the other one's URL. In HTML, for instance, this is done with the *<a>* or *<link>* tags.

New powerful and flexible mechanisms for the unambiguous expression of such relationships would definitely contribute to enable a semantically more precise and more automated manipulation, and navigation, of resources.

In this context, a novel mechanism was developed by the authors, for declaring semantic relationships between digital resources, building upon the MPEG-21 standard. Said mechanism, whilst adopting the MPEG-21 principles by using structured descriptors, expands that technology through a specific use of semantic metadata expressed in RDF/OWL. This novel tool provides the necessary capabilities to enable a full, and flexible, expression of the relational context of MPEG-21 Digital Items, on a domain specific basis. Validating the relevance and innovativeness of the mechanism in scope, is its acceptance as an extension to part 3 of the MPEG-21 standard [8].

This paper explains the mechanism's concept and presents two of its possible novel exploitation venues, thus demonstrating its added value. It is structured as follows: section 2 exposes the growing need for unambiguous expression of inter-resource relationships, highlighting the need for new tools to perform that expression; section 3 presents the MPEG-21 DI concept and exposes its limitations for expressing the mentioned relationships; section 4 presents the mechanism developed to overcome the identified limitations; section 5 presents two novel exploitation venues for the mechanism, with a special focus on its employment in a cloud-based blog support system; section 6 summarizes the overall added value of the mechanism in scope; and section 7 contains the final conclusions.

2 Relationships Expression and Assessment

2.1 State of the Art

Current, successful, human-centric Internet-based initiatives greatly depend on their ability to efficiently exploit the relationships that bind together their user community and the resources they provide. For instance, Facebook collects and exploits friendship or "*Liking*" relationships that users declare; IEEEXplore uses the cited references of science papers that authors supply; and Amazon deduces and exploits relationships between sold goods based on users' purchase patterns. The assessment of the relational context of real and digital resources is thus becoming a concrete necessity.

Generally speaking, two main alternative strategies exist to obtain this type of knowledge:

- implicitly, through information extraction and deduction by artificial intelligence mechanisms;
- explicitly, through the use of mechanisms that allow resource producers to insert semantically-rich inter-resource relationships.

The first strategy has great potential, as it does not depend on the existence of relational metadata or on its quality. It typically employs content-based mechanisms that extract low-level features and infer relationships between analysed content [11]. However this process is still costly and poorly achieved. In fact, artificial reasoning, of human-level reliability, is of dubious achievement and is difficult to verify [22].

The second strategy implies a greater investment at resource production time, for the expression of the resources' relational context. However, after that information is produced, it can be unambiguously and universally interpreted and exploited, through simple means. Currently, it is possible to find some information of this type on tags assigned by users to resources they publish or consume [10]. However, tagging information is far from being reliable and it is highly subjective and context-dependent.

The simplest way of declaring a relationship between two digital documents is to insert in one of them, the other one's URL. In HTML this is possible with the <a> and <link> tags. Such specifications 3 include some support for the definition of the nature/type of the declared relationship. However these provide a very limited and still semantically ambiguous relationship type pool. This is why such information is seldom used, and generally left unexploited by browsers and search engines [12].

The Microformats [14] approach emerged as an evolution to the concept of embedding machine-readable metadata into HTML pages. Recently replaced by similar stronger initiatives, such as the W3C RDFa [15] and the Google-Yahoo-Bing schema.org [16], this trend, when consolidated, will allow machine parsing and understanding of web content.

In parallel, the Semantic Web and Linked Data [13] initiatives, strive to promote usage and supply of structured semantic information, by humans, in a manner that renders it understandable by machines. They suggest representing content in fundamentally more structured ways than HTML. By careful usage and linkage of precise metadata, machine-assisted inference, automated information mashups and semantic search become possible and useful.

In-line with the Linked Data initiative, the OAI-ORE standard [21] provides the means for the declaration of inter-resource relationships within resource aggregations in a semantically meaningful way.

Although the latter proposals (W3C RDFa, Linked Data and OAI-ORE) start to weave a global, structured metadata mesh to describe real and digital resources and their relationships, their current implementation or employment levels are still at an early stage. Furthermore they are meant as an informational patch to be added to current web pages to provide semantically more precise meaning to their contents. Their employment is not organic to the objects that they are meant to enrich, resulting in, either a profuse intermixing of data and metadata (as in RDFa), or in a loss of self-

containment (as in OAI-ORE), considering that data and metadata objects are to be distributed separately.

It can thus be concluded that there is still a lack of efficient mechanisms that are able to support in a reliable and practical manner, the semantically precise expression and assessment of the relational context of real and digital resources. Such mechanisms should ensure that the diffusion of these semantic relationships is easily done in a coherent way, synchronized to that of the corresponding resources.

2.2 The Way Forward

The first of the above-mentioned strategies imposes greater demands and is unlikelier to obtain superior results than the latter, in the short term.

For that reason and for the fact that the focus of the work here presented, is the explicit semantic linking between self-contained digital objects, we pursue the second strategy (referred in section 2.1), for the acquisition of inter-resource relationship information.

The success of this strategy requires a comprehensive and, simultaneously, versatile tool for the structuring and description of information objects. Such a tool should enable information producers to: 1) perform a simple and semantically precise expression of inter-resource relationships; and 2) bind the descriptive data to the described objects and package and distribute this set in a reliable, transparent and coherent way.

An additional factor that contributes to the success and rapid adoption of any new technology concerns the use of open standards. Standardisation of the developed technologies is normally an added value, especially near the industry, to encourage wider take-up.

Accordingly, the approach that was adopted in this work to attain such a tool whilst taking into consideration the mentioned aspects, was to extend the MPEG-21 standard, taking advantage of domain-specific ontologies expressed using also standardized technologies, notably RDF/OWL. This tool empowers the MPEG-21 to support rich and flexible expression of relationships, between digital objects, within the context of domain specific ontologies of such relationships. This standard was chosen, over others, because, unlike them, it provides a globally coherent and extensive framework to structure and describe information objects, as well as to create self-contained packages.

With the incorporation of our solution into the MPEG-21 standard, information producers will be able to use the Digital Item (DI) concept as a structured container to declare and encapsulate all the media components of their creations, together with relevant metadata, including the digital item's relational context. As we aim at demonstrating in the following sections, this will contribute to the enhancement of the Internet as a whole, as the extended MPEG-21 standard provides useful means to enable a

more semantically aware manipulation of digital resources. Accordingly, before actually presenting the novel inter-resource relationship expression mechanism, a brief description of the MPEG-21 standard together with a detailed view on its shortcomings concerning the capability to describe a DI's relational context, is provided in the next section.

The remainder of the article explains the developed mechanism, which is now part of the MPEG-21 DII standard [9], demonstrating its added value in enabling MPEG-21 to address the needs identified in section 2, and goes beyond that, into the definition of two novel employment venues for it, at the service of two real-world usage scenarios.

3   MPEG-21

3.1   Contextualization

MPEG-21 is a powerful multipurpose content declaration standard that enables the multidimensional definition, structuring and packaging of sets of media resources into metadata rich, self-contained, information objects. These are designated as "Digital Items" (DIs) [5] in MPEG-21 part 2, DID [6]. MPEG-21 DIs, within the context of MPEG-21's DID Language (DIDL), are expressed in XML and carried in Digital Item Declaration (DID) documents. The Digital Item Identification (DII) specification [8] provides the tools to uniquely identify DIs and link them to related information.

We think these capabilities are extremely synergistic with the idea of extending content-centeredness to include "resources" such as real world objects, services and people, and to support the semantically precise description of their relational contexts.

3.2   Limitations of MPEG-21 for Relational Context Definition

The core technologies of the DI are the DID and the DII. However, these tools enable only a rigid, limited and generally implicit description of inter DI relationships.
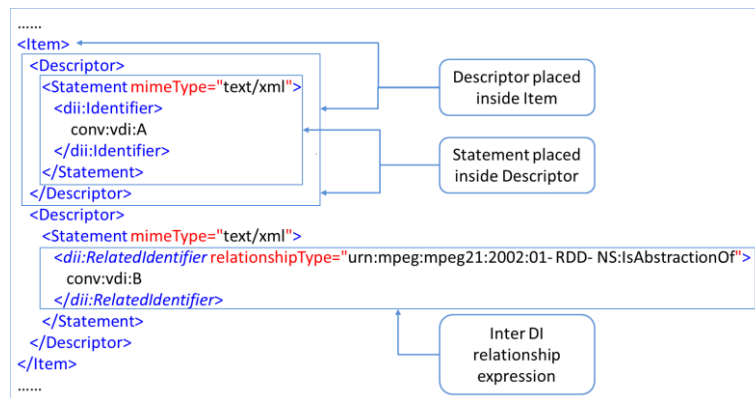

Fig. 1. Expressing Relationships in DID and DII

The DID supports only the implicit expression of containment/composition relationships between DIs by placing XML elements within elements, according to a strict linear hierarchy (as explained at the top of Fig. 1). Alternatively, DID elements may explicitly include other elements employing the XInclude [23] tool. However, XInclude only enables the specification of *inclusion* relationships, and of no other type (such as, for instance, a correction relationship as in *DI_A corrects DI_B*), thus being quite poor at the semantic level.

The original DII specification enabled the explicit expression of a limited set of inter-DI relationships, through the employment of *dii:RelatedIdentifier* elements and their *relationshipType* attribute (as indicated at the bottom of Fig. 1). It was more versatile than the DID, however, it was still insufficient as a fully-fledged tool for the expression of the relationships that normally occur.

The standard was then extended through an amendment [7] to MPEG-21 Part 3 [8], whereby the *dii:RelatedIdentifier* element was semantically enriched with the capability to express the nature of the relationship. Still, the available pool of relationship types is very small and, though it may be extended, such extension requires the registration of the new type with the Registration Authority for ISO/IEC. This is an unpractical scheme for the definition of relationship taxonomies and is also completely divorced from the RDF/OWL universe, which would be the natural venue for the expression of the logical connections in scope.

The tools defined in DID and DII are thus inadequate for the expression of semantically-rich inter-DI relationships. Their use would lead to contrived metadata constructs that would be hard, or impossible, to query through common and widespread technologies such as SPARQL. Furthermore, they do not allow the expression of relationships of higher complexity such as non-binary relationships, or property endowed relationships (e.g. a relationship with an *occurrenceTime* property).

4   A Novel Semantically-rich Resource Relating Mechanism

4.1   Rationale

In light of the conclusions drawn in the previous section, a relating mechanism was developed to allow resource creators to easily express static and explicit inter-resource relationships. It employs RDF/OWL constructs, expressed in XML, thus being aligned with the Linked Data initiative. Resources must however be expressed according to the standardized MPEG-21 DI format, as these constructs have been defined embedded in the MPEG-21 specification, in line with the analysis made in section 2.2. The following sections explain this process and the details of the mechanism.

4.2 Definition

In practical terms, the developed mechanism consists of the addition of a new DII element, named *dii:Relationships*, to the current specification of the MPEG-21 DII standard.

```
<xs:element name="Relationships">                                    a
  <xs:complexType mixed="true">
    <xs:sequence>
      <xs:element
        xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax- ns#"
        ref="rdf:RDF" minOccurs="1" maxOccurs="1"/>
    </xs:sequence>
  </xs:complexType>
</xs:element>
```

```
<rdf:RDF xmlns="mpeg21:corerelationalontology"                       b
  xmlns:mpeg21coreRelOnt="mpeg21:corerelationalontology"
  ...............>

  <owl:Ontology rdf:about="mpeg21:corerelationalontology"/>

  <!-- Classes -->
  <owl:Class rdf:about="mpeg21:corerelationalontology;DigitalItem"/>
  <owl:Class rdf:about="&owl;Thing"/>

  <!-- Object Properties -->
  <owl:ObjectProperty rdf:about="&mpeg21coreRelOnt;interDIRelationship">
    <rdfs:domain rdf:resource="&mpeg21coreRelOnt;DigitalItem"/>
    <rdfs:range rdf:resource="&mpeg21coreRelOnt;DigitalItem"/>
  </owl:ObjectProperty>
</rdf:RDF>
```

```
<Descriptor>                                                         c
  <Statement mimeType="text/xml">
    <dii:Relationships mimeType="application/rdf+xml">
      <rdf:RDF xmlns="individual:example:ontology#"
        xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#">
        <DigitalItem rdf:about="DI_A">
          <rdf:type rdf:resource="&owl;DigitalItem"/>
          <correction rdf:resource="DI_B"/>
        </DigitalItem>
        <DigitalItem rdf:about="DI_B">
          <rdf:type rdf:resource="&owl;DigitalItem"/>
        </DigitalItem>
      </rdf:RDF>
    </dii:Relationships>
  </Statement>
</Descriptor>
```

Fig. 2. MPEG-21 Extension Provisions

This element, whose schema definition is presented in Fig. 2.a, carries the relationship expressing information. It should be contained within a *did:Descriptor*. Each *did:Descriptor* binds to its parent *did:Item* element (which represents a DI), the relationship expressing information.

The *dii:Relationships* element carries an RDF/OWL metadata payload, describing the relationships in which the DI participates with other DIs. Relationships are expressed as RDF *subject-predicate-object* triples. In each triple DIs may be either the subject or the object, and the predicate denotes the relationship between them.

The MPEG-21 Core Ontology for DI Relationships (MPEG-21 CODIR), which is presented in Fig. 2.b, was also defined within the context of the work here presented. It was submitted to MPEG-21 together with the definition of the *dii:Relationships* element and thus, is now part of the same amendment to MPEG-21. It provides the concepts for both the *DigitalItem* and the base *interDIRelationship*.

The actual relationships, to be declared in the description of the relational context of digital resources, should be defined within custom, domain specific, ontologies. These ontologies are defined by adopting and extending the basic ontological model, provided by CODIR. As such, extended ontologies do not have to be defined or registered with any central authority. Instead, they just have to be defined to address domain specific needs and need only to be known within that domain. This independence avoids, for instance, the cumbersomeness involved when extending the pool of relationship types of element *dii:RelatedIdentifier* (as mentioned earlier), which in fact requires the registration of the new type with a Registration Authority.

In agreement with the Linked-Data proposal, the solution here presented:

- uses RDF to express structured data and publish it on the Web;
- uses RDF links to logically connect data from different data sources (interrelate different DIs);
- identifies all resources by means of URIs.

The URIs employed for the identification of resources, are not necessarily HTTP ones, as in the Linked-Data proposal. Still, we believe that this is not an issue but actually, an improvement, as it renders this solution more location independent, and thus, more content-centric oriented.

This solution is also perfectly compatible with MPEG-21 as it represents only the addition of an extra XML element, carrying XML expressed metadata (RDF/OWL), into an XML expressed standard (MPEG-21).

An example of the mechanism's employment is the following: if *DI A* is a correction to *DI B*, then, it may be said *DI A* is logically bound to *DI B* through the relationship of "*correction*". *DI A* is the subject, *DI B* the object and "*correction*" is the predicate of the relationship. Based on, and extending the MPEG-21 CODIR, a domain specific ontology can be defined (not presented in Fig. 2), containing the concept of

correction relationship. Consequently, the explicit relationship between *DI A* and *DI B* may be expressed using the MPEG-21 standard, as presented in Fig. 2.c. The descriptor, presented in Fig. 2.c, would, typically, be located within the subject DI, i.e. *DI A*.

5 Exploitation for Increased QoE

The relevance of the proposed mechanism, for providing increased QoE, may be demonstrated by analysing its use in different situations of multimedia resource consumption, among those that are flourishing on the Web. These scenarios can show the added value that can be obtained, in terms of objective and subjective user satisfaction.

In the next two sections we present two examples of multimedia resource consumption scenarios, where an innovative use is made of the developed mechanism. In them we highlight its relevance for an added QoE when searching, accessing and consuming multimedia content. Greater objective user satisfaction results from the mechanism's ability to assist the system in performing a more precise and automated resource discovery. This is a direct consequence of the mechanism' capability of enabling a flexible, precise, and system-independent definition of logical relationships between digital objects. Greater subjective user satisfaction results from the increased flexibility, transparency and un-intrusiveness of the process. Users do not need to be aware of the existence of related content – the system automatically notifies them of varying types of resources with immediate interest to them. Accordingly, this leads to better meeting users' natural expectations and to enabling easier interaction with content.

Section 5.1 presents a simple exploitation example within the context of an e-Learning application, whereas section 5.2 presents a more complete exploitation example pertaining to the support of Weblogs (blogs).

5.1 e-Learning System

An advanced, distributed, Publish-Subscribe based, e-Learning platform is maintained by a college to assist its students in their individual learning activities: "*Employing that system Professor Watson publishes an InfoObject (IO) containing the material she lectures in each class, as a Lecture IO. Each such IO is made available for subscription to all students. Tom, a student who subscribed to the course, is automatically notified about the release of new lecture episodes. While watching the latest episode, he comes up with a doubt about a lectured subject. He produces a new Question IO with a question pertaining to the subject he has doubts about, and makes it public. The system then pushes the Question IO to all subscribed students. Mary, another student, knows the answer. She publishes it as a Response IO. Eventually Tom receives her answer*".

e-Learning systems already exist that provide similar functionalities to those described above. However they typically depend on rigid and implicitly defined relationships, which are "hardcoded" into the system's operational procedures. Expanding these solutions to deal with new situations is therefore complex, and often unfeasible without the realization of development work on them.
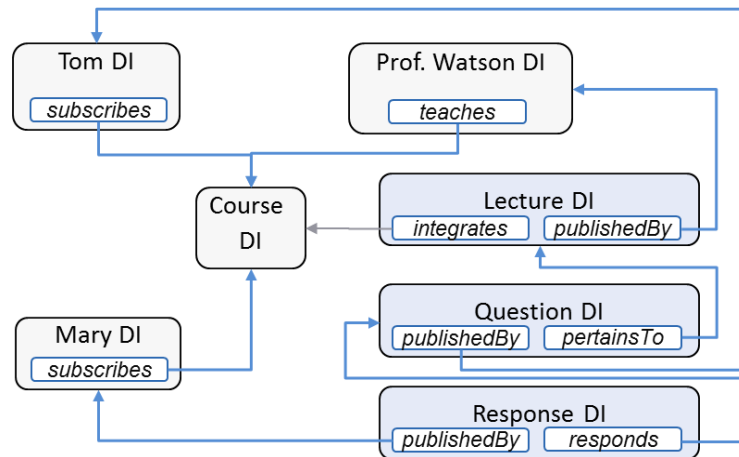
The fact that, in the above presented solution, relational metadata is provided in an explicit way, enables even simple e-Learning systems, which do not resort to Artificial Intelligence (AI) mechanisms, to profit from this knowledge so as to provide a more dynamic and expandable set of services to users.

For instance, in the above-described e-Learning scenario, one may identify that relationships-wise, Tom's *Question IO* has a "*pertainsTo*" relationship, linking it to the *Lecture IO*; or that Mary's *Response IO* has a "*responding*" relationship to Tom's *Question IO*.

Employing MPEG-21, enriched with the mechanism that we have defined, this reality can easily be described and, thus, knowledge about it can also easily be preserved and shared. With this expanded tool it is possible to perform the declaration, packaging and inter-relating of all involved information assets and the identification and inter-relating of all involved real world entities (i.e. people, courses, etc.).

Specifically: all the relevant system entities (courses, students and teachers), are represented by an updateable DI, which identifies them and declares their relevant relationships (e.g. Tom's DI declares that it maintains a "*subscription*" relationship with the course, which is represented by the *course DI*); all IOs are constituted as MEPG-21 DIs (within DID documents). Each such DI contains, besides its actual informational content, the specification of its relationships to the relevant other DIs. (e.g. the *Question IO* DI declares that it maintains a "*pertainsTo*" relationship with *Lecture IO* DI).

The resulting scenario of DIs and their logical relationships is depicted in Fig. 3.

Fig. 3. Augmented e-Learning Relational Context

This way, the reliably and explicitly declared relational context described before, tells an e-Learning system, for instance, that: Tom and Mary are course students, whereas Dr. Watson is the course teacher; specific resources, wrapped up as DIs are the lectures that make up the course; other specific DIs are questions pertaining to the lectures whereas others are responses pertaining to those questions.

Knowledge of this information makes it easier for these systems to manipulate resources and interact with users. It grants them a higher flexibility as they may adapt to new situations by having their ontology updated, and may, therefore, more easily be expanded.

Furthermore, more powerful systems that do resort to AI tools, may build on this information (which should, in principle, be error-free), to infer additional knowledge, establishing other types of relationships and thus offer users an enhanced service – they may, for instance, infer that *Question IOs* should be directed to the corresponding course teacher and that, if said teacher is replaced, they should be directed to the replacing teacher.

This simple, pertinent and semantically clear expression of relationships thus enables systems, as the one described above, to accurately and automatically detect interentity and inter-resource relationships and to operate on that knowledge, by means of programmed or intelligently derived directives, to infer subtle, semantic aspects of the reality in scope that would otherwise elude them. These systems are thus capable of a much more flexible adaptation to reality and to its changing demands and may, more easily, employ intelligent means to objectively increase the relevance and accuracy of the resources they bring to user attention. This enables them to diminish the effort required from users to find, or filter through irrelevant information until reaching, useful content, thus increasing subjective user QoE.
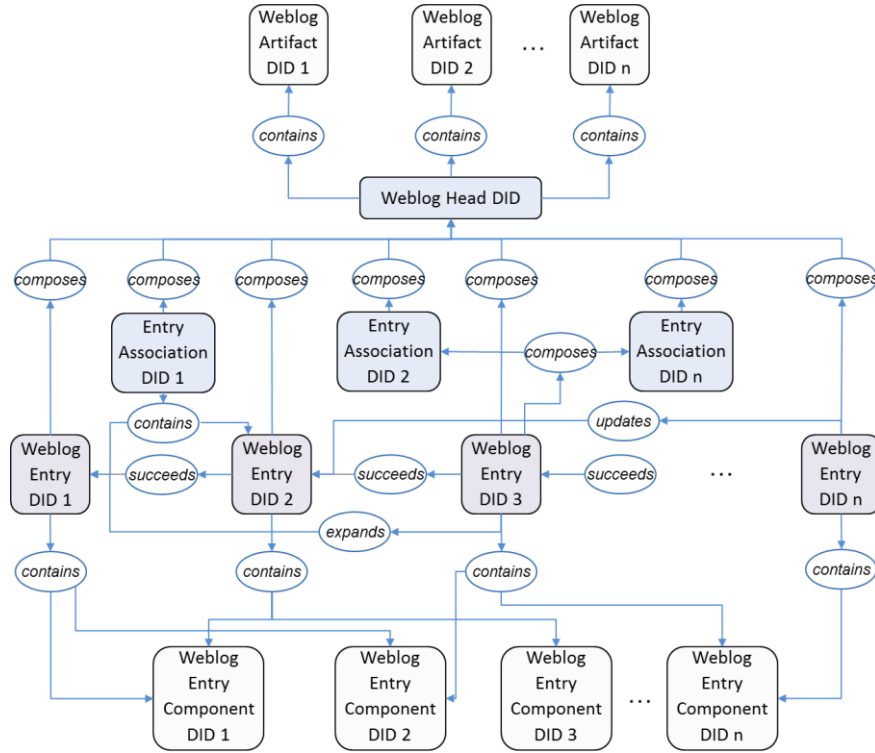
5.2   Cloud Based Weblog Hosting

Introduction.
Using the MPEG-21 specification, extended with our mechanism, it is possible to declare and package digital resources together with metadata defining their relational context. This information does not need to be extracted from the resources, nor stored and managed by a central repository. It is created, by resource producers, on a per DI basis and distributed alongside the resource/DI. This information is appropriate for distributed handling and thus, can be used to semantically empower cloud-based systems such as a cloud-based weblog hosting service. The Blog-Cloud is a hypothetical system delivering such a service.

Blog-Cloud performs a completely distributed storage and management of resources. This is achieved in a location independent way (hence, content-centric oriented), taking advantage of the semantically precise declarations of inter resource relationships (instead of http links). It can thus more efficiently and intelligently manage and distribute the resources that make up the weblogs.

Weblog Data Structure.

The content of a Blog-Cloud weblog is divided through multiple DIDs of different types. The DIs, declared by these DIDs, carry different types of payloads and exhibit various types of relationships amongst themselves, as depicted in Fig. 4 and described below.

Fig. 4. Data Structure of a Blog-Cloud Weblog

The employed types of DIDs and the relationship declarations that they comprise are the following:

- *WH-DID*, the Weblog Head DID, which is the main DID. The *WH-DID* represents the whole blog and contains all the metadata (authors' identification, blog's main subjects, etc.) that pertains to the entire blog. It declares "*containment*" relationships to the Weblog Artefact DIDs (*WArt-DID*);
- *WArt-DID*, representing and containing a specific artefact (graphical, audio, audio-visual, etc.), employed in the graphical interface(s) of the weblog;
- *WE-DID*, the Weblog Entry DID, representing a single user (author) entry into the blog and contains only relevant metadata. It declares "*containment*" relationships towards the Weblog Entry Component DIDs (*WEC-DID*), which carry the actual text, image or video components of the entry. *WE-DIDs* also declare a "*composition*" relationship towards their corresponding *WH-DID* and may also declare a

number of other relationships towards each other. They may declare, for instance, that *WE-DID2 succeeds WE-DID1* or that *WE-DIDn updates WE-DID2*;

- *WEC-DID*, a Weblog Entry Content DID containing a single media component (text, image or video). It may be the target of "*containment*" relationships by many different *WE-DIDs*;

- *EA-DID*, the Entry Association DID, representing a specific association of blog entries. It may, for instance, represent the group of *WE-DIDs* which were added during a specific day, or that address a specific subject. When an *EA-DID* is created it will declare a "*containment*" relationship towards all the entries that are encompassed by it and a "*composition*" relationship towards the *WH-DID*. Any *WE-DID* that logically belongs to it, which is added after the *EA-DID* has been published, will declare a "*composition*" relationship towards the latter.

All of these DIDs will be independently distributed within Blog-Cloud and likewise independently retrieved by the client side, when necessary.

When addressing relationships, we freely refer to their targets as either DIDs or DIs. Even though they are not equal, given that each of the previously presented DIDs contains only a single DI, the two terms may be applied interchangeably.

Weblog Ontology.

The weblog ontology (depicted in Fig. 5) expands the CODIR ontology to define the concepts corresponding to the identified basic set of relationships between the weblog's composing DIDs. It is defined on an exemplifying basis to demonstrate that there is a real interest in having a precise and explicit description of relationships between resources in this field of application. It is not meant as a final word for an ontology for said field.

System Structure and Operation.

The system is composed by a distributed set of server nodes that cooperatively serve client side requests.

Within the universe of possible system DIDs, one may identify three different information spaces:

- the space composed by the all the DIDs (more accurately, by their identifiers), S_{DID};

- the space composed by all the semantic relationships declared by all the system DIDs, S_{rel};

- the space composed by all the other semantic concepts declared in the semantic characterization metadata contained in the *WEC-DIDs* and *EA-DIDs*, S_{char}.

Each of the server nodes is responsible for the storage and delivery of the information contained in a segment of each of the three spaces. The mapping of responsi-

bility over these segments to specific individual servers is done in accordance with some algorithm known by all servers.

Each server node, (sn_i), will thus be responsible for storing and delivering:

- a group of DIDs ($S_{DID}^{sn_i}$), corresponding to a segment in S_{DID} ;

- a subset of inter-DI relationships information, ($S_{rel}^{sn_i}$), corresponding to a segment in S_{rel} ;

- and a subset of DIs semantic characterization data, ($S_{char}^{sn_i}$), corresponding to a segment in S_{char} .
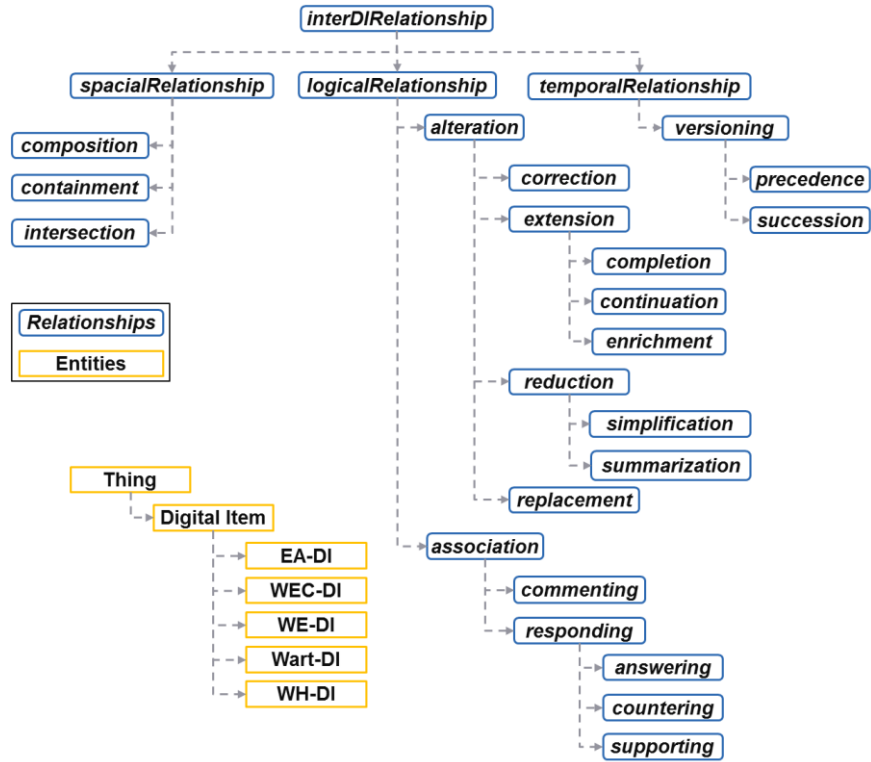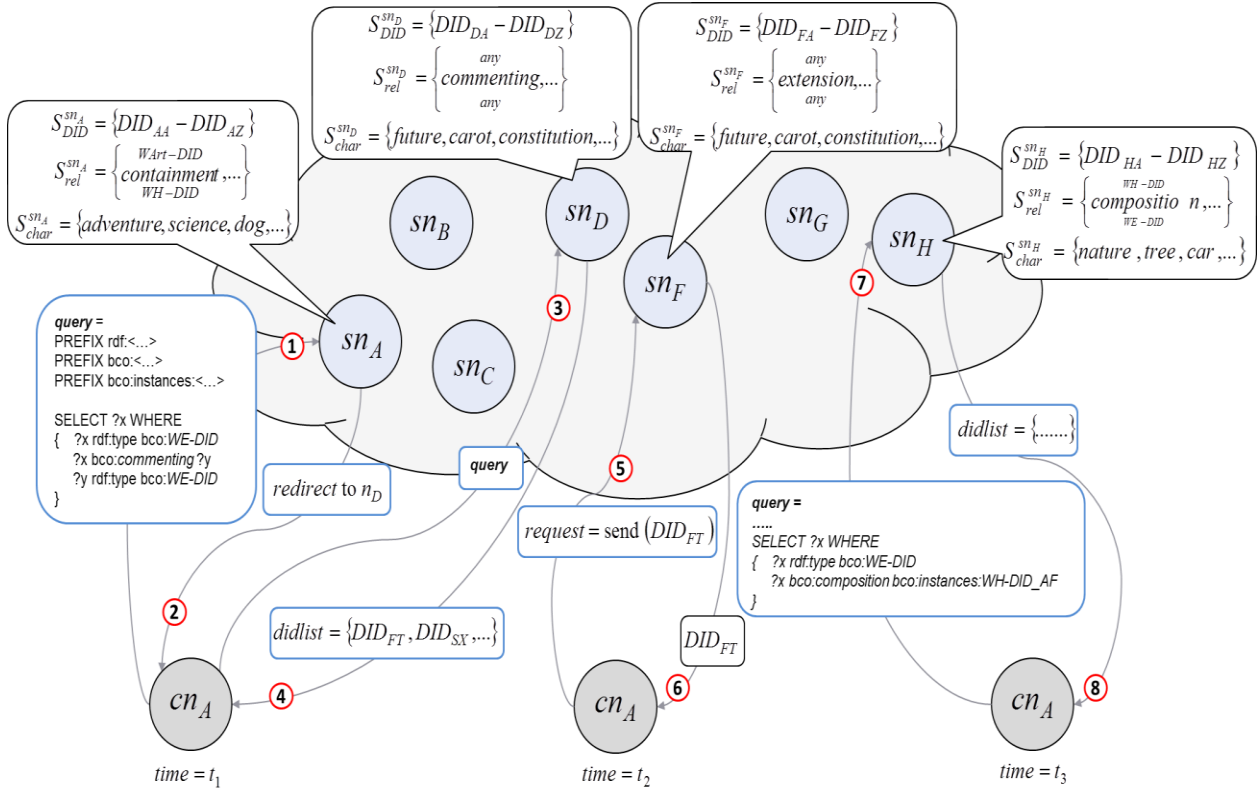

Fig. 5. Weblog Ontology

As all server nodes know the mapping algorithm. They, thus, know which server nodes to contact to obtain specific segments. Client nodes may not know that mapping. They will contact any known server node to obtain the desired information. If the latter is not the adequate one, it will redirect the client to the appropriate server.

Client nodes will interact with server nodes to request such operations as: DID searching (based on various specific criteria) or DID retrieval.

The subsections below present the system's main operations and capabilities, based on generic examples that highlight the possibilities offered by the employment of MPEG-21 DID and relationship declaration.

Content Search.

A client node, (acting on behalf of a user), may perform different types of searches for DIs. The traditional search will be based on the DIs semantic characteristics. That is, searches operating on the S_{char} information space. However, the typical links between documents have been replaced and semantically enhanced by explicitly declared relationships using the extended MPEG-21 standard. Thus, there will also be searches operating on the $S_{rel}^{n_i}$ information space.

Fig. 6. Depiction of the System's Operation

Fig. 6 depicts a search operation of the latter type. In that figure, the top rounded rectangular callouts indicate the segments of the information spaces for which each

server node is responsible. For illustration purposes the S_{DID} space of system DIDs is defined as an interval of DID identifiers, where such identifiers are expressed as two-letters-long foot indexes (e.g. $DID_{FA} - DID_{FZ}$). In this regard, server node sn_d, for instance, is responsible for the following segments of the system's information spaces:

- $S_{DID}^{sn_D} = \{DID_{DA} - DID_{DZ}\}$ – storing all DIDs contained in the set from DID_{DA} to DID_{DZ};

- $S_{rel}^{sn_D} = \left\{ \overset{any}{\underset{any}{commenting}}, ... \right\}$ – storing the declaration of all *comenting* relationships between any two DIDs;

$S_{char}^{sn_D} = \{future, carot, constitution, ...\}$ – storing the mapping of all attributes (keywords) in the set $\{future, carot, constitution, ...\}$, to all the DIDs that they pertain to.

The bottom, rounded, rectangles indicate the content of exchanged messages. Some user, hosted at client node cn_A, issues a request for WE-DIs (blog entries), which are comments to other WE-DIs (regardless of any other aspect). cn_A compiles the adequate SPARQL query and sends it to one of its known server nodes, sn_A (arrow 1). However, $S_{rel}^{sn_A}$ does not cover the targeted "commenting" relationship. It thus responds to cn_A, redirecting it to sn_D (arrow 2). cn_A re-sends the query to sn_D (arrow 3), which undertakes the searching procedure over its $S_{rel}^{sn_A}$, as the latter includes $\overset{any}{commenting}$ (registry of all "*commenting*" relationships from any subject DI to any $\underset{any}{}$ object DI). Information about all WE-DIs satisfying the search criteria, is compiled into a list ($didlist = \{DID_{FT}, DID_{SX}, ...\}$) and sent to cn_A. Finally, cn_A displays the obtained list to the user.

This procedure thus enables a new type of search to be performed over DIs, focusing on the relationships between on-line resources (such as the "*commenting*" relationships between blog entries exemplified above). These added search capabilities objectively enable a much more efficient and precise discovery and traversal of such logical connections. This makes it easier for users to navigate the mesh of posts and comments in a semantically precise manner, spending less time to find their desired information items, and, thus, contributing to a greater subjective user satisfaction and, hence, QoE.

Search Optimization.

Search engines operating on MPEG-21 DIDs, no longer have to "deduce" the types of logical connections binding digital resources, because these are formally specified. Smarter searches, i.e., searches with higher precision and recall (classic evaluation metrics of search techniques from the information retrieval science), may thus be achieved. Exploiting the semantic preciseness of ontologically defined inter-DI relationships and the semantic proximity or distance between them easily attains this. This proximity is expressed by linking two inter-DI relationships (from the same or different ontologies), with the *rdf:subPropertyOf*, *owl:sameAs* or *skos:narrower* connections. For instance the "*correction*" relationship may be an *rdf:subPropertyOf* the "*commenting*" relationship and the "*grammarcorrection*" relationship may be *skos:narrower* than the "*correction*" relationship.

High precision is obtained by searching for DIs bound to other DIs by a specific semantic inter-DI relationship. High recall is obtained through the reasoning over inter-DI relationships that fall within a single ontology or which are interconnected over multiple ontologies.

This way, to increase their precision, search queries for DIs that maintain a relationship r_A with a target DI, DI_A, may be reformulated to search for DIs maintaining other inter-DI relationships with DI_A: these may be relationships from the same ontology, in which case they will have an *rdf:subPropertyOf* connection with r_A; or from different ontologies, in which case they will have an *skos:narrower* connection to r_A.

```
//Original Query
PREFIX rdf: <..................................>
PREFIX bco: <................................>
PREFIX bco:instances <.................>

SELECT ?x WHERE
{   ?x rdf:type bco:WE-DID
    ?x bco:commenting ?y
    ?y rdf:type bco:WE-DID}

//Reformulated Query
PREFIX rdf: <..................................>
PREFIX bco: <................................>
PREFIX bco:instances <..................>

SELECT ?x WHERE
{    ?x rdf:type bco:WE-DID
    ?x bco:responding ?y
    ?y rdf:type bco:WE-DID}
```

Fig. 7 Depiction of Query Recall Increase Procedure

Searching for DIs related to DI_A through further inter-DI relationships, other than r_A, broadens its semantic target and may very well increase a search query's recall. Such additional relationships will be connected to the latter through *owl:sameAs* or *skos:broader*.

In light of the above, the SPARQL query sent by cn_A to the cloud (in arrows 1 and 3 of Fig. 6), may have its recall increased by broadening its target relationship from "*commenting*" to "*responding*", in the manner depicted in Fig. 7

By enabling a higher searching precision and recall, this optimization mechanism evidently contributes to better satisfying user expectations regarding content discovery and, hence, increases user QoE.

Content Retrieval.
After a search has been performed, the user may choose to "consume" the WE-DID, identified as DID_{FT} (DID identified as FT). For this, cn_A will proceed to retrieve DID_{FT} from the cloud (as depicted in Fig. 6, arrow 5). If cn_A knows the identity of the server node responsible for delivery of the desired DID (sn_F), it sends a request directly to it (the case depicted in Fig. 6). If not, it sends it to whatever server node it knows, and the latter will redirect it to the appropriate one. sn_D is the server node to contact, since $S_{DID}^{sn_F} = \left\{ DID_{FA} - DID_{FZ} \right\}$ and thus includes all DIDs from identifiers FA to FZ.

sn_F will then deliver DID_{FT} to cn_A. The latter node will inspect the contents of DID_{FT} and discover the identity of its components *WEC-DIDs* (as DID_{FT} contains declarations of its relationships of "*containment*" towards said *WEC-DIDs*). cn_A will then proceed to retrieve all the *WEC-DIDs* (the actual media objects), that make up DID_{FT} (through the same procedure as the one just presented). It will then build the blog entry graphical presentation and expose it to the user.

Imagining that the user liked the selected blog entry, one can further suppose that she would like to retrieve the entire content of the blog. For that to occur, cn_A obtains the id of the entry's parent blog (*WH-DID*), from DID_{FT} (as it declares a "*composition*" relationship to it), which is, for instance, DID_{AF}. It then obtains from the cloud a list of all the *WE-DIDs* (entry DIDs), belonging to the blog. To do so, cn_A requests

from the server node (sn_H) responsible for the segment of S_{rel} that contains $composition_{WH-DID}^{WE-DID}$ (information about all the "composition" relationships from *WE-DIDs* towards *WH-DIDs*), a list containing all of the *WE-DID* maintaining a "*composition*" relationship towards DID_{AF} (arrows 7 and 8 of Fig. 6). cn_A may then proceed, as explained earlier in this section, to retrieve all the *WEC-DIDs* that each of the *WE-DIDs* contain.

The described content retrieval procedure enables a flexible distribution of the content delivery workload throughout the server nodes, and makes it easier for client nodes to locate and retrieve individual digital items. This, then, contributes to a speedier acquisition of the desired items, by users, and thus, to an increased user satisfaction and QoE.

Content Browsing.
Exploiting the inter-DI relationships the Weblog's user interface could be designed to include links to relevant DIs, offering users a complete view of related resources and allowing them to explore content efficiently. Ontology-enhanced user interfaces [17], promote adaptability and personalization of the user interface according to user needs, and provide a highly interactive environment to browse through content, all of which fosters an increased user QoE. For instance, a Weblog user might select to view resources with a narrower or broader relationship to a blog post – e.g. DIs with "*responding*" instead of the narrower "*commenting*" relationship.

Ranking.
Inter-DI relationships can also be exploited to provide ranking among Weblog's entries in a similar way to Google's PageRank [18]. DIs with most incoming links will be promoted and users will benefit from the ranking functionality. Moreover, the explicit declaration of the inter-DI relationship type eliminates the need to deduce the link type and can be exploited to enable sophisticated ranking methods, e.g. methods setting weights to the type of the relationship [19] or leveraging the relationship graph between the resources [20].
Such more sophisticated ranking methods will provide better search results making it easier, and faster, for users to find relevant information resources, thus increasing their QoE.

Final Remarks.
Blog supporting systems are not new and already enable users to enter, search for and access content. However, such systems are typically maintained by highly centralized facilities, which enforce custom and rigid site and data structures and relationships (i.e. links). As such, the relationships between the delivered resources are "hardcoded" into these systems' operational procedures, and the nature of such relationships is not precisely defined in terms of its meaning.

This lack of semantically precise, explicitly expressed, inter-resource relationships means that there is little opportunity to enforce a more automated or "intelligent" mode of operation or evolution. Expanding such systems is thus a complex task as it implies "hardcoding" new relationships and relationship interpretations or processing procedures.

The above sections demonstrate that the employment of relationally extended MPEG-21 innovatively enables digital resources and their metadata to be packaged into self-contained objects. It also enables the semantically precise definition of the relationships that bind such objects (in the mentioned metadata), webbing them into a coherent relational context.

These factors enable the definition of a distributed, server-side, architecture for supporting the WebLog system, which performs a dynamic sharing of the content distribution workload, as the responsibility over segments of the system's information spaces may be dynamically adapted.

Furthermore, they permit the enforcing of a more versatile and intelligent interpretation of the resource relationships and exploitation of the corresponding resources, enabling, for instance, smarter searches presenting higher precision and recall. All of this translates into an empowerment of users, in their ability to search, retrieve and access interesting content, and thus, into an increase in naturalness of interaction and quality of experience for the end user.

6   Contribution

The developed mechanism enables and stimulates the explicit and unambiguous declaration of the interior and exterior relational context of digital resources, clearly binding that information to the resources. Interior relational context is the set of relationships that exist between the inner components of a Digital Item, whereas exterior relational context is the set of relationships that exist between the DI itself and other DIs or other entities, (e.g. users). It does this in a manner that is aligned with emergent Semantic Web initiatives, notably the Linked-Data proposal, whilst being incorporated in a well-known open specification, the MPEG-21 standard, which natively provides extensive support for the digital representation and transaction of complex media resources.

It enables a semantically richer linking between digital resources. Unlike what happens, for instance, with HTML (*<a>* and *<link>* tags) and the typical existing tagging solutions, the connections it establishes between resources are semantically precise, and their type is clearly defined and inscribed within a contextualizing ontology. Furthermore, as such connections are no longer static links between object locations (the typical http links), but are, instead, semantic bonds between universally identified digital objects, there is a decoupling of content and location. This enables a more dynamic resource placement, at the server side, and greater resource availability for the client side.

The total ontology pool (and, as such, the total pool of relationships) is unlimited – such ontologies are always extendible, bridgeable or may always be defined anew for any specific application case, if necessary. The mechanism in scope thus enables the development of an unlimited mesh of globally coherent ontologies, and the establishment of a global relational-semantic fabric involving all web resources in a self-contained way (within the resources themselves). This simplifies the emergence and operation of more complex systems and makes it easier to evolve and expand them.

Given the semantically rich and coherent nature of this fabric, our mechanism, thus, enables the execution of novel and smarter searches, which profit from the interrelations between DIs and from the logical relationships between the interrelations themselves (as they are defined within an ontological context). It enables, for instance, the exchange of query precision for recall and vice-versa, to be performed by automated provisions.

The OAI-ORE standard provides some similar capabilities, as it also employs RDF for the description of inter-resource relationships. However, this standard mainly focusses on the definition of digital resource aggregations. That is, it defines sets of digital resources as part of a specific body (the aggregation), specifying the relationships between resources, as internal relationships of that aggregation. Furthermore, the OAI-ORE standard does not seek to enable the self-containment of the concerned information objects, as it presupposes a separate distribution of the related resources and relating metadata.

The mechanism presented and exploited in this paper, is more versatile, as it operates from a more flexible and universal perspective, whilst providing a more self-contained definition of the relational context of resources. In opposition to OAI-ORE, it does not require the definition of any aggregations nor is there any base granularity for an object to be able to describe inter-resource relationships. Every DI, defined in MPEG-21, may employ our mechanism to describe its relationships with the entire "digital world". Given that the mechanism is part of the standard itself, those descriptions may, thus, be contained within the very DIs that they pertain to. Furthermore, whereas OAI-ORE is limited to the description of the interior relational context of an aggregation of resources, our extended MPEG-21 based tool is able to describe the entire, internal and external, relational context of any resource or DI, without any need to conceive or define it, as an aggregation of resources. So, if, for instance, one intends to declare that a DI is a "*comment*" to another DI, there is no need to define any artificial aggregation between the two resources. One simply defines the relationship and stores it in whatever DI it is relevant to store it.

For all of the above, it can be seen that the defined mechanism empowers MPEG-21 to become an optimal tool for the weaving of the relational tissue that permeates all digital resources. It enables the development of a complex, yet precise and easily interpretable relational fabric, which optimizes the handling, discovery and delivery of

digital resources, in a flexible and self-contained manner, by complex and distributed systems, as those presented in section 5. It, thus, makes it possible to provide users with a speedier, more precise and automated access to content of their immediate interest, which corresponds to an increase in their QoE, within the context of their interaction with digital information resources on-line

7   Conclusions

The growing interpenetration of today's real and virtual worlds demands more coherent and universal ways of declaring resources as digital objects as well as more precise and semantically meaningful tools to specify the relationships between such resources.

Declaring resources as digital objects is exemplary accomplished by MPEG-21, but a framework, and tools, for expressing and managing semantically rich relationships are presently missing.

In this paper we have exposed MPEG-21's need for the addition of such relational capabilities. We then defined and described an extension, to MPEG-21, that enables it to declare, in an open, simple and efficient manner, semantically precise, inter-resources relationships.

Furthermore, we went on to conceive two novel exploitation venues for the defined mechanism, with a special focus on an exemplifying a cloud-based content delivery system. We thus showed how different systems and services might use the extended MPEG-21 to optimize their operation, rendering it more semantically aware and "intelligent" and enabling them to provide users with better experiences.

Our novel mechanism was scrutinized, approved and incorporated into ISO/IEC MPEG-21 standard, to provide it with the new, extended capabilities we have described.

The approached exploitation venues, however, represent only a fraction of all that is possible. There is ample room for development regarding the exploitation of MPEG-21's relational capabilities. New forms of automatically exploiting the developed tools, in different application areas can be conceived and experimented. Also, technologies and languages currently used for the relationships' declaration can be advanced or enriched.

8   Acknowledgment

References

1. Aitenbichler E, Behring A, Bradley D, Strufe T (2010) Shaping the Future Internet. Ubiquitous Computing and Communication Journal, Vol. 5(2)
2. Zahariadis T, Daras P, Bouwen J, Niebert N, Griffin D, Alvarez F, Camarillo G (2010) Towards a Content-Centric Internet. In: G. Tselentis et al. (Eds.) Towards the Future Internet - A European Research Perspective, IOS Press
3. W3C's HTML Working Group (1997) W3C HTML 4.01 Specification, Section 12–Links, http://www.w3.org/TR/html401/struct/links.html
4. CONVERGENCE Consortium (2011) CONVERGENCE project home site, http://www.ict-convergence.eu
5. Burnett I, Davis S, Drury G (2005) MPEG-21 Digital Item Declaration and Identification – Principles and Compression. IEEE Transactions on Multimedia, Vol. 7(3)
6. MPEG-21 (2005) ISO/IEC FDIS 21000-2:2005(E) MPEG-21 - Part 2: Digital Item Declaration. ISO
7. MPEG-21 (2006) ISO/IEC 21000-3:2003/FDAM 1:2006(E) MPEG-21 - Part 3: Digital Item Identification, AMENDMENT 1: Relates identifier types. ISO
8. MPEG-21 (2002) ISO/IEC FDIS 21000-3 MPEG-21 - Part 3: Digital Item Identification. ISO
9. MPEG-21 (2011) ISO/IEC 21000-3:2003/PDAM 2, MPEG-21 - Part 3: Digital Item Identification, AMENDMENT 2: Digital item semantic relationships. ISO
10. Gupta M, Li R, Yin Z, and Han J (2010) Survey on social tagging techniques. SIGKDD Explor. Newsl. 12, 1, 58-72
11. Chang C-H, Kayed M, Girgis MR, and Shaalan K (2006) A survey of web information extraction systems. IEEE Transactions on Knowledge and Data Engineering, vol. 18, no. 10, pp.1411 - 1428
12. Shadbolt N, Berners-Lee T, and Hall W (2006) The Semantic Web Revisited. IEEE Intelligent Systems, Vol. 21(3), pp. 96-101
13. Linked-Data Organization HomePage, http://linkeddata.org
14. Khare R (2006) Microformats: the next (small) thing on the semantic Web?. Internet Computing, IEEE , vol.10, no.1, pp. 68- 75, Jan.-Feb
15. RDFa Use Cases: Scenarios for embedding RDF in HTML. http://www.w3.org/TR/xhtml-rdfa-scenarios
16. Official blog for schema.org. http://blog.schema.org/2011/10/w3c-web-schemas-group-is-our-new-public.html
17. Paulheim, Heiko, and Probst F (2010) Ontology-Enhanced User Interfaces: A Survey. International Journal on Semantic Web and Information Systems (IJSWIS)
18. Page L, Brin S, Motwani R, Winograd T (1998) The PageRank Citation Ranking: Bringing Order to the Web. In Stanford InfoLab
19. Aleman-Meza B, Arpinar IB, Nural MV, and Sheth AP (2010) Ranking Documents Semantically Using Ontological Relationships. In Proceedings of the 2010 IEEE Fourth International Conference on Semantic Computing (ICSC '10)
20. Kahng M, Lee S, and Lee S (2011) Ranking objects by following paths in entity-relationship graphs. In Proceedings of the 4th workshop on Workshop for Ph.D. students in information & knowledge management (PIKM '11)
21. Open Archives Initiative Object Reuse and Exchange: http://www.openarchives.org/ore
22. Clocksin W. F. (2003) Artificial Intelligence and the Future. Philosophical Transactions Royal Society London, A 2003 361

23. Marsh J, Orchard D and Veillard D (2006) XML Inclusions (XInclude) Version 1.0, http://www.w3.org/TR/xinclude/, W3C