The Role of Foundational Ontologies for Domain Ontology Engineering

Article *in* International Journal of Information System Modeling and Design · August 2012

DOI: 10.4018/jismd.2010040101

4 authors, including:

Giancarlo Guizzardi
Universidade Federal do Espírito Santo

Fernanda Araujo Baião
Universidade Federal do Estado do Rio de Janeiro (UNIRIO)

Ricardo de Almeida Falbo
Universidade Federal do Espírito Santo

The Role of Foundational Ontologies for Conceptual Modeling and Domain Ontology Representation

Giancarlo Guizzardi
Computer Science Department, Federal University of
Espírito Santo, Vitória, Brazil
Laboratory for Applied Ontology (LOA),
ISTC-CNR, Trento, Italy

Abstract— **In recent years, there has been a growing interest in the development and use of domain ontologies, strongly motivated by the Semantic Web initiative. However, as we demonstrate in this paper, an approach for ontology representation uniquely based on the modeling languages adopted in the Semantic Web is insufficient to address a number of semantic interoperability problems that arise in concrete application scenarios. The main objective of this paper is to advocate in favor of an approach for conceptual modeling, in general, and domain ontology representation, in particular, in which lightweight modeling languages such as OWL and standard UML are complemented by modeling languages and methodologies based on theoretically principled Foundational Ontologies.**

I. INTRODUCTION

Since the word ontology was mentioned in a computer related discipline for the first time [1], ontologies have been applied in a multitude of areas in computer science. The first noticeable growth of interest in the subject in mid 1990's was motivated by the need to create principled representations of domain knowledge in the knowledge sharing and reuse community in AI, which motivated the creation of forums such as the conference series FOIS (Formal Ontology and Information Systems)[1]. Nonetheless, an explosion of works related to the subject only happened in the past five years, highly motivated by the growing interest on the Semantic Web, and by the key role played by ontologies in that initiative. Just to illustrate this point, the paper submission rate from the first International Semantic Web Working Symposium (SWWS) in 2001 [2] to the 4th edition of the International Semantic Web Conference (ICSW) [3] has increased by around 300%.

In the scope of the Semantic Web, ontologies are represented using a family of description logics-based languages which includes the languages RDF (Resource-Description Framework), DAML (DARPA Agent Modeling Language) and, more recently, the W3C recommendation OWL (Ontology Web Language). These languages, as well as most other languages which are used for conceptual modeling, in general, and ontology representation, in particular (e.g., UML, ER, LINGO), are based on a very simple meta-conceptualization, namely, the one of set-theory. For this reason, they are named here *lightweight ontology* languages, and the models produced using them are named, accordingly, *lightweight ontologies*.

In this paper, we demonstrate the insufficiency of these *lightweight ontology* languages to tackle a number of semantic interoperability problems that can arise in an open and dynamic scenario (such as, for instance, the Semantic Web). We then advocate that these languages should be complemented by a language and methodology based on a *Foundational Ontology*, i.e., a domain-independent common-sense theory constructed by aggregating suitable contributions from areas such as descriptive metaphysics, philosophical logics, cognitive science and linguistics.

In the remaining of this article we will make use of a running example to demonstrate the defended argument. In section 2, we present a number of lightweight ontologies used in a typical Semantic Web application, and review some semantic interoperability problems in the modeling and integration of these ontologies reported in the literature. In this section, we also show that lightweight ontologies equivalent to the ones discussed exist and are in active use in the Semantic Web, thus, demonstrating the expedience of the problem discussed here. Section 3 exemplifies how a modeling language and methodology rooted in a theoretically principled *Foundational Ontology* can be used to tackle the problems discussed in section 2. Section 4 elaborates on some final considerations of this article.

II. LIGHTWEIGH ONTOLOGIES IN A SEMANTIC WEB APPLICATION SCENARIO

Ríos [4] proposes an architecture for an ontology-based context-aware service platform[2]. This platform, depicted in Figure 1 below, employs distributed and concurrently

---

[1] http://www.fois.org/

[2] For a complete description of the proposed architecture and the WASP platform one should refer to [4].

developed ontologies to define the semantics of syntactic items which are used to compose the messages exchanged by the platform and its environment. These messages include both context-aware applications service subscriptions and context-information supplied by external (context) providers.

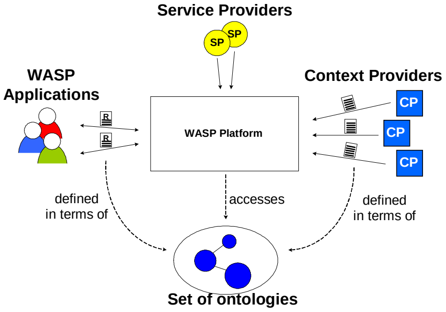

Fig. 1. An Ontology-Based version of the WASP platform

As demonstrated by Ríos, the use of ontologies in this version of the WASP platform brings a number of important benefits to the original proposal [5]. These benefits include:

1. More intelligent behaviour and the ability to reason about context information;
2. Reusability: the platform can (re)use already existing ontologies for the modeling of context information;
3. Flexibility: in contrast to the original proposal, the platform is not closed w.r.t. a pre-defined set of context modeling concepts.

Due to these benefits, ontologies are being considered in practically all the architectural evolutions of this platform [6].

In spite of these benefits, Ríos discusses the insufficiency of Semantic Web languages (and the produced lightweight ontologies) to prevent interoperability problems when different ontologies are integrated in such a scenario. This author proposes an illustrative example on the integration of five independent domain ontologies. The first ontology (whose fragment is depicted in Figure 2) is a Spatial ontology that defines the concepts of Spatial Location and Physical Object and their corresponding properties (e.g., Spatial Location includes attributes such as latitude and longitude coordinates). This ontology might be considered as a very simple generic ontology, as it does not define knowledge related to any specific domain. Thus, it can be referred or imported by different domain ontologies. For example, this ontology could be used by a GPS sensor agent to provide a service to track the location of physical objects in a context-aware platform.

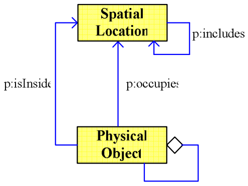

Fig. 2. Fragment of a Spatial Ontology (from [4])

There are two axioms defined for the *Spatial* ontology:

1. For every two arbitrary physical objects X and Y, if there are two spatial locations A, B, such that X occupies A, Y occupies B, and A is equal to B, then X and Y are the same physical object.

This axiom helps the users of this ontology to identify an object in a given time instant (*synchronic identity*). However, it cannot distinguish if two physical objects X and Y at different spatial locations in different time instants are the same objects (*diachronic identity*). For this reason, the ontology prescribes the following axiom.

2. For every two arbitrary physical objects X and Y, X is equal to Y if and only if they have the same parts, i.e., the *identity criteria* for physical objects is determined by the sum of its parts (*extensional identity criteria*).

A second ontology presented in the example regards the *Medical* domain. (see fragment on Figure 3) This ontology defines some medically related concepts such as *Human Organ* or *Human Being* and *Surgery Room*. Ríos presents a situation, in which the *Medical* ontology imports the concepts of *Spatial Location* and *Physical Object* from the *Spatial* ontology (symbolized by the i: character in the name of the class representing these concepts). The idea is to allow for the possibility of defining applications for checking location of patients, locate organs for transplants, and so forth.

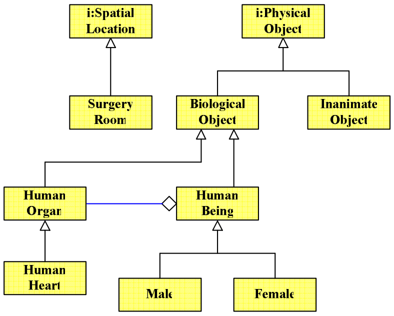

Fig. 3. Fragment of a Medical Ontology (from [4])

A third presented ontology, i.e. a fragment of a *Legal* ontology, is shown in Figure 4. This fragment represents legal aspects of people that can be used by bureaucratic applications. This ontology imports the concepts of *Human Being*, *Male* and *Female* from the *Medical* ontology. This import allows, for example, legal applications to refer to the medical histories of people; to have access to their personal data (e.g., blood type, skin color, fingerprints, height, weight); to differentiate people by sex; and to maintain a record of living and deceased people in a community.

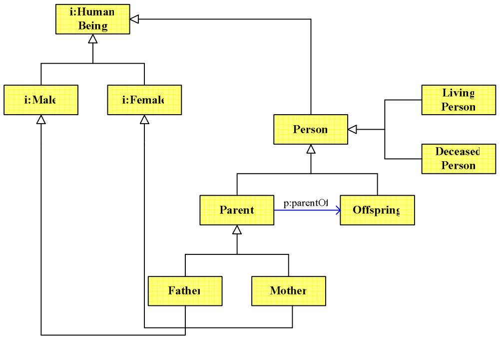

Fig. 4. Fragment of a Legal Ontology (from [4])

Figure 5 shows a fragment of a *Museum* ontology, which imports the *Spatial* and the *Medical* ontology to respectively define spatial locations like galleries within a museum, or inanimate objects like statues. These imported ontologies allow for applications to locate objects within the museum (e.g., statues, paintings) using the *Museum* ontology.

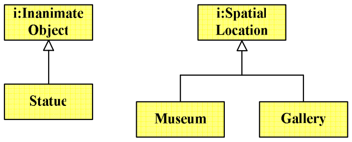

Fig. 5. Fragment of a Museum Ontology (from [4])

Finally, Figure 6 represents a fragment of a *Musical* Ontology containing some related concepts. Ríos defines as an application for the complete ontology (to which this fragment belongs) an *Event Advisor*, which notifies users about upcoming events that match their personal interests. The Music ontology imports from the Legal ontology concepts like person (and its possible attributes, like name, age, sex, etc.).

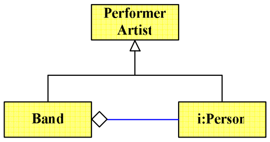

Fig. 6. Fragment of a Music Ontology (from [4])

A. *Concrete Examples in the Semantic Web*

In the original scenario proposed by Ríos, the used ontologies are created for the purpose of exemplification, with the aim of representing stereotypical problematic cases found in the WASP platform. The objective of this section is to demonstrate that real ontologies exist in current Semantic Web efforts which are structurally similar to the ones proposed by the author, and which are used by practitioners in concrete semantic web applications. Moreover, we also show that there are concrete efforts to unify these separate lightweight ontologies in context-aware applications in a manner analogous to the one described in the scenario proposed by Ríos.

A fragment of a *Music Ontology* such as the one presented in Figure 6 can be found in practice in the *MusicBrainz II Metadata proposal*[3]. A simplified version of the MusicBrainz database structure is presented in Figure 7. In this model, like in the one of Figure 6, the intention might be to represent that Artists can be either people or groups. However, due to the semantics of the subsumption relation, what actually is represented is that every person is an artist (and analogously, every group is an artist).

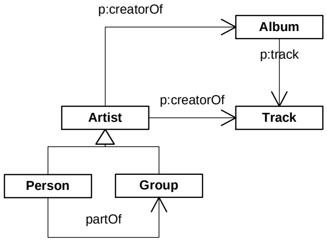

Fig. 7. Fragment of the MusicBrainz metadata proposal

In fact, the model excerpt depicted in Figure 7 can be seen as an extension of a more general pattern found in the *FOAF (Friend-of-a-Friend)* ontology[4] shown in Figure 8 below. The FOAF ontology is a proposal for capturing concepts related to the representation of personal information and social relationships. Its purpose is to serve as a basis for developing computational support for online communities. The FOAF ontology is also used by the SOUPA ontology (see discussion below) to support the expression and reasoning about a person's profile and social connections in pervasive computing applications.

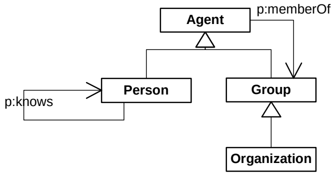

Fig. 8. Fragment of the FOAF (Friend of a Friend) Ontology

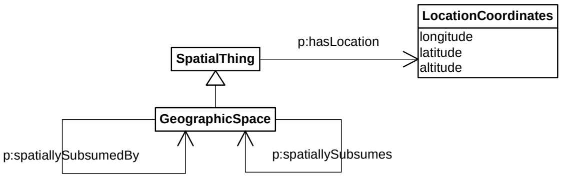

Fig. 9. Fragment of the SOUPA (Standard Ontology for the Ubiquitous and Pervasive Applications) dealing with spatial concepts and relations

The conceptualization modeled in the fragment of Figure 8 has an analogous representation in the SOUPA (Standard Ontology for the Ubiquitous and Pervasive Application)[5]

---

[3] MusicBrainz is a large database of music metadata (http://www.musicbrainz.org/).

[4] See The FOAF Project (http://www.foaf-project.org/) and the FOAF Vocabulary Specification (http://xmlns.com/foaf/0.1/)

[5] http://pervasive.semanticweb.org.

Ontology [8]. The SOUPA ontology also includes a Spatial Ontology (such as the one of Figure 2), whose fragment is presented in Figure 9.

SOUPA is a proposal for a standard ontology for supporting pervasive and ubiquitous computing application. It integrates parts of several other ontologies such as FOAF, DAML-Time [9], OpenCyC[6] and OpenGIS [10] (Spatial Entities), Rei Policy ontology [11], and COBRA-ONT [12], but also an ontology defining agent related concepts named MoGATU BDI ontology (see Figure 10)[7].

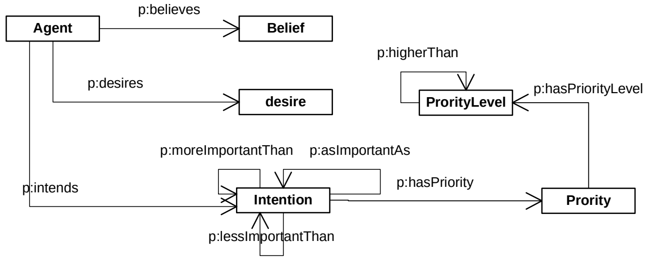

Fig. 10. Fragment of the MoGATU BDI ontology

B. Interoperability Problems

In [4], Ríos highlights the following problems that can occur with the integration of these ontologies:

1. “An application using the Medical ontology can derive the following wrong information: if a human being receives a heart transplant, he/she becomes a different human being. This is due to the extensional identity criteria, which is defined for physical objects in the Spatial ontology. If the identity of an object is defined by the sum of its parts, then changing one of the parts changes the identity of the object. Similarly, consider a tourist route planner application that plans a route including tourist points of interest or events never seen by the user of the application. Due to an accident, a human statue known by the user has lost a hand. [Having in mind that such application commits to the spatial ontology of Figure 2, it] will consider this statue different from the one the user visited; therefore it will be included in the route plan by error. This example uses a physical object (statue) for the purpose of illustration of the problem, but an analogous situation can be imagined with events such as a play or a concert”;
2. “Suppose an application for the obituary section of a music newspaper, which sends information about artists who die. It uses the Musical ontology, which imports the Legal ontology (to reuse the concept of person). The application will malfunction and it will send information about every person who dies, since [according to the ontology of Figure 6] every person is a performer artist. The intention in the ontology represented in Figure [6] is to represent that either persons or bands are performer artists. However, as a side effect, the ontology also states that every person is a performer artist”;
3. “Since Musical ontology imports the Legal ontology, which imports the Medical ontology, the heart (and all other parts) of a person can be inferred to be part of a band, due to transitivity of the “partOf” relation, which can cause undesirable inferences to be derived”.

III. ADDRESSING THE SEMANTIC INTEROPERABILITY PROBLEMS WITH AN ONTOLOGICALLY WELL-FOUNDED ONTOLOGY REPRESENTATION LANGUAGE

The main objective of this section is to illustrate how a foundational ontology can help in: (i) making explicit the underlying ontological commitments of the ontologies used in the examples of the previous section as well as the fragments of MusicBrainz II, FOAF, SOUPA and MoGATU BDI presented above; (ii) producing an adequate conceptual model representation that integrates these same ontologies. To reach our objective, we make use of the ontologically well-founded version of UML, as well as the modeling techniques proposed in [7]. Thus, from now on, we will apply concepts of these language and methodology. Such concepts will be defined along with the presentation of this running example, to facilitate their comprehension.

The result of redesigning the integrated model is shown in Appendix A. In producing this conceptual specification, we were forced to make a number of assumptions, giving particular interpretations of the represented concepts. This is due to the lack of information provided by some the integrated ontologies w.r.t. the real-world semantics assigned to these concepts. We emphasize, nonetheless, that the goal here is to demonstrate the suitability of the proposed modeling language. Thus, the underlying conceptualization which results from the set of assumptions made is of lesser relevance.

A. Principles of Identity

The problem (1) discussed by Ríos in the previous section originates from the assumption that there is one single principle of identity which all entities should obey. A *principle of identity* is a principle that supports the judgment whether two particulars are the same, i.e., in which circumstances the identity relation holds. Moreover, it defines which changes an entity can undergo and still be considered the same. In this case, particularly, it is assumed that all entities can obey an extensional principle of identity, i.e., two entities are the same iff they are composed of the same parts.

One famous puzzle in the philosophical literature that can be used to illustrate the notion of principle of identity is the puzzle whether, for instance, a certain statue is *the same* as the material it is made of. Take a statue of the Dalai Lama and a portion of metal that constitutes this statue in a given circumstance. Suppose that the statue is created out of this portion of metal at instant t_1. Additionally, suppose that at t_2 an accident causes the right hand of this statue to be destroyed. Alternatively, we can suppose that at t_3 an

---

[6] http://www.opencyc.org/.

[7] http://mogatu.umbc.edu/bdi/.

accident causes the statue to be melted preserving the exact portion of metal but completely altering its shape. Now, how can we answer the following questions: "Is the entity e_2 that we have in t_2 (t_3) the same as the entity e_1 the one we had in t_1?" A response to such question can only be given if we determine the principle of identity that e_1 (e_2) should obey, and consequently, which are its essential properties, i.e., the properties that this entity must have in all possible circumstances. These are, in turn, determined by the *kind* of thing an entity is. The notion of Kind used here is a technical notion which is fully described in a complete theory of identity and of conceptual modeling classifiers developed in [7, chap.4]. In summary, kinds are classifiers that are instantiated by its instances necessarily (in the modal sense) and are able to supply the unique principle of identity obeyed by these instances. Examples of kinds in the model of Appendix A include Person, Statue and Physical Objects (*quantities*) such as portion of metal, and lump of clay. Other types of classifiers admitted by this theory includes *roles*, *phases*, *mixins* and *categories*.

The kinds statue and portion of metal supply different principles of identity. For instance, while the latter supplies an extensional principle of identity, the former does not. As a consequence, we have that for an instance of portion of metal, all its parts are essential, since it cannot change any of its parts without altering its identity. In contrast, for an instance of statue some of its parts can be inessential (e.g., the hand). These two kinds are characterized not only by different types of essential parts but, more generally, by different types of essential properties. So, while for a statue its shape is essential, this is not the case for portion of metal. Thus, we must conclude that the statue and the material it is made of are numerically distinct entities.

Now, let us return to the model of Appendix A. We assume that, in contrast with physical objects, Biological Entities, Persons and Inanimate Entities do not carry extensional principles of identity. Therefore, we acknowledge the ontological distinction between the instances of these types and the physical objects that constitute them. We recognize that, for example, an Inanimate Object, such as a statue, and the raw material that constitutes it (e.g., a lump of clay) obey different principles of identity and, consequently, the relation between them is not one of identity but one of constitution. Likewise, we differentiate a Biological Entity, such as a heart, from the quantity of cellular tissue that constitutes this entity, and a Person from her body.

A Person is composed of a number of Biological Entities that amount to the person's body and its constituent parts. A Person has the spatial location of its body, which in turn is derived from the spatial location of its constituent Physical Objects. The same holds for Inanimate Entities. However, by separating in the model of Appendix A the classifiers that carry different principles of identity, we avoid the problems mentioned by Ríos in (1), i.e., it is not the case that the identity of a person is altered by replacing any of her body parts, and the statue of the Dalai Lama at t_2 is the same entity even if constituted by a different portion of metal.

The foundational ontology adopted in this work incorporates both the theory of identity aforementioned, as well as a theory of classifiers that acknowledges the ontological distinctions among different types of classifiers previously discussed. In the great majority of conceptual modeling and ontology representation languages (the semantic web languages included), all these important ontological distinctions (as well as the constraints derived from them) collapse into one single notion of types, with a semantics which is basically that of a unary predicate. Here in contrast, these constraints and distinctions are represented in the UML profile employed in the model of Appendix A.

B. Roles with Disjoint Allowed Types

The problem described by Ríos in (2) is a recurrent and much discussed problem in role modeling in the literature and is known as the problem of Roles with Disjoint Allowed Types. However, it can be rephrased as the problem of *roles* that can be played by instances of multiple *kinds*. For example, take the roles student, husband and exporting agency. The individuals that can play these roles are supplied by the kind Person, in the first two cases, and Organization, in the last one. For each of these roles, there is always a single kind supplying the instances that can play that role. This is not always the case. Take, for example, the role Customer, illustrated in Figure 11. Instances of Customer can be both of the kind Person or the kind Organization. The problem is how to represent the relation between the role Customer and the kinds Person and Organization? At first, two possible alternatives are the ones presented in parts (a) and (b) of Figure 11.

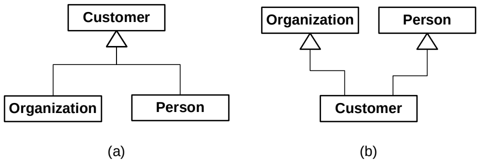

Fig. 11. Problems of modeling roles with disjoint allowed types

In the model of Figure 11(a), the role Customer is defined as a supertype of Person and Organization. This modeling is ontologically incorrect since: (i) not all persons (or organizations) are customers, i.e., it is not the case that the extension of Person is necessarily included in the extension of Customer; (ii) an instance of Person (or Organzation) is not necessarily a Customer, i.e., whilst Person (or Organization) are instantiated by its instances necessarily, Customer is only instantiated contingently. Together with the semantics of the subtyping relation, (ii) leads to a logical contradiction. In the model of Figure 11(b), the extension of Customer is empty, since, according to this model, every instance of customer is both a Person and an Organization.

By employing the theory of universals mentioned in section *A* above, we propose an *ontological design pattern* capturing a standard solution to this problem. The adequacy of this design pattern is demonstrated by several examples in [7]. In Figure c below we illustrate how this design pattern is used to solve this problem both for the case of the role Customer just discussed as well for the case of the Performer Artist type in Appendix A. The application of this design pattern solves the problem mentioned in (2): Performer Artist has as instances individuals that obey incompatible principles of identities, namely, Bands (which are kinds of Organizations) and Individual Artists (which are Persons). However, each of subtypes Individual Artist and Band that partition this type have extensions populated by individuals of one single kind.

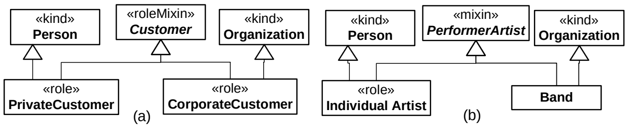

Fig. 12. Application of an Ontological Design Pattern

C. Transitivity of Parthood

Let us focus now on the problem earlier mentioned in (3), i.e, the issue raised by the transitivity of the *partOf* relation. The problem of transitivity of part-whole relations is a much debated topic not only in conceptual modeling but also in the linguistic and cognitive science literatures. In many conceptual modeling languages (e.g., UML), part-whole relations are always considered transitive. However, as discussed in [7, chap.5], examples of fallacious cases of transitivity among part-whole relations abound.

In [7], we developed a foundational theory of conceptual part-whole relations, which among other things address the problem of transitivity. We show that if we consider a unique general sense of parthood, transitivity cannot be said to hold unrestricted, but only with respect to certain *contexts*. The delimitation of contexts, however, typically requires extensive knowledge of the domain being modeled. In order to provide methodological assistance to the conceptual modeler in this task, we derive from this theory a number of language elements and methodological tools. For instance, we define a typology of different sorts of part-whole relationships (*subQuantityOf*, *subCollectionOf*, *memberOf*, *componentOf*) and demonstrate which combinations of these different types of relationships are transitive. Moreover, we were able to define a number of *visual patterns*, whose correctness is formally proven, that can be used to identify and delineate contexts of transitivity in class diagrams for the most complex and also most common sort of part-whole relation, namely, the *componentOf* relation. These visual patterns are depicted in Figure 13.

Figure 14 presents an excerpt of the specification of Appendix A that focuses on the meronymic relation between a Human Heart of a Band Member on one Band, and the relation between a Band Member and a Band, on the other. As this figure shows, this model is an exemplification of the pattern of Figure 13(d). As a consequence, the alleged derived relation between Human Heart and Band does not exist. This solves the problem (3) above for this case.

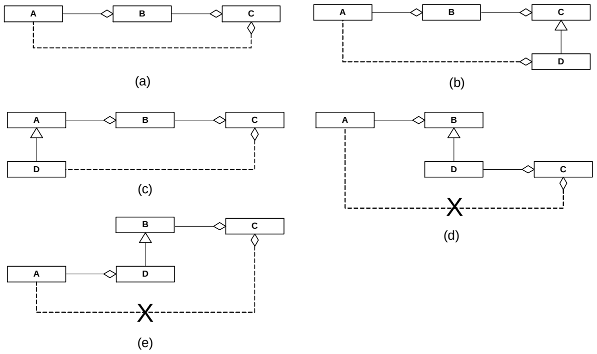

Fig. 13. Visual Patterns for isolating contexts of transitivity in functional part-whole relations: the patterns of figures (a), (b) and (c) represent cases in which a derived transitive parthood relation can be inferred. Intransitive cases are shown in figures (d) and (e).

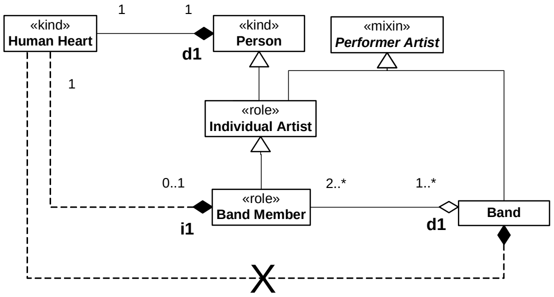

Fig. 14. Instantiation of the pattern that exemplifies situations in which transitivity does not hold across functional parthood relations.

The problem of transitivity is also manifested in Figure 15, which shows another fragment of the integrated model of Appendix A. A Person can be a member of a Band. For example, Eric Clapton is a member of the British Guitar Players. However, Clapton's hands are not members of this band. That is, also in this case, transitivity does not hold across the two represented meronymic relations, since the combination of *componentOf* and *memberOf* relations is never transitive [7, chap.5].

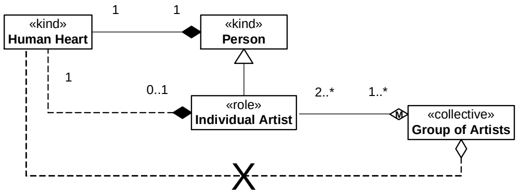

Fig. 15. Although a human heart can be part of an individual artist, which in turn is part of a group of artists, a human heart is never a part of a group of artists. This is because the combination of *componentOf* and *memberOf* parthood relations is never transitive.

D. Moments and Quality Structures

In our theory, we make a fundamental distinction between two different categories of concrete entities, namely, *substantials* and *moments*. Substantials are entities which can exist by themselves, i.e., independently of other entities. Examples of substantials include a person, a car, the moon, an electron. Moments, in contrast, are entities which are existentially dependent on other entities, in the way, for example, the electric charge of a conductor depends on the conductor, John's headache depends on John, Paul's knowledge of greek depends on Paul, the color of an apple depends on the apple. In other words, moments are entities which are parasitic to other entities named their bearers. The relation of existential dependence can be used to derive a further specialization within the category of moments: moments which are dependent on a single individual are named *qualities* or *modes*; moments which are dependent on multiple individuals are named *relators*. We will not elaborate on the distinction between qualities and modes here. For a comprehensive theory on the subject, the reader should refer to [7]. Relators are discussed in subsection *E*.

The special type of existential dependence relation connecting a moment to its bearer (named *inherence*) is a functional relation. Thus, for example, if we have two red objects x and y, the color c_x of x and the color c_y of y are two numerically distinction entities. However, c_x of c_y can be qualitatively indistinguishable when considered in a given *measurement structure*. We, therefore, distinguish between a moment inhering in an object, and the value of this moment when projected into a given measurement structure. These structures, named here *quality dimensions*, are endowed with geometrical and topological properties (e.g., ordering properties) that organized the possible values that can be assigned to individuals of a particular moment type. A quality dimension can be a one-dimensional structure (e.g. the possible weight values of an object), or a multi-dimensional structure ultimately composed of other primitive dimensions (e.g, the color structure, which is composed of the quality dimensions of hue, saturation and brightness. Measurement structures such as this one are named here *quality domains)*.

In the model of Appendix A, we separate physical spaces such as Surgery Room, Gallery and Museum from their spatial location (as informed by a GPS system). This is analogous to the approach adopted in the SOUPA ontology (Figure 9), as opposed to the solution used in the model of Figure 2. One reason behind this choice is the fact that some Geographical Spaces (e.g., Museum) can have several Location Coordinates, but also because different Geographical Spaces can be associated with a particular set of Location Coordinates in different circumstances. Thus, whereas the physical spaces are represented by the general category of spatial thing, the spatial location of these physical spaces is modeled here as a quality domain composed of the quality dimensions *latitude*, *longitude* and *altitude*.

Examples of intrinsic moments in Appendix A are the different types of Mentals States. Mental States are existentially dependent entities. For example, a Belief depends rigidly on a specific bearer active agent, i.e., a particular Belief cannot exist without inhering one (and always the same) active agent. By explicitly representing (objectifying) intrinsic moments, we can also represent their attributes and the relations in which they participate. For instance, every Intention has an inhering *priority* quality, which is modeled in Appendix A via an attribute function that maps Intentions to a priority value in a Priority Level quality dimension. This quality dimension is a finite set of values ranging from *level 0* to *level 10* and totally ordered under the *higherThan* and *lowerThan* relations.

As the discussion above shows, a modeling language based on a foundational ontology in which these categories are considered can allow for the representation of much more elaborated and semantically precise structures. Moreover, the explicit identification of existentially dependent entities has a direct consequence even for design and implementation, since existential dependence relations between objects in a conceptual level will typically give rise to life-cycle coupling between objects in design and implementation. Although these problems have not been identified by Ríos in his analysis, these are criteria which can be used to argue in favor of an ontologically well-founded conceptual modeling and ontology representation language.

E. Formal and Material Relations

In the foundational theory adopted in this work, relations are divided into two broad categories, called *material* and *formal* relations. *Formal relations* hold between two or more entities directly without any further intervening individual. A special type of formal relation considered here are the relations of *comparison* such as is taller than, is older than, knows more greek than. Comparison relations are logical constructions which are completely reducible to intrinsic moments, the values these moments take in a certain quality structure, and the relations between these values induced by the properties of these structures. For instance, the relation heavier-than between two atoms is a formal relation that holds directly as soon as the relata (atoms) are given. The truth-value of a predicate representing this relation depends solely on the atomic number of each atom: an atom **a** is heavier-than an atom **b** iff the atomic number of **a** (the projected value of **a**'s weight in the weight dimension) is bigger than the atomic number of **b**. As discussed in [7, chap.6], since comparison formal relations are founded in intrinsic moments of a certain type, the formal meta-properties of these relations can be derived from the properties of the quality structure associated with those moment types. In this example, the relation heavier-than is totally ordered because the weight quality dimension in which this relation is based is a total order.

Examples of formal relations of this type in Appendix A include the relations *moreImportantThan, lessImportantThan* and *asImportantAs*. These relations defined to hold between individual Intentions (*moments*) are completely reducible to the relations between the individual priority levels of these intentions. The first two relations are anti-symmetric and transitive. The last one is an equivalence relation. Here again these meta-properties are derived from the properties of PriorityLevel quality dimension.

Unlike formal relations, material relations are not founded on intrinsic moments of the involved relata. Material relations are induced by mediating entities called *relators*.

Thus, for a material relation to hold between two entities **a** and **b**, another entity needs to exist, namely, an instance of a relator which is existentially dependent on both **a** and **b**, hence, connecting the two. Take, for example, the relation *being married to* between John and Mary. This relation cannot be reduced to intrinsic properties of John and Mary. For this relation to hold, a certain wedding event involving John and Mary must have taken place which creates an individual relator *marriage* connecting the two. Similarly, we can say that Lisa works for the UN because there is an employment contract connecting them, and that Paul studies at the University of Twente because there is an Enrollment relator connecting the two. It is important to emphasize that a relator such the marriage m_1 between John and Mary is considered here a genuine ontological entity and can be thought as the aggregation of all social rights and responsibilities that John and Mary acquire by virtue of their participation in that relation.

There are several material relations represented in the model of Appendix A. As discussed in depth in [7, chap.6], the meaning of these relations is made evident by the explicit representation of their founding relators. Take for instance the relation *parentOf* between Parent and Offspring. In this case, we assume that parent is considered in this conceptualization in the legal not in the biological sense, and that in legal terms a person is a parent of another (offspring) iff the former is registered and legally recognized as such. Therefore, we explicit represent the Registration relators that connect parents to their offsprings. Another benefit of this approach is to allow for the unambiguous representation of the cardinality constraints of material relation, thus, avoiding what is known as the problem of collapsing single-tuple and multiple-tuple cardinality constraints [7]. This is the case for the *records* relation between the universals Performer Artist and Track. The multiplicity one-to-many from Track to Performer Artist leaves open several possible interpretations for the meaning of this relation. Does this multiplicity means that a track like *Georgia on my mind* can have several recordings (e.g., one by Ray Charles, and another by Jerry Lee Lewis)? By explicitly representing the Recording relator, the model makes clear that what is meant by a track is the result of specific recording. However, several artists can participate in one single Recording (e.g., both Clapton and B.B.King participate in the Recording of *Riding with the King*). Thus, the track *Georgia on my mind* recorded by Ray Charles, and the one recorded by Jerry Lee Lewis are different tracks.

Looking back at the ontologies previously exemplified in section II, the model of Figure 7 duplicates the relation creator between artist and album, and artist and track. The intention is to allow for the representation of tracks that are not parts of albums (i.e., that only exist as digital tracks). This situation is modeled in Appendix A by having AlbumTrack as a restriction of the type Track, in which the *restriction condition* is to be a part of an Album. However, it is unclear in the original model whether these two relations have the same real-world semantics. One interpretation is that, in case a track is part of an album, then the creator of the track must be same as the creator of that album. Still in this interpretation, in the case that different artists participate in the recording of different tracks of the same album (a song collection) they would all be considered creators of that album. The problem is that this interpretation does not allow for the situation in which an artist participates in the recording of one of more tracks of a given album, but is not considered an author of that album. Take, for example, U2's *Rattle & Hum*. Although B.B.King participates in the recording of *Angel of Harlem*, he is not considered an author of that album. We therefore assume that the relation between an artist and an album is one of *legal rights*. This is modelled in Appendix A by an Authorship relator universal.

IV. FINAL CONSIDERATIONS

The objective of this paper was to demonstrate the insufficiency of lightweight conceptual modeling languages (such as the representation languages typically in the realm of the Semantic Web), to address semantic interoperability problems that arise when one has to integrate concurrently developed conceptual models (or domain ontologies). In particular, we argue for the need of ontologically well-founded representation languages and modeling methodologies such as the one proposed in [7].

In other to make the case for our argument, we use some semantically interoperability problems highlighted by Ríos [4], which can happen in the integration of lightweight ontologies. These problems happen exactly because of the inadequacy of the used modeling language (OWL) in making explicit the underlying ontological commitments of the conceptualizations involved.

The conceptual modeling language employed here was proven useful in addressing these problems. First, by precisely representing the (modal) meta-properties of the underlying concepts, it allows for an explicit account of their ontological commitments. Second, by providing solutions to classical and recurrent problems in conceptual modeling (e.g., representation of roles with multiple allowed types, the problem of transitivity of parthood relations, the problem of collapsing single-tuple and multiple-tuple multiplicity constraints in the representation of associations, among others), it allows for the production of conceptually clean and semantically unambiguous integrated models.

This case study exemplifies the approach defended in [7, chap.3] for semantic interoperability of conceptual models, namely, that in a first phase of off-line meaning negotiation, an ontologically well-founded modeling language should be used. The main requirements for this language are *domain and comprehensibility appropriateness* [7, chap.2]. Once this meaning negotiation and semantic interoperation phase is complete, then a lightweight representation language can be used to express the results produced on this phase. The main requirements for such a language instead include computational efficiency in supporting automatic reasoning, machine-understandability, and easy mapping to standard design and implementation technologies.

ACKNOWLEGEMENT: The author would like to express his immense gratitude to Renata S.S.Guizzardi for her invaluable comments in previous versions of this article.

REFERENCES

[1].Mealy, G. H. ‘Another Look at Data’, Proceedings of the Fall Joint Computer Conference, November 14–16, Anaheim, California (AFIPS Conference Proceedings, Volume 31), Washington, DC: Thompson Books, London: Academic Press, 525–534, 1967.

[2].Cruz, I. ‘NSF Information and Data Management Program Report on the International Semantic Web Working Symposium (SWWS)’, Stanford University, California, USA, July 30 - August 1, 2001.

[3].Gil, Y.; Motta, E.; Benjamins, V.R.; Musen, M. (Eds.),4th International Semantic Web Conference, ISWC 2005, Galway, Ireland, 2005.

[4].Rios, D. ‘Using Ontologies in Context-Aware Service Platforms’, Master Thesis in Telematics, University of Twente, The Netherlands, 2003.

[5].Dockhorn Costa, P. ‘Towards a Service Platform for Context-Aware Applications’, Master Dissertation in Telematics, University of Twente, The Netherlands, 2003.

[6].Dockhorn Costa, P; Ferreira Pires, L.; van Sinderen, M.; Rios, D. ‘A services platform for context-aware applications’, 2nd European Symp. on Ambient Intelligence (EUSAI 2004), LNCS 3295, 8-10 November 2004, Eindhoven, The Netherlands, 2004.

[7].Guizzardi, G. ‘Ontological Foundations for Structural Conceptual Models’, ISBN 90-75176-81-3, Universal Press, The Netherlands.

[8].Chen, H. L.; Perich, F.; Finin, T.; Joshi, A. ‘SOUPA: Standard Ontology for Ubiquitous and Pervasive Applications’, In Proceedings of the First Annual International Conference on Mobile and Ubiquitous Systems: Networking and Services (Mobiquitous 2004), Boston, MA, August 22-26, 2004.

[9].Hobbs, J. R. ‘A DAML ontology of time’, [online: http://www.cs.rochester.edu/_ferguson/daml/daml-time-20020830.txt], 2002.

[10]. Cox, S.; Daisey, P.; Lake, R; Portele, C.; Whiteside, A. ‘Geography markup language (gml 3.0)’. In OpenGIS Documents. OpenGIS Consortium, 2003.

[11]. Kagal, L; Finin, T.; Joshi, A. ‘A Policy Based Approach to Security for the Semantic Web’. In 2nd International Semantic Web Conference (ISWC2003), September 2003.

[12]. Chen, H. L.; Finin, T.; Joshi, A. ‘An Ontology for Context-Aware Pervasive Computing Environments’, Special Issue on Ontologies for Distributed Systems, Knowledge Engineering Review, 18(3):197–207, 2004.

APPENDIX A