Towards Ontology Reuse

H. Sofia Pinto
Instituto Superior Técnico
Departamento de Eng. Informática
Grupo de Inteligência Artificial
Av. Rovisco Pais, 1049-001 Lisboa, Portugal
Tel: (351-1) 8417641 Fax: (351-1) 8417472
sofia@gia.ist.utl.pt

Abstract

Integration as the process of building an ontology reusing other ontologies which are parts of the resulting ontology has already been acknowledged as an essential process in building ontologies. Unfortunately little work has been done in this area. In this article we characterize the integration process, we provide some guidelines to that process, and we present a living classification of integration operations. The aim is to establish guidelines to an integration methodology and establish which integration operations should be provided to build ontologies by assembling modular ontologies.

Introduction

It is agreed that the strategy to build large ontologies should be the same as in software engineering: divide and conquer (Borst 1997). Therefore domain knowledge should be divided into "small, manageable pieces with strong internal coherence but relatively loose coupling" (Borst 1997). These pieces of domain knowledge should be specified in separate ontologies and kept in libraries as building blocks to construct larger ontologies. Mechanisms to combine ontologies (by assembling these building blocks into large ontologies) are required. Combining ontologies can be seen from two points of view:

- building an ontology reusing (by assembling, extending, specializing and adapting) other ontologies which are parts of the resulting ontology;
- building an ontology merging different ontologies on the same subject[1] into a single one that "unifies" all of them.

This paper is focused on the first point of view. This combining mechanism is ontology integration (Pinto, Gómez-Pérez, & Martins 1999).

Some of the available methodologies to build ontologies include an integration step (Uschold & King 1995; Gruninger 1996; Fernández, Gómez-Pérez, & Juristo 1997) but leave it more or less undefined how integration could be performed. In order to do ontology integration, operations that specify how knowledge from an imported ontology is going to be integrated in the resulting ontology are required.

[1]By subject we mean what the ontology deals about. We avoid the term domain since it is only used to describe domain ontologies and integration can be used to build general ontologies.

So far, few operations have been identified (Borst 1997; Farquhar, Fikes, & Rice 1996). Some ontology building descriptions allow us to abstract other operations (Dalianis & Persson 1997). Our own experience in building ontologies pointed us the need to abstract some other operations. We have built the Reference ontology (V. *et al.* 1998), a living domain ontology about ontologies that gathers, describes and has links to existing ontologies (a kind of yellow pages of ontologies), and we have incorporated this ontology into the $(KA)^2$ ontology (Benjamins & Fensel 1998). We were involved within a team in an effort to build an Environmental Pollutants ontology reusing other ontologies like Chemical-Elements (Fernández 1996) and Standard-Units (Gruber & Olsen 1994).

In this article we characterize the integration process and we present a living classification of integration operations. We begin by presenting previous work in, or related to, the integration problem. Then we present our own experience in building and integrating ontologies. We characterize the integration process and provide guidelines to it. Finally, we present our proposal for integration operations.

Previous Work

In the Ontolingua Server (Farquhar, Fikes, & Rice 1996), an ontology development environment for collaborative ontology construction, users are allowed three integration operations (Farquhar, Fikes, & Rice 1997): inclusion, polymorphic refinement, and restriction (specialization). *Inclusion* is used when the ontology is included (from the library of ontologies kept by the tool) and used as it is. The inclusion relations between ontologies may be circular, so one concept in one ontology can point to a concept in another ontology that, again, points to another concept in the first ontology. *Polymorphic refinement* extends one operation so that it can be used with several kinds of arguments. *Restriction* makes simplifying assumptions that restrict the included axioms. The Ontolingua Server also provides facilities for local symbol renaming. This facility enables ontology developers (1) to refer to symbols from other ontologies using names that are more appropriate to a given ontology and (2) to specify how naming conflicts among symbols from multiple ontologies are to be resolved.

PhySys (Borst 1997; Borst, Akkermans, & Top 1997), which is an ontology on Physical Systems, is based on general ontologies like Mereology, Topology, Systems Theory, Component and Process, that were implemented in Ontolingua (Gruber 1993). For instance, Topology was built reusing Mereology (by extending it). The researchers that built PhySys concluded that to build a large ontology from smaller ones, the dependencies between concepts and relations in different ontologies should be formalized as ontology projections. The three projections (that correspond to integration operations) identified are: include and extend, include and specialize and include and map[2]. *Include and extend* is the operation according to which "the imported ontology is extended with new concepts and relations". *Include and specialize* is the operation according to which "an abstract theory is imported and applied to the contents of the importing ontology" (concepts are specialized). Finally, *include and map* is the operation according to which "different viewpoints on a domain are joined by including the views in the domain ontology and formalization of their interdependencies".

To the best of our knowledge these are the only ontology integration operations proposals in the literature. However, there are some ontology building descriptions in the literature that show the need for other integration operations. A good example is the description of the construction of an ontology for the electrical distribution network domain (Dalianis & Persson 1997) reusing an electrical transmission network ontology (Bernaras, Laresgoiti, & Corera 1996) and a topology ontology (Benjamin & Jansweijer 1995). In what concerns the reuse of the electrical transmission network ontology, one of its subontologies was edited to create a similar ontology for the electrical distribution network domain. Most of the needed concepts were already represented. The modifications to the ontology can be summarized as:

- inclusion of the important concepts that were missing;
- removal of the concepts that were not relevant;
- replacement of two different concepts by a single one since they played essentially the same function (in the electrical distribution domain);
- adaptation of some definitions;
- changes in the name of some concepts since (1) the existing terminology was not recognized in the new domain or (2) the existing terminology was not the most agreed-upon in the new domain;
- redefinition of one relation because the existing relation was being misused.

The topology ontology was specialized. For that:

- all concepts had their names changed;
- important missing concepts were introduced;
- a relation that was of no use in the resulting ontology (a topology ontology for distribution networks) was removed.

Besides reusing these two ontologies a new ontology which models customers and customer agreements had to be built from scratch. The ontology resulting from assembling the three building blocks constitutes an adequate ontology to build electrical distribution network applications.

If we just want to improve or slightly modify the integrated ontology then it may be possible to mistakenly take integration for maintenance activities. In the case of Chemicals (Fernández 1996; Fernández, Gómez-Pérez, & Juristo 1997; Fernández *et al.* 1999) (an ontology that defines the chemical elements of the periodic table and their crystalline structures) the length centimeter ("cm") unit, a length unit commonly used in Europe (but not in the USA), was needed. The Standard-Units ontology (Gruber & Olsen 1994) available in the Ontolingua Server library did not include such unit when Chemicals was implemented in the Ontolingua Server. The solution found, with the operations available, was to develop a new ontology which included Standard-Units and add to it the needed unit. However the right solution, adopted latter, was the inclusion of this unit in the Standard-Units ontology kept in the library. This is the appropriate solution since this unit is not a specific purpose one, but a world wide generally accepted one and it applies to all domains that may reuse the Standard-Units ontology. Chemicals reuses other ontologies from the Ontolingua Server Library by simple inclusion: representation ontologies (van Heist, Schreiber, & Wielinga 1997), like the Frame ontology (Gruber 1993), KIF-sets, KIF-numbers and KIF-lists (Genesereth & Fikes 1992), and domain ontologies (van Heist, Schreiber, & Wielinga 1997), like Standard-Dimensions (Gruber & Olsen 1994).

Our Own Experience

For us, an ontology consists of a set of concepts (classes and instances), relations, functions, axioms, etc. that we refer to as knowledge pieces. Each knowledge piece in an ontology is associated to a name, documentation and a definition.

Building the Reference Ontology

We have built the Reference ontology (V. *et al.* 1998): a domain ontology about ontologies that plays the role of a yellow pages of ontologies. The Reference ontology was built at the knowledge level (Newell 1982) using METHONTOLOGY framework (Fernández, Gómez-Pérez, & Juristo 1997; Blázquez *et al.* 1998; Fernández *et al.* 1999) and the Ontology Design Environment (ODE) (Blázquez *et al.* 1998; Fernández *et al.* 1999). We have incorporated the Reference ontology into the reengineered version (Blázquez *et al.* 1998) of the $(KA)^2$ ontology (Benjamins & Fensel 1998).

The development of the Reference ontology was divided into three phases. In the first phase the conceptual structure was developed, its main concepts, taxonomies, relations, functions and axioms were identified. The knowledge source was a living taxonomy of features that characterizes ontologies from the user point of view (V. *et al.* 1998). Since one of the research topics in the KA community is ontologies it was decided to incorporate the Reference ontology into the $(KA)^2$ ontology. Before the incorporation step the $(KA)^2$ ontology was studied, analyzed, evaluated (verification and validation) and assessed (Gómez-Pérez, Juristo, &

[2]In earlier versions it was named include and project.

Pazos 1995) for reuse. In the second phase the conceptual model of the Reference ontology was incorporated in the restructured conceptual model of $(KA)^2$ (Blázquez *et al.* 1998). While knowledge represented in the $(KA)^2$ ontology describes the field of ontologies and differentiates ontologies from other fields of research, knowledge represented in the Reference ontology characterizes each ontology and differentiates one ontology from another. In this phase the knowledge sources used were the conceptual model of the Reference ontology, the restructured conceptual model of $(KA)^2$, and the set of properties identified for the Research-Topic ontology which were established during KEML workshop held at Karlsruhe, on January 23, 1998 and distributed by R. Benjamins to the KA-coordinators-list. In the third phase knowledge about specific ontologies that act as instances in the Reference ontology is added to the ontology. Ontology developers enter such knowledge by filling in a form distributing the effort of collecting information about specific ontologies.

For the purpose of this paper the second phase is the most relevant. The advantages of reusing the $(KA)^2$ ontology were: (1) some important knowledge needed by the Reference ontology was already available; (2) the knowledge represented in $(KA)^2$ reflects a consensus among a large number of researchers. The design criteria used to build and incorporate the Reference ontology into the $(KA)^2$ ontology were: modularize, specialize, diversify each hierarchy, minimize the semantic distance between sibling concepts, maximize relationships between taxonomies and standardize names of relations (see (V. *et al.* 1998) for detailed information). Another criterion was that modifications in the $(KA)^2$ structure should be kept to a minimum. To comply to this criterion it was decided that the knowledge about ontologies of the Reference ontology should be associated to the class `Ontology` of the Research-Product ontology and not as a separate subontology. Another criterion was that knowledge already represented in the ontology should not be removed and preferably not modified. The changes introduced in the restructured conceptual model of $(KA)^2$ can be categorized as:

- Some classes that are important to describe ontologies were missing in the $(KA)^2$ ontology and were therefore introduced. In the same category of changes is the addition of subclasses of existing or new classes. For example, the class `Languages`, subclass of `Computer-Support` at the Research-Product ontology, was introduced as well as its subclass, `Ontology-Languages`.
- Some instances that are important to describe knowledge in the Reference ontology, more precisely instances, were missing in $(KA)^2$ and were therefore introduced. For example, `Ontolingua` an instance of `Ontology-Languages` is important to describe the ontologies that were implemented in that language.
- Some of the relations and properties that are important to describe research products in general, identified as important by our Reference ontology, were missing in the $(KA)^2$ ontology. Therefore they were introduced for `Product` and not only for `Ontology`. For example, the relation `Distributed-by` between a `Product` and an `Organization` or the property `Product-Name`.
- Some of the relations and properties that are important to describe ontologies and research-topics were missing and were introduced. For example, the relation `Ontology-Formalized-in-Language` between an `Ontology` and an `Ontology-Language` or the property `Type-of-Ontology`.

As a summary, 4 classes were added, 22 new relations were defined and 71 new properties were introduced in the $(KA)^2$ ontology. From the study, analysis, evaluation and assessment that we did of the ontology we concluded that:

- Some small highly reusable ontologies should be developed and introduced in the $(KA)^2$ ontology namely about URL(s) and e-mail(s). This knowledge is important not only to describe ontologies but to describe persons, publications, etc. This was not done since it violated the minimal structural changes design criterion.
- In the Person subontology, one of its subclasses, `Administrative-Staff`, is inheriting inadequate binary relations, such as `Member-of-Program-Committee`, from class `Person`. These relations should be defined only for the adequate arguments and multiple inheritance mechanisms will make them inheritable only by the appropriate subclass, `PhD-Students`. The change could be done by removing the inappropriate relations from class `Person` and introducing them appropriately (defining them for the appropriate concepts). This was not done since it violated the minimal modifications criterion.

These changes were not done, but were suggested to the KA-coordinators.

Building an Environmental Pollutants Ontology

We were involved within a team in an effort to build an Environmental Pollutants ontology. Experts from various fields are involved in this project: Environmental experts (soil and water), Chemistry experts, Biology experts (public health), Geology experts, etc. This ontology should represent a unified, complete and consistent set of concepts needed to build environmental applications. This ontology reuses other ontologies such as Chemical-Elements (the subontology of Chemicals that defines the chemical elements of the periodic table) and Standard-Units. The first step to build this ontology was to look for ontologies to be reused. One of the elected ones was Chemicals. In the second step to build this ontology, Knowledge Acquisition meetings with Chemical experts were held and they were asked to evaluate Chemicals. The conclusions from the experts were:

- Only Chemical-Elements (C-E) should be reused because knowledge about the crystalline structure of the elements is not relevant.
- Some of the terminology used in C-E is not the most usual or adequate one. It should be changed.

- Some of the terminology used in C-E is not the standard terminology recommended by the International Union of Pure and Applied Chemistry (IUPAC). It should be changed.
- Some of the documentation provided for concepts represented in C-E should be revised since it came from outdated or less reputable sources.
- Some of the definitions represented in C-E should be revised (classes, instances, attributes, functions, relations, taxonomies, axioms).
- Some environmental properties of the elements are missing in C-E.
- Some of the knowledge should be revised since the knowledge sources used were not the most reputable ones. For instance, the values of the properties of the elements that were taken from (Barbor & Ibarz 1979) should be changed for the values in (Greenwood & Earnsshaw 1986).
- Some of the knowledge should be revised since all values for each property should be taken from one single source for all the elements. Therefore, if no value is provided in the knowledge source for one particular element it should not be looked for in another source.
- Some of the properties represented in C-E aren't of use for the environmental ontology but they should be kept since they characterize elements in their pure state.
- Some properties characterizing elements in their pure form are missing in C-E and should be included.
- Some properties of the elements that are not most adequately defined and that are not very important to characterize the elements should be removed.
- The values of the properties of the elements that are represented in less standard (or usual) scales should be substituted.

We were involved in the validation, verification and revision of C-E according to the expert's suggestions. The experts provided the sources from which the new knowledge should be taken. All unusual and non-standard terminology was verified and revised. All definitions were verified and revised. All documentation was revised. All properties had their values verified and corrected according to the knowledge sources and method suggested by the experts (some properties maintained their sources). Typo and introduction mistakes were corrected. Some values of some properties were missing and were introduced. Some values of some properties were removed since they came from more than one knowledge source. As a summary, from an universe of 103 elements, the values of 10 properties had their knowledge sources changed, the values of 7 properties were verified and the values for a new property were correctly introduced (for the time being, 4 properties were not verified). Besides being evaluated by chemical experts, the ontology was assessed by ontologists. We were also involved in the unification of the several versions of Chemicals (see (Gómez-Pérez & Rojas-Amaya 1999)) and helped in structuring the Monoatomic Ions ontology (Amaya 1998). In the Monoatomic Ions ontology knowledge (about the ions formed from the chemical elements in the periodic table) is represented from both a general chemical and an environmental point of view, paying special attention to soil, water, air and human health issues. Chemicals is included in the Monoatomic Ions ontology. The next step will be building the Poliatomic Ions ontology. Both the Monoatomic and the Poliatomic Ions ontologies will be some of the building blocks of the future Environmental Pollutants ontology.

Discussion

We would like to stress that all of our work was developed at the knowledge level. Since we had access to the knowledge level representations of both $(KA)^2$ and Chemicals the integration process was simplified. If the knowledge level representation of an ontology is not available, an ontological reengineering process (Blázquez *et al.* 1998) can be applied.

An important issue in reuse of ontologies is the language in which they are available. If the ontology is available in the required language and the ontology is going to be reused as it is the task is greatly simplified. The translation of ontologies is in itself a very important and difficult problem to be solved in order to allow more generalized reuse of ontologies. As discussed in (Uschold *et al.* 1998), translation is far from being a fully automatic process in the near future.

However the availability of the ontology in the required language is not enough. From our experience, we can say that integration is a process that takes place along the entire ontology building life cycle. We believe that the integration process and consequently the ontology building process can be greatly simplified if integration starts as early as possible in the ontology building life cycle. Therefore more integration effort is needed at the earlier stages of ontology building. In our case it began as early as the conceptualization phase.

In our case the task of finding and choosing the ontologies to be reused was greatly simplified since the reused ontologies were publicly available at the Ontolingua Server library, were published and we had previous knowledge of them. An important issue in the reuse of ontologies is the choice of the ontology to be reused. In a step towards simplifying this process we (V. *et al.* 1998) present a taxonomy of features and a WWW-broker to help users select the most adequate and suitable ontology.

One important conclusion that we would like to point out, from an integration point of view, is that sometimes knowledge must be relocated. For instance, in the case of the Reference ontology some knowledge (in this case properties and binary relations) was moved upwards in the hierarchy to class `Product` since it characterized products in general and not only ontologies specifically.

Another important point is that ontologies that are going to be reused should be evaluated (verified and validated) by domain experts (usually helped by ontologists) and assessed by ontologists in order to know their faults and strong points. Based in our experience, we identified the following (among others) criteria to which the experts should pay special attention to when analyzing the ontology for integration: (1) what knowledge is missing (concepts, classification criteria, relations, etc), (2) what knowledge should be removed, (3) what knowledge should be relocated, (4) which knowledge sources changes should be performed, (5) which documentation changes should be performed, (6) which terminology changes should be performed, (7) which definition changes should be made, in order to reuse an ontology that satisfies the requirements demanded of it.

Also based in our experience, we identified the following (among others) criteria to which the ontologists should pay special attention to when analyzing the ontology for integration: (1) the general structure of the ontology (one hierarchy, several hierarchies, a graph, etc.) to assess whether the ontology has an adequate (and preferably well-balanced) structure, adequate and enough modules, adequate and enough specialization of concepts, appropriate use of inheritance mechanisms, adequate and enough diversity, similar concepts are closer whereas less similar concepts are represented further apart, etc; (2) the basic distinctions (classification criteria made of the concepts described in the ontology) to assess whether they are relevant and exactly the ones (quantity and quality) required, (3) the privileged relation upon which the ontology is structured is the required one[3], (4) the names of the knowledge pieces follow standardization rules, (5) the definitions follow unified patterns, are simple, clear, concise, consistent, complete, correct (lexically and syntactically), precise and accurate, (6) the documentation is clear, helpful and adequate, (7) all and only the appropriate knowledge pieces are represented (or included).

If the ontology has not the adequate structure the changes to be made can be so extensive that it may be more cost effective to build an ontology from scratch. Changing the privileged relation according to which the ontology is organized can have profound consequences. The whole knowledge would, most probably, have to be revised, since the new relation organizes knowledge in a completely different way. Knowledge about a given domain that should be represented using one relation has nothing to do with what should be represented using another relation. Probably it is preferable to build a new ontology from scratch (if none of the available ones meets our needs).

Changing the basic distinctions (usually represented at the top-levels of the ontology) upon which the ontology is based can also imply a vast revision of the ontology. For instance, changing the top-level of C-E according to natural state at normal pressure and temperature conditions instead of the group classification proposed by Mendeleev will cause profound changes in the whole ontology. The only knowledge that can be reused are the instances representing the elements. Usually in this case it is preferable to build another ontology.

Integration Process

In integration there are at least two ontologies involved: one or more ontologies that are integrated, and one ontology resulting from the integration process. The integrated ontology(ies) are those that are being reused. They are a part of the resulting ontology. The ontology resulting from the integration process is what we want to build and although it is referenced as one ontology it can be composed of several "modules", that are (sub)ontologies. This happens not only in integration but when building an ontology from scratch.

The domain of the integrated ontology is different from the domain of the resulting ontology but there may be a relation between both domains. The integrated concepts can be, among other things, (1) used as they are, (2) adapted (or modified), (3) specialized (leading to a more specific ontology on the same domain) or (4) augmented by new concepts (either by more general concepts or by concepts at the same level). The domains of the different integrated ontologies usually are different among themselves (either from the resulting ontology or the various integrated ontologies). In integration, the domain in which we are building the new ontology should be such that there is no similar ontology already built, otherwise one should simply reuse the existing one.

The reused ontologies are chosen from those available in ontology libraries that meet a series of requirements, such as, domain, abstraction, type (van Heist, Schreiber, & Wielinga 1997), generality, modularity, evaluation, just to name a few. The resulting ontology should have all the properties of a good ontology. Not only should it be clear, coherent, extensible, comply to the principle of the minimal ontological commitment and to the minimal encoding bias as proposed in (Gruber 1995) but it should also be complete, concise, non-ambiguous (related to coherence), have an adequate level of detail, be built upon the appropriate basic distinctions, have been evaluated, etc.

The problem is how to integrate several existing ontologies within a new one that is being built. Problems such as consistency of the resulting ontology, level of detail throughout the whole ontology[4], etc, have to be dealt with. The solution seems to be the specification of a set of integration operations that specify how knowledge in the soon to be integrated ontology is going to be included and combined with the knowledge in the ontology that is being built. Integration operations can be viewed as composing, combining or assembling operations. However these operations should only be performed if the integrated ontologies have a series of features. Not only will the features assure that the integrated ontology is the most appropriate one but also that the integration operations can be successfully applied and that the resulting ontology will have the desired characteristics. In (V. *et al.* 1998) we present a series of features (and a WWW broker), not specifically for integration purposes, that can help the search for suitable ontologies.

As the development of an ontology should follow an evolving prototyping life cycle, the ontology may be considered for integration in specification, conceptualization, formalization, implementation and maintenance. That is, we

[3]An ontology can be thought of as structured or organized according to one privileged relation, for example, ISA, part-of, etc.

[4]That is, the ontology doesn't have "islands" of exaggerated level of detail and other parts with an adequate one. It should be stressed that none of the parts should have less level of detail than the one required or else the ontology would be useless, since it would not have sufficient knowledge represented.

can have different integration procedures for the same ontology but in different states of the ontology building process. As we have an evolving prototyping ontology building life cycle the same ontology can be used in the same state in integration activities more than once. These procedures and activities form the overall process of integration. The integration process needs to be further studied, namely, the integration procedures and activities need to be defined and a larger set of integration operations needs to be identified, specified and defined. In the next section we identify a larger set of operations.

Integration Operations

We would like to remark that all integration operations proposed in this article were identified at the knowledge level. Therefore symbolic or implementational issues were not taken into consideration. The objective of our work is to specify integration operations independently of a specific methodology.

Broadly speaking, reusing an ontology by means of integration can require changes in the terminology, changes in the documentation and changes in the definitions of all knowledge pieces in the ontology. Moreover it must be possible to introduce and/or remove knowledge pieces to/from the ontology. We will refer to removal, introduction and composition of removal and introduction operations as changes. In the overall, operations should be provided to change:

- which knowledge pieces are represented in the ontology,
- the documentation of any knowledge piece represented in the ontology,
- the terminology of any knowledge piece,
- the definition of any knowledge piece.

Changes in the documentation are usually performed to update it or increase its clarity. Changes in terminology are usually performed to comply to naming standardization rules, or introduce standard or more usual terminology. Changes in the definitions are usually performed in order to more accurately, precisely, simply, clearly, concisely, correctly, completely and better represent knowledge of a given domain. When the changes introduced in an ontology by integration operations are profound and extensive, they imply the creation of a new ontology on the same domain.

Integration operations can be classified into:

- operations on whole ontologies: inclusion of an ontology.
- operations on the constituents of an ontology: change name, documentation, and definition of any knowledge piece. All these operations should be used parsimoniously. If we are going to remove all concepts and introduce new ones, we are not really reusing the ontology but building another one. Operations that relate new pieces of knowledge to existing knowledge should be provided (that is, operations to manipulate the relations that structure knowledge). For instance, the inclusion of a concept usually needs the specification of the place in the hierarchy where the concept should be introduced (for instance, by stating the parents of the class). The removal of a class may imply or not the removal of the whole hierarchy underneath it. If the underneath hierarchy should not be removed then the dangling hierarchy should be kept in the ontology as another hierarchy. Removing a relation (a structuring one or a non structuring one) implies removing all tuples representing the relation from the ontology. Changes in the definition include, among others:
  - in the case of concepts (classes or instances):
    * changes in the relations in which the concept is involved including the privileged relations that structure the ontology, changes in the properties/attributes defining the concept, etc.
  - in the case of attributes/properties of concepts:
    * changes in the value of the attribute, changes in the cardinality of the attribute, changes in the class of values of the attribute, changes in the value of any facet of the attribute (default value, value range, etc), etc.
  - in the case of relations (and functions since they are special relations):
    * changes in the class of values to which its arguments must belong to, changes in the arity of a relation, changes in the tuples that hold (the extensive definition of the relation), etc.
  - axioms, constants, etc

The operations identified above are simple integration operations. They manipulate the knowledge pieces that form the ontology. They are the basis for more complex operations. For example, ontology specialization can be performed by composing the application of simple integration operations, as previously described. Other operations, such as mapping and extension, can also be defined by composing simple integration operations. Further research in integration operations is within our plans for the near future. Future work in integration operations will include identifying a broader set of integration operations and its specification.

Conclusions and Future Work

Ontology integration needs to be more thoroughly studied. In this article we present our experience in ontology integration, we characterize this process and give some guidelines to it. Integration is based on a set of operations. In this article we present a living set of integration operations. This is just an initial set of integration operations, therefore the set may still increase. The set of integration operations needs to be more thoroughly/formally specified (requirements, how do they work?, etc.). In the near future we hope to stabilize the initial set of integration operations, identify the requirements of those operations, better and formally specify those operations and better specify the integration process.

Acknowledgments

We thank Prof. João P. Martins for his supervision. We thank all members of GIA at Instituto Superior Técnico for their support, specially António Leitão, João Cachopo, and

Ana Cachopo. The work described in "Building the Reference Ontology" and "Building an Environmental Pollutants Ontology" has been done during our stay at the Knowledge Sharing and Reuse group at the Lab. de Inteligencia Artificial of the Universidad Politecnica de Madrid (UPM). We thank Prof. Asunción Gómez-Pérez for the opportunity to work within her group. We thank Mariano Fernández for his comments about Chemicals, METHONTOLOGY and ODE. We thank Adolfo Lozano and Asunción Gómez-Pérez for our joint collaborative identification of the living taxonomy of features that was the starting point of the Reference ontology. We thank the vast team involved in building the Environmental Pollutants ontology, in particular M.D. Rojas Amaya with whom we jointly worked in the Configuration Management, Validation and Verification of Chemicals, and in structuring the Monoatomic Ions ontology. This work was partially supported by JNICT grant No. PRAXIS XXI/BD/11202/97 (Sub-Programa Ciência e Tecnologia do Segundo Quadro Comunitário de Apoio) and Project PRAXIS XXI 2/2.1/TIT/1568/95.

References

Amaya, M. D. R. 1998. Ontologia de Iones Monoatómicos en Variables Fisicas del Medio Ambiente. Proyecto Fin de Carrera, Fac. de Informática, UPM.

Barbor, J. A., and Ibarz, J. 1979. *Quimica General Moderna.* Martin.

Benjamin, J., and Jansweijer, W. 1995. The KACTUS Library of Ontologies v.1.0. Technical report, ESPRIT Project 8145 KACTUS.

Benjamins, R., and Fensel, D. 1998. The Ontological Engineering Initiative $(KA)^2$. In Guarino, N., ed., *Formal Ontology in Information Systems*, 287–301. IOS Press.

Bernaras, A.; Laresgoiti, I.; and Corera, J. 1996. Building and Reusing Ontologies for Electrical Network Applications. In Wahlster, W., ed., *ECAI96 Proceedings*, 298–302. John Wiley&Sons.

Blázquez, M.; Fernández, M.; García-Pinar, J. M.; and Gómez-Pérez, A. 1998. Building Ontologies at the Knowledge Level Using the Ontology Design Envirnoment. In *Knowledge Acquisition Workshop.*

Borst, P.; Akkermans, H.; and Top, J. 1997. Engineering Ontologies. *International Journal of Human Computer Studies special issue "Using Explicit Ontologies in KBS Development"* 46(2/3):365–406.

Borst, P. 1997. *Construction of Engineering Ontologies for Knowledge Sharing and Reuse.* Ph.D. Dissertation, Tweente University.

Dalianis, H., and Persson, F. 1997. Reuse of an Ontology in an Electrical Network Domain. In *AAAI97 Spring Symposium Series, workshop on Ontological Engineering.*

Farquhar, A.; Fikes, R.; and Rice, J. 1996. The Ontolingua Server: A Tool for Collaborative Ontology Construction. In *Knowledge Acquisition Workshop.*

Farquhar, A.; Fikes, R.; and Rice, J. 1997. Tools for Assembling Modular Ontologies in Ontolingua. In *AAAI97 Proceedings*, 436–441.

Fernández, M.; Gómez-Pérez, A.; Sierra, A. P.; and Sierra, J. P. 1999. Building a Chemical Ontology Using METHONTOLOGY and the Ontology Design Environment. *IEEE Intelligent Systems* 14(1):37–46.

Fernández, M.; Gómez-Pérez, A.; and Juristo, N. 1997. METHONTOLOGY: From Ontological Art Towards Ontological Engineering. In *AAAI97 Spring Symposium Series, workshop on Ontological Engineering.*

Fernández, M. 1996. CHEMICALS: Ontologia de Elementos Químicos. Proyecto Fin de Carrera, Fac. de Informática, UPM.

Genesereth, M., and Fikes, R. 1992. Knowledge Interchange Format. Technical Report Logic-92-1, Computer Science Department, Stanford University.

Gómez-Pérez, A., and Rojas-Amaya, D. 1999. Ontological Reengineering for Reuse. In *11th European Workshop on Knowledge Acquisition, Modeling and Management, EKAW99.* Springer Verlag. Forthcoming.

Gómez-Pérez, A.; Juristo, N.; and Pazos, J. 1995. Evaluation and Assessement of the Knowledge Sharing Technology. In Mars, N., ed., *Towards Very Large Knowledge Bases*, 289–296. IOS Press.

Greenwood, N. N., and Earnsshaw, A. 1986. *Chemistry of the Elements.* Pergamon Press.

Gruber, T., and Olsen, G. R. 1994. An Ontology for Engineering Mathematics. In J. Doyle, E. S., and Torasso, P., eds., *KR94 Proceedings*, 258–269. Morgan Kaufmann.

Gruber, T. 1993. A Translation Approach to Portable Ontology Specifications. *Knowledge Acquisition* 5:199–220.

Gruber, T. 1995. Towards Principles for the Design of Ontologies for Knowledge Sharing. *International Journal of Human Computer Studies* 43(5/6):907–928.

Gruninger, M. 1996. Designing and Evaluating Generic Ontologies. In *ECAI96's workshop on Ontological Engineering.*

Newell, A. 1982. The Knowledge Level. *Artificial Intelligence* 18:87–127.

Pinto, H. S.; Gómez-Pérez, A.; and Martins, J. P. 1999. Some Issues on Ontology Integration. In *IJCAI99's workshop on Ontologies and Problem Solving Methods: Lessons Learned and Future Trends.* Forthcoming.

Uschold, M., and King, M. 1995. Towards a Methodology for Building Ontologies. In *IJCAI95's workshop on Basic Ontological Issues in Knowledge Sharing.*

Uschold, M.; Healy, M.; Williamson, K.; Clark, P.; and Woods, S. 1998. Ontology Reuse and Application. In Guarino, N., ed., *Formal Ontology in Information Systems*, 179–192. IOS Press.

V., J. A.; Gomez-Perez, A.; T., A. L.; and Pinto, H. S. 1998. $(ONTO)^2$Agent: An Ontology-Based WWW Broker to Select Ontologies. In *ECAI98's workshop on Application of Ontologies and Problem Solving Methods.*

van Heist, G.; Schreiber, A. T.; and Wielinga, B. J. 1997. Using Explicit Ontologies in KBS Development. *Internacional Journal of Human-Computer Studies* 45:183–292.