

A Synthesis of State of the Art Enterprise Ontologies

Work in Progress

Chris Partridge
The BORO Program
partridge@BOROProgram.org

Milena Stefanova
LADSEB CNR, Italy
stefanova@ladseb.pd.cnr.it

Abstract

This paper presents a report on work in progress of a Synthesis of (selected) State of the Art Enterprise Ontologies (SSAEO) – which aims to produce a Base Enterprise Ontology to be used as the foundation for the construction of a Core Enterprise Ontology (CEO). The synthesis is intended to harvest the insights from the selected ontologies, building upon their strengths and eliminating – as far as possible – their weaknesses. One of the main achievements of this work is the development of the notion of a person (entities that can acquire rights and obligations) enabling the integration of a number of lower level concepts. In addition, we have already been able to identify some of the common ‘mistakes’ in current enterprise ontologies – and propose solutions.

1. Introduction

This paper results from a collaboration between two projects: the BRont (Business Reference Ontologies)¹ and European IKF (Intelligent Knowledge Fusion)² projects. The BRont project is part of the BORO Program, which aims to build ‘industrial strength’ ontologies, that are intended to be suitable as a basis for facilitating, among other things, the semantic interoperability of enterprises’ operational systems.

This European IKF project has as an ultimate goal the development of a Distributed Infrastructure and Services System (IKF Framework) with appropriate toolkits and techniques for supporting knowledge management activities. The following countries participate in the IKF project; Italy, UK, Portugal, Spain, Hungary and Rumania. The project will last 3.5 years, and started in April 2000.

There are a couple of vertical applications whose domain is the financial sector. One of these, IKF/LEX – a part of the Italian IKF project – has been selected to undertake a pilot project. IKF/IF-LEX is lead by ELSAG BankLab

SpA and its goal is to provide semi-automatic support for the comparison of banking supervision regulations.

There will be two kinds of ontologies developed within the IKF project:

- A *Reference Ontology* composed of a *Top Level Ontology* and several *Core Ontologies* (see [1]). The top level ontology contains primitive general concepts to be extended by lower-level ontologies. The core ontologies span the gap between various application domains and the top level ontology. The IKF/IF-LEX and the BRont projects are collaborating on developing a *Core Enterprise Ontology* (CEO) that IKF will use on this and its other applications in the enterprise domain.
- *Domain Ontologies*. The vertical applications will build ontologies for their specific domains. For example, the IKF/IF-LEX project is building an ontology for bank supervision regulations, focusing on money laundering.

2. SSAEO work plan

The scope of the SSAEO work is large – and so the work has been divided into more manageable chunks.

As [1] states, a core ontology contains “the categories that define what a field is about.” A first rough intuitive guess of what these categories might be has proved a useful tool in:

- helping clarify the scope focus on the important aspects for the CEO, and
- acting as a basis for segmenting the work.

The selected categories are:

- *parties (persons)* which may enter in
- *transactions* (composed of agreements and their associated activities), involving
- *assets*.

The ontologies to be analysed were selected according to:

- the relevance of their content to the Core Enterprise categories, and
- the clarity of the characterisation of the intended interpretations of this content (see [7], [11]).

This gave us the following list:

- Toronto Virtual Enterprise - TOVE (see [6], [13]),
- AIAI’s Enterprise Ontology - EO (see [3], [14]),

¹ <http://www.BOROProgram.org>

² http://www3.eureka.be/Home/projectdb/PrjFormFrame.asp?pr_id=2235

- Cycorp's Cyc® Knowledge Base - CYC (see [2]),
- W.H. Inmon's Data Model Resource Book - DMRB (see [10]).

The work proceeds by analysing one category in one ontology at a time, and then re-interpreting the previous results in the light of any new insights. Initially, the work focuses on individual ontologies but as it proceeds there is enough information to start undertaking comparisons between ontologies. The final analysis will encompass analyses of both the individual ontologies and comparisons between them.

In each of the ontologies, the concepts and relations relating to the category being considered are examined for the clearness and uniformity of their descriptions and formalisations. Further, each concept is analysed for its coverage and extendibility in cases where the coverage is not complete. Relations between concepts that are not explicitly described, but clearly exist, are identified as well. In addition, for the sake of a clear interpretation, we have found it necessary to consider the top concepts (whether or not they are explicitly described).

An important part of the analysis is testing each concept and its relations against a number of standard examples and more specialized concepts. Further, a check is made against a number of standard difficult cases. Both these checks help to identify weaknesses in the coverage of the ontologies.

A key concern in the analysis is to understand how the various concepts interlink with one another, to better understand the unifying structure of the Enterprise ontology.

At various stages during the analysis an interim ontology is synthesised from the strengths found in the analysis, in such a way as to eliminate the known weaknesses – and itself analysed. In the final synthesis, all the categories in all the ontologies are combined into a base CEO ontology.

At this time, the SSAEO work is concluding the analysis of the Parties (Persons) category for the EO and TOVE ontologies – and early drafts of synthesised ontologies are being reviewed. There is still substantial work that needs to be done in determining the precise relations between concepts, such as *LEGAL ENTITY* and *OWNERSHIP* within the EO.

3. Initial Findings

Though both the ontologies have many important insights and provide much useful material – our most general findings, at this stage, are that none of the ontologies:

- adequately meet our criteria of clear characterisation, or
- really share a common view of what an organisation is.

Taken together, these findings mean that the creation of the synthesised base CEO ontology cannot just be a simple merging of the common elements of the selected ontologies.

We now illustrate these findings with examples. We also show how we synthesised a resolution to some of these problems - for the two ontologies we have analysed.

3.1. Clear Characterisation

With an unclear characterisation it can be difficult to work out the intended interpretation – in the worst case, impossible to decide between competing interpretations. There are many different ways in which the characterisation can be unclear – as we show below.

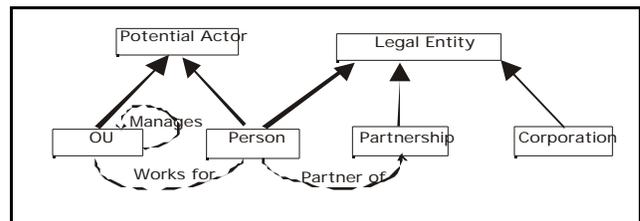


Figure 1. Simplified EO overview

In both TOVE and EO we found no clear overview of the structure – so we developed graphical representations to help us understand it. Figures 1 & 2 provide simplified versions of these.

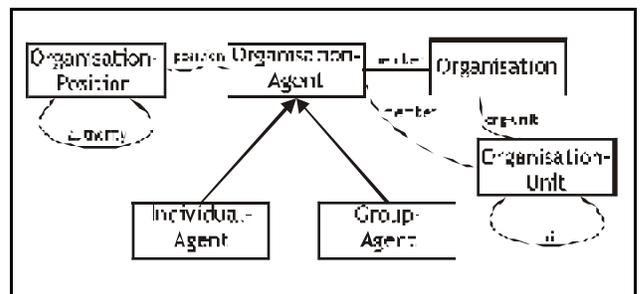


Figure 2. Simplified TOVE overview

Both TOVE and EO make use of a number of top concepts. A top ontology – or top concepts – can provide a useful structure for defining and using domain concepts and relations – segmenting the enterprise and other domains into general categories. However, if this is not done properly it can have the opposite effect.

Some of the problems we encountered with the top concepts and the domain analysis are:

- *Insufficient characterisation of the disjointness of top concepts.* For example, in the informal EO the relationship between the top concepts *ENTITY*, and *ROLE* is not clear – in particular, whether *ROLES* can be *ENTITIES* or not, and so whether they can enter into *RELATIONSHIPS*.

- The same lack of care in characterising disjointness (and overlapping) exists at the domain level in both TOVE and EO. We found this can make it impossible to definitely determine the intended interpretation. For example, in TOVE the formalisation allows an *ORGANISATION-UNIT* to be an *ORGANISATION* – though this seems counter-intuitive, and probably not what the authors intended.
- *Not applying top concepts.* TOVE states that a fluent is "a [type of] predicate or function whose value may change with time". But it does not identify which predicates in its ontology are fluents – leaving this to the readers, who have to make their own judgements. Supplying such information would have helped not only the users of the ontology but also its creators and designers. For example, the TOVE's creators end up (probably unintentionally) having to regard *ORGANISATION* as a fluent – when in the normal (commonsense) use of the concept it is not.
- *Messy formalization trajectories.* EO formalizes its concepts in logical systems (Ontolingua and KIF), which rely on their own (different) top concepts. An attempt for a clear formalisation trajectory has been made (see [14]), but unfortunately this does not match very well with the informal specification. For example, in the informal EO it is stated that each *RELATIONSHIP* is also an *ENTITY*, but is not defined as such in the formalization. Furthermore some *RELATIONSHIPS* are defined in the formalization as *classes* and others are defined as *relations* without explaining what the motivations for these choices are (e.g., *SALE* is a *RELATIONSHIP* formalized as a *class*, *HAVE-CAPABILITY* is a *RELATIONSHIP* formalized as a *relation*). This becomes a more serious problem if the formalisation is meant to be taken as the more accurate version.
- *Failing to use general concepts to achieve uniformity.* Both TOVE and EO fail to use top concepts to describe in a uniform way core relations and concepts. This hampers understanding. Typical examples are the *part-of* relation, used in describing the decomposition of organizations into smaller units, and the relation, which shows the different ways for *participation* in organizations. For example, TOVE introduces two kinds of *part-of* relations: *org-unit* (between *ORGANISATION* and *ORGANISATION-UNIT*), and *unit* (between two *ORGANISATION-UNIT*s). These relations express *ORGANISATION* and *ORGANISATION-UNIT* decompositions, but are not explicitly unified under a common relation. In the EO several ways of participating in a company are considered, as a partner (*partner_of* relation between *PERSON* and *PARTNERSHIP*), as an employee (*works_for* relation between *PERSON* and *OU*), as a shareholder in a corporation (only in the informal EO specification, see [14]). These ways of participation are not unified in the EO.

- *Insufficient analysis.* As an example consider the EO concepts of *OWNERSHIP* and *SHAREHOLDING* (see [14]) which are formally unrelated, while *SHAREHOLDING* as evident from its informal and formal definitions represents the ownership relation between a *CORPORATION* and its owners.

3.2. Common view of an organization

Figures 1 & 2 give a broad picture of the concepts included in the analysis of TOVE and EO. As even a cursory glance can tell there are significant differences.

There are many examples in both TOVE and EO of how a better analysis would have led to more similar views:

- *Insufficient analysis.* In TOVE, for example, it seems that an *ORGANISATION* is not an *AGENT*, but has *AGENTS* as members. Yet there are many examples of organisations (such as the EU or NATO), which have other organisations as members.
- *Missing Links.* In the EO, the relation between the concepts *OU* and *LEGAL ENTITY* is unclear. All that we are told is that a *LEGAL ENTITY* "may correspond to a single *OU*" (see [14]). No further analysis (informal or formal) of the link between these two concepts is given.
- *Implicit context dependencies.* In the EO, the concept *LEGAL ENTITY*, is not well thought out – having several (informally inconsistent) descriptions. It seems that the intended meaning actually depends on a particular jurisdiction (in this case on the current UK jurisdiction) – though it is not clear that the authors recognise this. This dependence is inappropriate in the modern global economy – and it raises potential problems should the UK jurisdiction change. For example, the *LEGAL ENTITY* concept would no longer be the "union of *PERSON*, *CORPORATION*, and *PARTNERSHIP*".

3.3. Unifying the Core Concepts: Person

Part of the synthesis work is to analyse the ontologies in preparation for a synthesised common view. A vital missing element from both the ontologies is a unifying core category.

To resolve this, we have introduced the concept *PERSON* (*PARTY*), which can be a *NATURAL PERSON* or *SOCIALLY CONSTRUCTED PERSON* (*SOCIAL PERSON* in short). This acts as the catalyst for transforming the ontologies into ones with similar characteristics. The next step (which we will undertake soon) is to merge them into a single synthesised ontology.

The result of introducing *PERSON* into the EO ontology is shown in Figure. 3. A comparison of this with Figure 1 shows how *PERSON* has unified the taxonomy.

To give the reader some idea of how the transformation was effected, we describe the steps we went through. The

EO concepts *LEGAL ENTITY* and *OU* are generalized into the concept *PERSON*. The EO concept *PERSON* (human being) is renamed into *NATURAL PERSON*. *OU* becomes *SOCIAL PERSON*, while *LEGAL ENTITY* is taken completely out and substituted with the context independent notion of *LEGALLY CONSTRUCTED PERSON* (*LEGAL PERSON* in short).

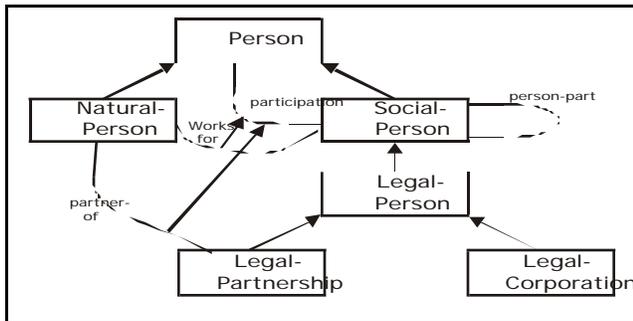


Figure 3. EO transformation

Note that *LEGAL PERSON* is not the same concept as the EO *LEGAL ENTITY*, since it is intended to represent parties which are constructed according to a legal jurisdiction, but not necessarily recognised by it as legal persons (in EO terms, *LEGAL ENTITIES*). For example, in UK a partnership is not legally recognised as a person (it cannot sign contracts in its name) but it is a *LEGALLY CONSTRUCTED PERSON*, because there are legal constitution rules for partnerships. Finally the two participation relations, *partner_of* and *works_for* are consolidated under a general *participation* relation, and the relation *manages* is renamed into *person-part* (which is a particular kind of *part_of* relation).

The result of introducing *PERSON* into the TOVE ontology is shown in Figure 4. As before, a comparison of this with Figure 1 shows how *PERSON* has unified the taxonomy. The transformation steps between Figure 2 and Figure 4 are similar in many respects to those between Figures 1 and 2.

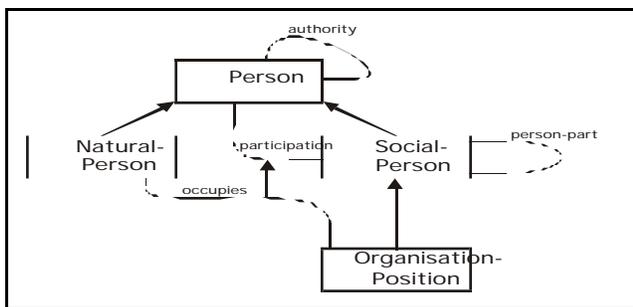


Figure 4. TOVE transformation

4. Conclusion

Even at this early stage our work has revealed the need for a substantial improvement in enterprise ontologies to

bring them up to 'industrial strength'. Hopefully, our work will go some way towards realising this.

5. Acknowledgements

We would like to thank to the IKF Project in general and to ELSAG SpA in particular for making this research possible. Furthermore we would like to thank to Alessandro Oltramale, Claudio Masolo, and Nicola Guarino for the numerous fruitful discussions we had on topics related to ontologies and organisations.

6. References

- [1] Breuker J, Valente A., Winkels R. *Legal Ontologies: A Functional View* in P.R.S. Visser and R.G.F. Winkels, Proceedings of the First International Workshop on Legal Ontologies, 1997
- [2] CYC - <http://www.cyc.com/publications.html>
- [3] EO, <http://www.aii.ed.ac.uk/project/enterprise/enterprise/ontology.html>
- [5] Fox, M.S., Chionglo, J., Fadel, F., *A Common-Sense Model of the Enterprise*, Proceedings of the Industrial Engineering Research Conference 1993.
- [6] Fox, M.S., Barbuceanu, M., and Gruninger, M., *An Organisation Ontology for Enterprise Modelling: Preliminary Concepts for Linking Structure and Behaviour*, Computers in Industry, Vol. 29, pp. 123-134.
- [7] Gruber, T., *Toward Principles for the Design of Ontologies Used for Knowledge Sharing*, in *Formal Ontology in Conceptual Analysis and Knowledge Representation*, Nicola Guarino and Roberto Poli ed., 1993.
- [8] Guarino, N., *Semantic Matching: Formal Ontological Distinctions for Information Organization, Extraction, and Integration*. In M. T. Pazienza (ed.) *Information Extraction: A Multidisciplinary Approach to an Emerging Information Technology*, 1997.
- [9] Hay David C., *Data Model Patterns: Conventions of Thought*, Dorset House, 1997.
- [10] Inmon W.H.. *The Data Model Resource Book: A Library of Logical Data and Data Warehouse Models*, John Wiley and Sons, 1997.
- [11] Partridge, C. *Business Objects: Re-Engineering for Re-Use*, Butterworth-Heinemann, 1996.
- [13] TOVE - <http://www.eil.utoronto.ca/tove/>
- [14] Uschold, M., King, M., Moralee, S., and Zorgios, Y. *The Enterprise Ontology*, AIAI, The University of Edinburgh, 1997
- [15] Uschold, M., King, M., Moralee, S., and Zorgios, Y. *The Enterprise Ontology*, in *The Knowledge Engineering Review*, Vol. 13, eds. M. Uschold and A. Tate, 1998