Synthesising an Industrial Strength Enterprise Ontology

Chris Partridge¹, and Milena Stefanova²

¹*The BORO Program*, ²*LADSEB CNR, Italy*, partridge@BOROProgram.org

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 an ‘industrial strength’ 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) (BORO, <http://>) and European IKF (Intelligent Knowledge Fusion) (EUREKA, <http://>) 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,

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: [10.1007/978-0-387-35621-1_43](https://doi.org/10.1007/978-0-387-35621-1_43)

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 Bank-Lab 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* (Breuker, et al, 1997). 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 synthesis work is large – and so the work has been divided into more manageable chunks.

As Breuker, et al, (1997) 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 (Guarino, 1997, Gruber, 1993, Partridge, 1996).

This gave us the following list:

- Toronto Virtual Enterprise - TOVE (Fox, et al, 1993 & 1996, TOVE, <http://>),
- AIAI's Enterprise Ontology - EO (EO <http://>, Uschold, 1997 & 1998),
- Cycorp's Cyc® Knowledge Base - CYC (CYC, <http://>),
- W.H. Inmon's Data Model Resource Book - DMRB (Inmon, 1997, Hay 1997).

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 specialised 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.

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:

- 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.

3.1 Clear Characterisation

```
graph TD; OU[Organisation Unit] -- manages --> PA[Potential Actor]; P[Person] -- works for --> PA; P[Person] -- works for --> LE[Legal Entity]; P[Partnership] -- works for --> LE; C[Corporation] -- works for --> LE; OU[Organisation Unit] -- works for --> P[Partnership];
```

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

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.

Figure 2: Simplified TOVE overview

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 (common-sense) use of the concept it is not.
- *Messy formalisation trajectories*. EO formalises 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 (Uschold, et al 1997), 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 formalisation. Furthermore some *RELATIONSHIPS* are defined in the formalisation as classes and others are defined as relations without explaining what the motivations for these choices are (e.g., *SALE* is a *RELATIONSHIP* formalised as a class, *HAVE-CAPABILITY* is a *RELATIONSHIP* formalised as a rela-

tion). 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 organisations into smaller units, and the relation, which shows the different ways for *participation in organisations*. For example, TOVE introduces two kinds of *part-of* relations: *org-unit* (between *ORGANISATION* and *ORGANISATION-UNIT*), and *unit* (between two *ORGANISATION-UNITs*). 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, Uschold, et al, 1997). These ways of participation are not unified in the EO.
- *Insufficient analysis.* As an example consider the EO concepts of *OWNERSHIP* and *SHAREHOLDING* (Uschold, et al, 1997), 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 organisation

Figs 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*" (Uschold, et al, 1997). 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.

Figure 3: EO transformation

The result of introducing *PERSON* into the EO ontology is shown in Fig. 3. A comparison of this with Fig. 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 generalised 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).

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).

Figure 4: TOVE transformation

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

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 Allesandro Oltramale, Claudio Masolo, and Nicola Guarino for the numerous fruitful discussions we had on topics related to ontologies and organisations.

6 REFERENCES

- BORO, <http://www.BOROProgram.org>
- Breuker, J, Valente, A., Winkels, R. (1997), *Legal Ontologies: A Functional View* in P.R.S. Visser, and R.G.F. Winkels, Proceedings of the First International Workshop on Legal Ontologies.
- CYC, <http://www.cyc.com/publications.html>
- EO, <http://www.aiai.ed.ac.uk/project/enterprise/enterprise/ontology.html>
- EUREKA, http://www3.eureka.be/Home/projectdb/PrjFormFrame.asp?pr_id=2235
- Fox, M.S. Chionglo, J., Fadel, F. (1993), *A Common-Sense Model of the Enterprise*, Proceedings of the Industrial Engineering Research Conference.
- Fox, M.S. Barbuceanu, M. Gruninger, M. (1996), *An Organisation Ontology for Enterprise Modelling: Preliminary Concepts for Linking Structure and Behaviour*, Computers in Industry, Vol. 29, pp. 123-134.
- Gruber, T. (1993), *Toward Principles for the Design of Ontologies Used for Knowledge Sharing*, in Nicola Guarino and Roberto Poli, (Eds.) *Formal Ontology in Conceptual Analysis and Knowledge Representation*.
- Guarino, N. (1997), *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*.
- Hay, David C. (1997), *Data Model Patterns: Conventions of Thought*, Dorset House.
- Inmon, W.H. (1997), *The Data Model Resource Book: A Library of Logical Data and Data Warehouse Models*, John Wiley and Sons.
- Partridge, C. (1996), *Business Objects: Re-Engineering for Re-Use*, Butterworth-Heinemann.
- TOVE, <http://www.eil.utoronto.ca/tove/>
- Uschold, M. King, M. Moralee, S. and Zorgios, Y. (1997), *The Enterprise Ontology*, AIAI, The University of Edinburgh.
- Uschold, M. King, M. Moralee, S. Zorgios, Y. (1998), *The Enterprise Ontology*, in M. Uschold and A. Tate (Eds.) *The Knowledge Engineering Review*, Vol. 13.