Distributed Ontology Management

Prospects and Pitfalls on Our Way Towards a Web of Ontologies

Oliver Wendt; Tim Stockheim; Sven Grolik; Michael Schwind

JWG University Frankfurt
Chair of Economics, esp. Information Systems
Mertonstr. 17
D-60054 Frankfurt (Germany)
phone: ++49 (0) 69 798 23318
fax: ++49 (0) 69 798 28585
{wendt | stockheim | grolik | schwind}@wiwi.uni-frankfurt.de

Enabling communication between human beings as well as between software agents requires a minimal set of shared knowledge and language. The more the ontologies of the respective partners "overlap", the more efficient the communication process. This paper analyses requirements for the management of distributed ontologies and exemplifies the process of ontology integration for two logistics domains, namely *production logistics* and *hospital logistics*. By simulating an adoption process of standard ontologies in a multi-agent system, we show that the *topology* of the community networks significantly influences the diffusion processes of languages in societies, explaining why, despite of all benefits from standardization, stable clusters of specific ontologies and languages may co-exist, although being perfect substitutes.

1. INTRODUCTION

As social beings, human agents are not living in isolation but rather communicate via language, meant to convey parts of some conceptualization from the sender to a single recipient or a set of recipients. Communities of agents not only share a common *language* but also the individual *conceptualizations* of the world (real and abstract) overlap to a significant extent, allowing for efficient reference to whole conceptual structures like "the German constitution", "game theory" or "medical sciences". Although the agents' individual ontologies associated with these terms may still differ significantly, there is less structural ambiguity in these cases compared to discussing "the sense of life": At least some *core concepts* and their *relations* are agreed to be undisputed and thus offer the basis for dispute over individual, hypothetical extensions[1].

For "societies" of software agents, often called multi-agent systems (MAS), the situation is not quite as babylonian since although the software agents are meant to act individually (and also have a private state and private knowledge) in most cases they are assumed to refer to

---

[1] While all physicians will agree on the heart being connected by blood vessels to other organs and its main function being to keep up blood circulation, considerable dispute arises when it comes to causes and appropriate therapy for a given malfunction.

one common ontology and thus do not have to deal with misunderstanding. However, the more these MAS become interconnected, the more this situation resembles the one described for societies of human agents even though the misunderstanding might be easier to detect when the different reference ontologies are made explicit and published.

Unfortunately, the World Wide Web emerged as a web of linked documents, with a markup language well-suited for specifying *layout* but not *semantics* of the represented information. The term *Semantic Web* therefore refers to the next generation web, trying to overcome these defects by using markup information along with background information on *terms* and *content*, thus allowing machines to "browse" the web's structure and efficiently retrieve the information needed. Therefore, the Semantic Web might rather be imagined as a huge *distributed data base* with interconnected *entity relationship models* [Chen 76] (combining and renaming remote and local entities with remote and local relations) than a "flat" web of hyperlinked documents.

While the term *ontology* itself has been used with different connotations, ranging from simple *taxonomies* to *theories* specified in formal logics, it seems undisputed that the ontology is a *specification of a conceptualization* [Gruber 93] [Sowa 00]. Of course, the fact that a more or less formal interpretation of these two terms leads to a more or less fuzzy structure of represented knowledge, should not surprise after all and, ironically, points to a major challenge a system of distributed ontology management has to cope with. [OWL 02] specifies its main elements to comprise:

- *Classes* (general things) in the many domains of interest
- The *relationships* that can exist among things
- The *properties* (or attributes) those things may have

> "The next element required for the Semantic Web is a Web ontology language which can formally describe the semantics of classes and properties used in web documents. In order for machines to perform useful reasoning tasks on these documents, the language must go beyond the basic semantics of RDF Schema." [OWL 02]

Six use cases are given by the W3C Web Ontology working group to motivate such an endeavor:

- *Web portals* have to provide a starting point for (human of automated) access to various other sources of information with different ontologies.
- *Multimedia collections* require the specification of meta data and "part-whole structures" for efficient access.
- *Corporate web sites* and *Content Management Systems* are integrating information from divers domains and requiring parametric search.
- *Design documentation* usually has a hierarchical structure, but the structure of cross-linked documentation sets may differ; nevertheless they document common objects. (The same is certainly true for the medical domain, in which almost all documents refer to a restricted number of common object classes and their interrelations.)

- *Intelligent (software) agents* are using information on personal preferences to retrieve information on appropriate offers for services and goods and initiate the matching processes. This will require integration of multiple separate ontologies across different domains and services.
- *Ubiquitous computing* requires devices to discover and identify other devices providing services needed and initiate *ad hoc networks* with these devices. When there is no central network provider available, the devices have to be able to autonomously negotiate their mode of interoperability.

While these use cases are certainly far from giving a complete coverage of the set of possible application domains, it should be clear that the agent-based systems developed by the research program "Intelligent Agents and Realistic Application Scenarios" for the two domains *production logistics* and *hospital logistics* certainly involves at least one of them.

2. CHALLENGES OF DISTRIBUTED ONTOLOGY MANAGEMENT

Although management of *distributed databases* has been a well-researched area [Couloris 01], identifying *scalability*, *availability*, *reliability* and *performance* as the main criteria of system performance, the process of compiling requirements for a language and tools for managing *distributed knowledge* in terms of *distributed ontologies* (without any ex ante schema integration) is a pretty recent process which just yielded a first specification of requirements by the W3C's Semantic Web activities [OWL 02]. In the sequel we will specify requirements for distributed ontology management from the perspective of our research program's community process (comprising the interaction of the programs *researchers* as well as the interaction of the *software agents* deployed to the respective application domains).

2.1 Interoperability of shared ontologies

Ontologies have to be made available *publicly* or at least to a group of users who may then

- *commit* to it in order to share its semantics
- *explicitly augment / extend* the ontology with additional conceptual structures and terminology for the intended application domain; this extension relation has to be transitive, i.e. whenever ontology C is an extension of ontology B which in turn extends ontology A, C is considered to be an extension of ontology A.
- *shadow* a subset of its concepts which have to be redefined for the application domain. However, the question of granularity of this shadowing needs further investigation, since the resulting "patchwork" ontologies could lead to interoperability problems due to homonyms.

Whenever a term has partial definitions in multiple schemas, RDF [RDFS 02] assumes the definition to be the union of all descriptions using the same identifier. In a distributed environment, comprising possibly incorrect or false definitions, the absence of a mechanism to specify which set of definitions a resource agrees on may cause serious problems.

As already mentioned, the web of ontologies cannot be expected to develop around a set of common core ontologies but rather by a growing interaction of loosely coupled individuals and communities of these individuals. The success of an Ontology Web Language (OWL) will strongly depend on the facilitation of this coupling process in terms of support for explicit *mapping of synonyms* and the constructs offered for *merging* and *differentiating* “packages” of knowledge as well as automatic *identification of isomorphic cores* of two or more ontologies, suggesting a generalization into a common super-ontology (thereby initiating an interaction process of the respective communities, which might not even have imagined talking about the “same”, since their knowledge has been disguised by different vocabulary).

To achieve these goals, at least the following requirements should be met:

- **Ontologies as distinct objects:** An OWL should allow for ontologies to be identifiable as distinct objects by unique identifiers (e.g. an URI reference).
- **Unambiguous term referencing with URIs:** While homonyms have the identical relative identifier within their respective ontologies, they must be given different global names by different absolute identifiers (e.g. an URI reference).
- **Class definition primitives** have to be provided by the language including sub-classing and Boolean combinations of class expressions. The language has to allow for explicitly treat *classes as instances* of an other class. This must not be confused with a subclass relation: While the classes teacher, lawyer and physician may be modeled as instances of the class profession for an unemployment statistics database, it is clear that no individual teacher “is a” profession. Exactly this fact would arise from making teacher a subclass of profession.
- **Property definition primitives:** Properties are comparable to the binary[2] relations called *slots* in frame-based meta ontologies. While *domain*, *range* and *cardinality* constraints must be available, *transitivity* and *inverse* properties should be provided as well as the possibility to specify user-defined “facets” or “tags” like measures of uncertainty or default values for a given property. The so-called *chaining* of properties (e.g. yielding the *grandfather* relation by chaining the *father* property with itself) should also be supported.
- **Complex data types:** It seems clear that the OWL should deal with primitive data types (based on XML Schema data types) as well as with structured complex ones like addresses or dates.
- **Procedures and functions**: The requirements draft is not very specific about abstract data types, i.e. objects whose properties may be methods or functions like in object-oriented programming. On the one hand, *arithmetic primitives*, functions for *string manipulation* (concatenation and simple pattern matching) and *aggregation and grouping* (comparable to SQL's grouping and counting commands) are required. On the other hand, the requirement of *procedural attachment* remains very vague about the form of integrating executable code: “Procedural attachments greatly extend the expressivity of the language, but are not well-suited to formal semantics. A procedural attachment mechanism for web ontologies should specify how to locate and execute the procedure. One potential candidate language would be Java, which is already well-suited to intra-platform sharing on the Web.”[OWL 02] This rises the question whether the languages and tools for the Semantic Web is meant to finally converge towards a global distributed software engineering, runtime and reasoning environment or whether interoperable execution of code should remain a “different story”?

---

2 Although every *n*-ary relationship over *n* domains may be broken down to a set of binary relations (e.g. over newly created dummy objects representing the tupels of the initial *n*-ary structure), this endeavor might lead to very inefficient data access compared to a relational data base. No clear strategy for specifying and handling relations of higher arity may be deduced from the current requirements specification of the W3C.

2.2 Ontology evolution

As a dynamic system, the web will constantly require changes of published ontologies. In order to avoid any "ontology not found" errors, it should be clear that a released ontology must never be overwritten but rather a revised version will have to be published. However, this leaves it up to the user, whether he builds his extensions on the most up-to-date version or rather relates to a specific one. Since intended meanings of terms may change without altering its formal description, authors will have to be given the possibility to explicitly mark such changes and also explicitly state backward-compatibility. The current method proposed by RDF Schema to use a separate URI for each version does not allow for retracting (incorrect) definitions.

Equivalence statements: To achieve this, the language must include explicit *equivalence statements* for classes, properties and individuals. The use of a standard URL is not enough, since, like a person may be identified by multiple ID cards or passports, the same is true for objects on the web and an identity relation is the most primitive help to identify what is not only an equal object but really the same one.

Local unique names assumptions: Although no name of a concept can be assumed to be unique it is very helpful to explicitly *state* that all names within a given name space refer to distinct objects whenever this is true for a given application.

Inconsistency detection: Due to the distributed nature of the interconnected ontologies, contradictions may not be identified ex ante. There should rather be mechanisms to detect the inconsistencies that may arise in any ontology specification or extension task und thus prevent automated agents from using any combination of incompatible specifications. Up to now, there is no way to express inconsistencies in RDF and RDFS.

2.3 Expressivity versus scalability

The larger the set of primitives and thus the expressivity of a language, the higher the *variety* of knowledge that may be represented. On the other hand, the need for *efficient automated reasoning* and scalability puts constraints to this variety, i.e. requiring the language to be decidable. A web ontology language must therefore find an appropriate balance, certainly more on the *expressive* side than current RDF.

It is not even clear, if this balance has to yield a unique standard language or whether it is preferable to just specify the interfaces and protocols allowing for easy incorporation of new representational features and interoperation of various models like RDF(S) [RDFS 02], OIL [Fensel et al. 01], DAML [Hendler 00], and DAML+OIL [Connolly 01] [Horrocks 02].

Default and uncertain property values: Although default values for properties might be considered useful in many contexts, it is known that this requires capabilities of non-monotonic reasoning. In a distributed environment like the Web new information is constantly added, updated or retracted in many locations, i.e. there is a need for research, how the findings from AI research may be extended to a heterogeneous network, having to map even different systems of default logic or uncertainty management.

Ability to state closed worlds: While the assumption that anything will be assumed false that cannot not be derived to be true certainly is a bad one for a distributed growing information system, it may be appropriate for closed subsystems, assumed not to change their structure any more (and thus allowing to know for sure, that an object does *not* have a given property, if not *specified* and thus allowing to draw conclusions accordingly).

2.4 Ease of use

A related trade-off concerns the user, now having to act as a "knowledge engineer" rather than a "layout engineer". Even if the availability of front-end tools might help the user to get a graphical access to the ontology, the *language* should be independent from syntax and will have to have clear primitives. Nevertheless, it should be easy to learn and tools should be usable by domain experts and business analysts as well as by "ontologists".

The editor's *easy of use* conflicts with the *expressiveness* of the formal languages supported. The social incentives for coding one's knowledge into some "machine readable" form must be high enough to guarantee adoption. Whenever the individual costs exceed the individual benefits, the decision not to do it is clear, no matter how high the social benefits might be.

2.5 XML syntax

Since XML is widely-used for interoperability and information sharing between information systems, the language needs to have an XML serialization in order to provide the exchange of data and ontologies and in a standard format. The syntax of RDF already complies with this requirement.

2.6 Internationalization

Of course, global usability of an ontology and support for different natural languages requires alternative user-displayable labels for its objects to be specifiable as well as support the use of multilingual character sets. Although many formal conceptual structures are independent of cultural context und thus easily translated, for others the translation may be misleading, if there is no direct correspondence: While e.g. the spouse relation is 1:1 in the western countries, this constraint has to be relaxed in some Arabic countries. However, in contrast to spoken or written language the translated vocabulary will now still be referring to the original conceptual structure, allowing for a clear distinction.

2.7 Rule Markup Techniques

Most W3C activity relating to the Semantic Web is concerned with vocabulary (taxonomy) markup. *Rule markup techniques* should complement the OWL and explore the suitability of rule systems (like e.g. extended Horn logics) for web application. A recent Dagstuhl Seminar "Rule Markup Techniques"[3] was dedicated to contribute to rule-related issues in recent proposals such as Notation 3 (N3)[4], DAML-Rules[5], and the Rule Markup Language RuleML[6]

---

[3] refer to http://www.dfki.uni-kl.de/ruleml/rmt/

[4] refer to http://www.w3.org/2000/10/swap/Primer.html

[5] refer to http://daml.umbc.edu/tasktwo

[6] refer to http://www.dfki.uni-kl.de/ruleml/

and the combination of rules and taxonomies via sorted logics, description logics, or frame systems.

2.8 Security Management

Of course, all questions of security management already relevant in today's internet will also be relevant in a semantic web and although a more structured web might give rise to more specialized and thus more efficient security concepts, we do not see any plausible reason why security problems should become significantly worse or easier. A related question is how *trust* will be modeled and administered in a heterogeneous environment and under which restrictions it could be considered to be transitive.

3. TOWARDS A LOGISTICS ONTOLOGY: INTEGRATION VERSUS INTEROPERABILITY

Within the above-mentioned research program "Intelligent Agents and Realistic Application Scenarios" several research teams have modeled or are currently modeling the conceptual structures and ontologies their software agents will rely on. The following diagram shows the meta-structure, the teams involved with specific production logistics scenarios have already agreed upon:

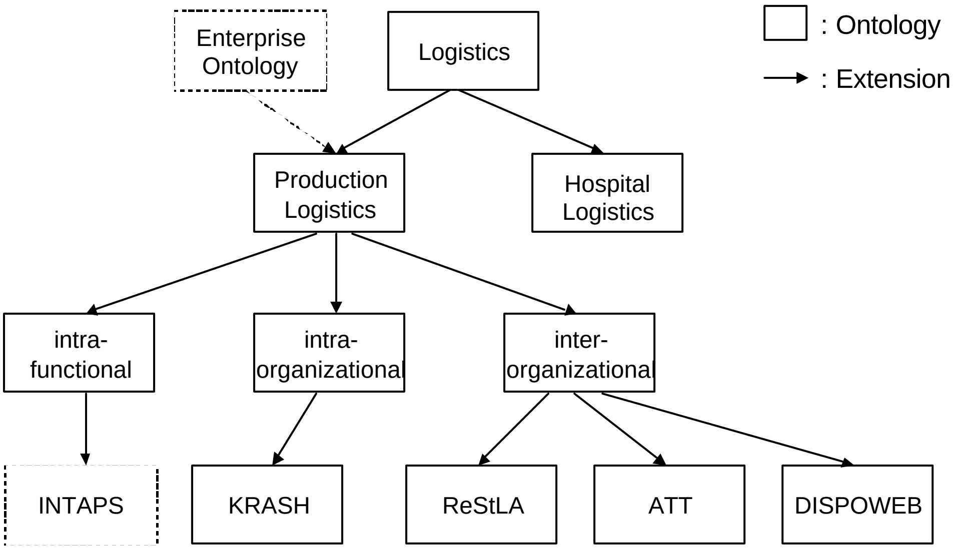

Figure 1: Logistics Ontology and Extensions

Our objective is the development of a consistent and axiomatized ontology of logistics, which can serve as a basis for ontologies in the domains *production logistics* and *hospital logistics*. In order to do achieve this goal, we identified three application areas within the domain of production logistics the individual SPP projects are dealing with. The INTAPS[7] project is

[7] IntaPS (Integrated Agent-based Process Planning and Production Control) is carried out by the Center for Computing Technologies (TZI) of the University of Bremen and the Institute of Production Engineering and Machine Tools (IFW) of the University of Hannover (http://intaps.ifw.uni-hannover.de).

dealing with problems on an intra-functional level, while the KRASH[8] project is dealing with application problems on an intra-organizational level and the projects ReStLA[9], ATT and DISPOWEB[10] on an inter-organizational level. For the development of a joint ontology in the production logistics domain, the individual projects ontologies were integrated into the Enterprise Ontology [Uschold 98] and their concepts and vocabulary compared against each other. In a first step this allows the analysis of differences as well as the identification of common characteristics of the sub-ontologies.

Since the Enterprise Ontology, which provides the most comprehensive set of classes and relations in the business domain, is available in Ontolingua, it is a straightforward choice to evaluate Ontolingua's aptitude to serve as a platform for distributed ontology management.

- **Ontolingua:** Developed in the early 90s at Standford's Knowledge Systems Laboratory, Ontolingua provides probably the most expressive knowledge representation support based on a frame language which is supporting the (first order logic) Knowledge Interchange Format KIF [Genesereth 92]. Ontolingua combines advances representational features with a theorem prover and reasoning capabilities. Although meeting most of the requirements of ontology merging and evolution Ontolingua comes with a major drawback: It uses a centralized server based on a (proprietary) Common Lisp implementation for hosting and accessing the ontologies, i.e. there is currently no way to use it for distributed ontology management. Although the Ontolingua server provides an OKBC interface, features like RDF exports are not available yet and it is not clear, if there will be an further development of the server at all.

- **Protégé 2000:** Also developed at Stanford University, Protégé's primary focus was not on ontology engineering but rather on facilitating user-friendly design of knowledge-acquisition interfaces. But due to a component based software architecture, several plugins have become available for specific extensions, like an OIL-Plugin with a FACT/Shiq reasoner. The semantics of OIL may be translated into SHIQ description logic [Horrocks 00], a very expressive concept language fully capturing the OIL core language. Since Protégé is freely available[11] for research purposes for Windows and Linux and (in contrast to Ontolingua) Protégé may be used as a generic editor for RDF/RDFS[12]. Unfortunately, although it comes with an OKBC interface, there is no way for an automatic translation of the enterprise ontology from Ontolingua to Protégé. While this could not be expected due to the fact that Protégé does not understand Ontolingua axioms specified in KIF, a lot of "manual" editing was necessary even for the basic skeleton of classes and relations.

---

8 KRASH (Karlsruhe Robust Agent Shell) is a joint project at the Computer Science Department of the University of Karlsruhe, Germany. Involved are the Institute for Process Control and Robotics and the Institute for Program Structures and Data Organisation (IPD), Information Systems Group (http://www.ipd.uka.de/KRASH/)

9 ReStLA (Reaktive Steuerung von Lieferketten mit Agentensystemen) is carried out by the DAI-Labor at Technische Universität Berlin and the Institute for Planning and Control of Production Systems - BIBA at the Universität Bremen.

10 DISPOWEB (Dispositive Supply Web Coordination) is carried out by the Institute of Information Systems at Frankfurt University and the Faculty IV Electrotechnics and Information Technologies at Technischen Universität Berlin (http://www.dispoweb.org)

11 protege.stanford.edu

12 Like for Ontolingua and (many other more primitive tools) internationalization or versioning support as well as scalability and performance for commercial applications is beyond the current scope of Protégé.

3.1 Integration of Production Logistics Ontologies

To facilitate interoperability, we agreed on using OIL (Ontology Inference Layer)[13] as a uniform ontology language for all projects. Within the scope of these translation processes some minor inconsistencies within the Enterprise Ontology were solved[14].

The integration and the comparison of the individual research projects' ontologies into the Enterprise Ontology was realized as follows: At first, one class was generated with the project name for each project-ontology. Afterwards, all classes of the respective project, which do not "compete" with any EO classes, were made sub-classes of the class with the corresponding project name (cp. Figure 2) and incorporated in the class hierarchy of the EO.

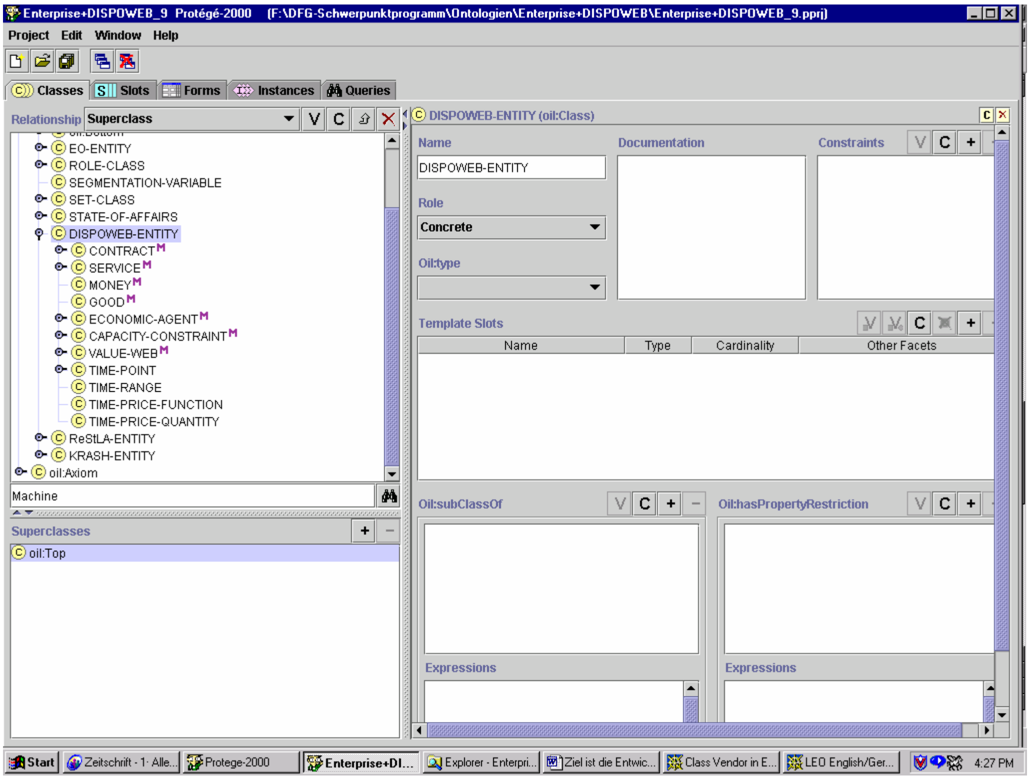

Figure 2: Ontologies of single projects

In the case of classes "competing" directly with EO classes or even having the same name as an EO class, the class name is preceded with the project name separated with a hyphen and integrated as a sub-class of the respective „competing" EO class (cp. figure 3). In addition, these classes were integrated into the respective projects' class hierarchy at their proper place.

---

13 www.ontoknowledge.org/oil

14 For example, on the one hand the class COMPETITOR is a sub-class of the class VENDOR, which is in turn a sub-class of the class QUA-ENTITY and therefore an indirect sub-class of the class QUA-ENTITY, but on the other hand a direct sub-class of the class QUA-ENTITY at the same time. A similar problem arises with the class SUB-PLAN, which is a sub-class of the class PLAN and therefore indirect a sub-class of the class QUA-ENTITY, but on the other hand a direct sub-class of the class QUA-ENTITY at the same time. For the solution of these problems the direct subclass relation of COMPETITOR and SUB-PLAN to the class QUA-ENTITY was deleted.

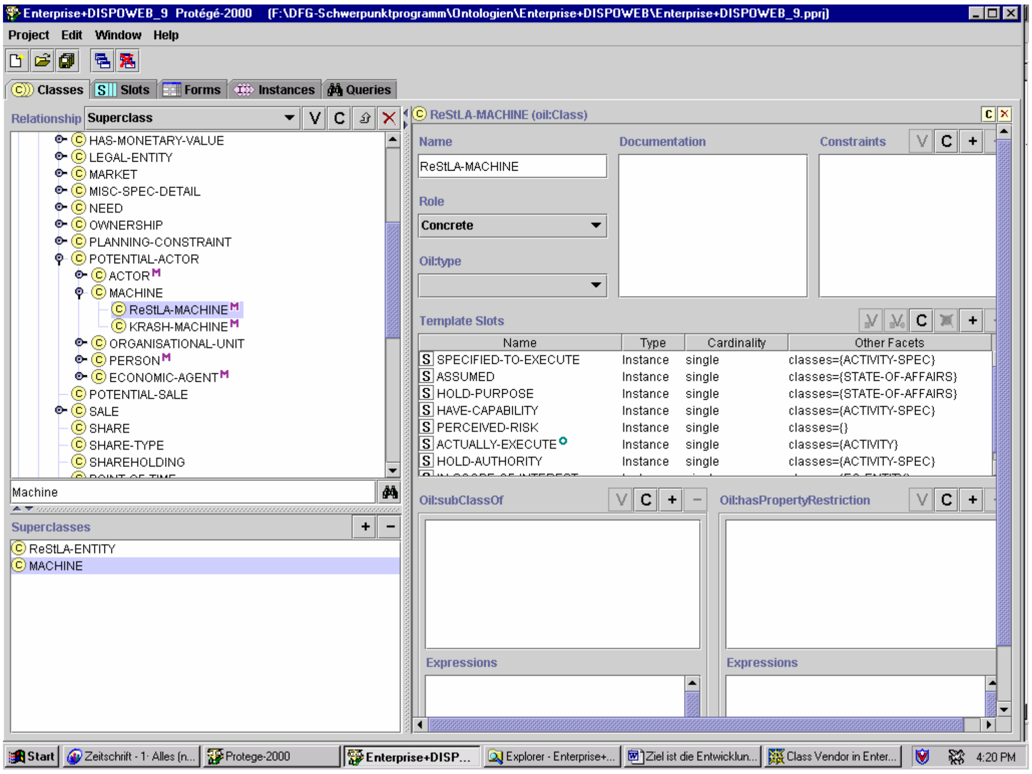

Figure 3: Project dependant MACHINE-class

After integration of the *classes* of the project ontologies the project specific *slots* were integrated at their corresponding positions. When doing this the slot name is preceded with the project name separated with a hyphen.

Another problem arises if there is the necessity to introduce a super-class in one of the sub-ontologies. E.g. the DISPOWEB ontology requires different types of contracts under a super-class for contracts. EO, however, only provides an EMPLOYMENT-CONTRACT, which is a direct subclass of EO-ENTITY and not subsumed under any more generic contract super-class. Two approaches are conceivable to resolve this conflict: Because we defined EO as a superset of the project ontologies, the introduction of a super-class DISPOWEB-CONTRACT into EO by the sub-ontology can be considered as a first solution method, we call the ***integration-approach*** (cp. figure 4). Alternatively, the EO can be left untouched and the EMPLOYMEMT-CONTRACT has to be duplicated as DISPOWEB-EMPLOYMEMT-CONTRACT. This process is called the ***interoperability-approach*** (cp. figure 5).

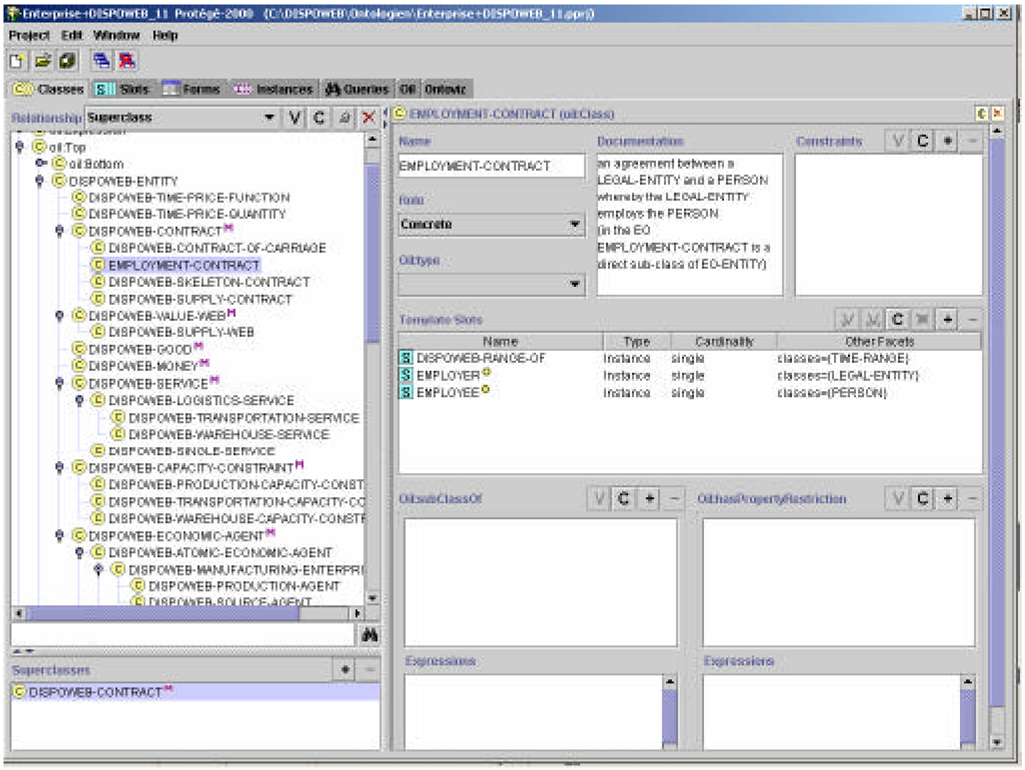

Figure 4: Insertion of a super-class in the EO by the *integration-approach*

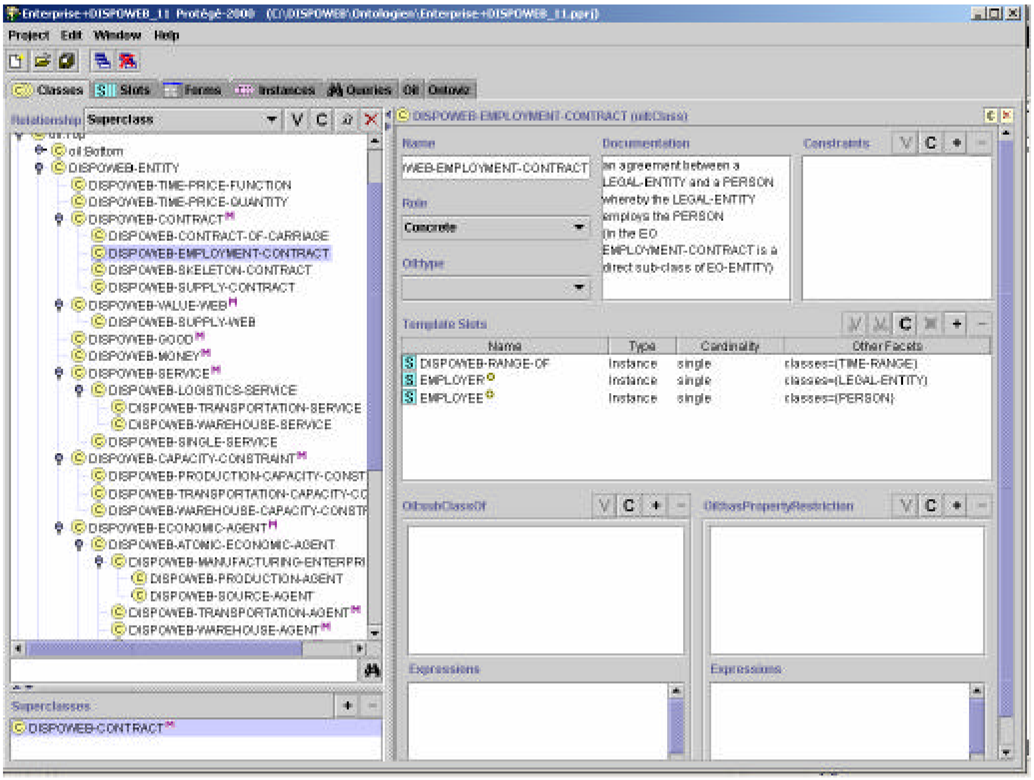

Figure 5: Creation of a project-specific class for DISPOWEB-ontology by the *interoperability-approach*

Since the integration of multiple ontologies can lead to semantic systems of large size, visualization tools seem to be helpful to tackle this problem. One of these tools is OntoViz[15], a Protégé plugin, which can depict classes and slots using tables an directed edges. Especially the ability to visualize sub-ontologies separately has been a valuable support. In figure 6 a view on the class tree of the DISPWEB sub-ontology generated by OntoViz is presented.

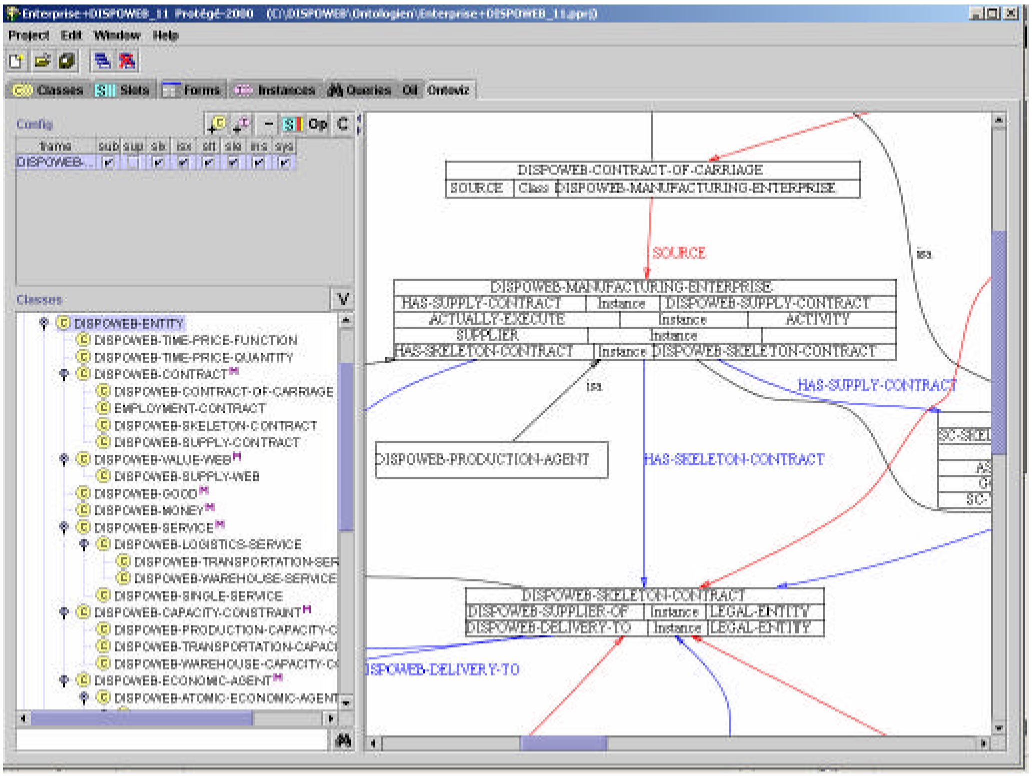

Figure 6: Representation of the DISPOWEB sub-ontology with the OntoViz tool

On the other hand, consistency checking gets of ontologies is possible by using the already mentioned FACT/SHIQ Reasoner, which also generates a latex output of the SHIQ description logic (cf. figure 7)

[15] http://protege.stanford.edu/plugins/ontoviz/ontoviz.html

$$
\begin{array}{rcl}
\text{DISPOWEB-LOGISTICS-AGENT} & \sqsubseteq_c & \text{DISPOWEB-ECONOMIC-AGENT} \\
 & & \sqcap\, \forall\text{HAS-SKELETON-CONTRACT.Class} \\
 & & \sqcap \leqslant 1\ \text{HAS-SKELETON-CONTRACT}.\top \\
 & & \sqcap\, \forall\text{ACTUALLY-EXECUTE.ACTIVITY} \\
 & & \sqcap \leqslant 1\ \text{ACTUALLY-EXECUTE}.\top \\
\text{DISPOWEB-LOGISTICS-SERVICE} & \sqsubseteq_c & \text{DISPOWEB-SERVICE} \\
\text{DISPOWEB-MANUFACTURING-ENTERPRISE} & \sqsubseteq_c & \text{DISPOWEB-ATOMIC-ECONOMIC-AGENT} \\
 & & \sqcap\, \forall\text{HAS-SUPPLY-CONTRACT.DISPOWEB-SUPPLY-CONTRACT} \\
 & & \sqcap \leqslant 1\ \text{HAS-SUPPLY-CONTRACT}.\top \\
 & & \sqcap\, \forall\text{SUPPLIER}.\bot \\
 & & \sqcap \leqslant 1\ \text{SUPPLIER}.\top \\
 & & \sqcap\, \forall\text{HAS-SKELETON-CONTRACT.DISPOWEB-SKELETON-CONTRACT} \\
 & & \sqcap \leqslant 1\ \text{HAS-SKELETON-CONTRACT}.\top \\
 & & \sqcap\, \forall\text{ACTUALLY-EXECUTE.ACTIVITY} \\
 & & \sqcap \leqslant 1\ \text{ACTUALLY-EXECUTE}.\top \\
\text{DISPOWEB-MONEY} & \sqsubseteq_c & \text{GOOD-SERVICE-OR-MONEY} \sqcap \text{DISPOWEB-ENTITY} \\
\text{DISPOWEB-PRODUCTION-AGENT} & \sqsubseteq_c & \text{DISPOWEB-MANUFACTURING-ENTERPRISE} \\
\text{DISPOWEB-PRODUCTION-CAPACITY-CONSTRAINT} & \sqsubseteq_c & \text{DISPOWEB-CAPACITY-CONSTRAINT} \\
\text{DISPOWEB-RESELLER-AGENT} & \sqsubseteq_c & \text{VENDOR} \\
 & & \sqcap\, \text{CUSTOMER} \\
 & & \sqcap\, \text{DISPOWEB-ATOMIC-ECONOMIC-AGENT} \\
 & & \sqcap\, \forall\text{HAS-SKELETON-CONTRACT.Class} \\
 & & \sqcap \leqslant 1\ \text{HAS-SKELETON-CONTRACT}.\top \\
\text{DISPOWEB-RETAILER-AGENT} & \sqsubseteq_c & \text{DISPOWEB-RESELLER-AGENT} \\
\text{DISPOWEB-SERVICE} & \sqsubseteq_c & \text{GOOD-SERVICE-OR-MONEY} \sqcap \text{DISPOWEB-ENTITY}
\end{array}
$$

Figure 7: Cutout of FACT/SHIQ modal logic representation of DISPOWEB ontology

3.2 From production logistics to a unified logistics ontology

Up to now, only ontologies from the production logistics branch of our ontology tree have been considered but core concepts of logistics are expected to be the same in the health care domain. Following the standardization effort of the Health Level Seven (HL7) Consortium (http://www.hl7.org) we try to identify differences as well as common concepts of manufacturing logistics scheduling and scheduling within the healthcare domain. The HL7 effort should not be misunderstood to provide a “plug and play” specification: the very nature of the diverse business processes that exist within the healthcare system prevents the development of either a universal process or a data model to support the definition of HL7’s target environments. In addition, HL7 does not make any assumptions about the architecture of healthcare information systems nor does it attempt to resolve architectural differences between healthcare information systems.

The current HL7 Version does not, in itself, attempt to define or even support the implicit and explicit relationships between persons such as patients, physicians, providers. Nevertheless, the HL7 Reference Information Model can provide a suitable framework for an ontology specification.

With respect to scheduling problems, the HL7 data model is very similar to traditional appointment books: for instance every column for a room is equivalent to a “schedule” for this resource which is divided into time slots of a predefined duration. If necessary, the “appointments” can occupy one or more slots.

In the sequel we describe some examples (mainly classes and roles), which are specific to the healthcare domain and their general class, which could be a class in a generalized logistics ontology:

General class	*Medical logistics subclass (HL7):*	*Production logistics subclass (EO):*
Supply: Supply orders and deliveries are simple acts that focus on the delivered product. The product is modeled as an attribute of the Supply Act. For general Supply Acts, the precise identification of the Material (manufacturer, serial numbers, etc.) is important. Most of the detailed information about the Supply should be represented using the Material class. If delivery needs to be scheduled, tracked, and billed separately, one can associate a Transportation Act with the Supply Act.	**Diet:** Diet acts are supply acts, with some aspects resembling Substance_administration acts: the detail of the diet is given as a description of the Material associated via an attribute "product". Medically relevant diet types may be communicated in the Entity.cd, however, the detail of the food supplied and the various combinations of dishes should be communicated as Material instances.	
Person: A human being described by its skills and possibly its affiliation to an organizational unit.	**Patients:** is of central concern to all processes this domain. Patients receive appointments from the clinic, and in contrast to the production domain they are objects of the scheduling process and customers of services and resources.	**Partner:** The Person in the Partner-Of Relationship; i.e. a Person who forms part of a Partnership.

For other concepts, similarities are evident in both ontologies neither of them is specified by explicit attributes / slots:

Medical logistics subclass (HL7):	*Production logistics subclass (EO):*
Material: A Material is an Entity that excludes Living_subjects and places. Manufactured or processed products are considered material, even if they originate in living matter. Parts (e.g. organs) derived from living subjects are Material that may need to be tracked through associations with the individual Living_subject from which they were obtained. Examples of Material are pharmaceutical substances (including active vaccines containing retarded virus), disposable supplies, durable equipment, implantable devices, food items (including meat or plant products), waste, traded goods, etc.	**Resource**: The Entity that is used or consumed in the Can-Use-Resource relationship.

Role:	**QUA-Entity:**
A categorization of competency of the Entity that plays the Role as defined by the Entity that scopes the Role. An Entity, in a particular Role, can participate in an Act. Note that a particular entity in a particular role can participate in an act in many ways. Thus, a Person in the role of a practitioner can participate in a patient encounter as a rounding physician or as an attending physician. The Role defines the competency of the Entity irrespective of any Act, as opposed to Participation, which is limited to the scope of an Act. Each role is 'played' by one Entity (the Entity that is in the role) and is usually 'scoped' by another. Thus the Role of 'patient' is played by (usually) a person and scoped by the provider from whom the patient will receive services. Similarly, the employer scopes an Employee role. The identifier of the Role identifies the Entity playing the role. This identifier is formally either issued by the scoping Entity or, at least, formally recognized by the scoping Entity as an identifier of the Role. Attributes of Role are those that are particular to the playing entity while in the particular role.	An EO-Entity that is defined in terms of the role it plays in one or more Relationships. Qua-Entity is the most general Role-Class, • every instance of Role-Class is a subclass of Qua-Entity, • every Qua-Entity is an instance of at least one Role-Class.
Act:	**Activities:**
An action of interest that has happened, can happen, is happening, is intended to happen, or is requested/demanded to happen. An act is an intentional action in the business domain of HL7. Healthcare (and any profession or business) is constituted of intentional actions. An Act instance is a record of such an intentional action.	Something done over a particular Time-Range. The following may pertain to an Activity (taken from the EO): • is performed by one or more Actual-Doers; • is decomposed into more detailed Sub-Activity(s); • Can-Use-Resources; • An Actor may Hold-Authority to perform it; • there may be an Activity-Owner; • has a measured efficiency.

In contrast to the Enterprise Ontology the HL7 effort is not only more detailed, but also defines patterns of an architectural design.

For a most up-to-date version of the ongoing integration process, we may refer the reader to: www.dispoweb.org/logistics-ontology/

4. DIFFUSION OF ONTOLOGICAL STANDARDS

The diffusion of ontologies and languages, like the diffusion of many innovations in the software market is strongly influenced by *positive network effects*, i.e. the willingness to adopt an innovation, learn a new language or switch to a new ontology positively correlates with the number of existing adopters.

On the one hand, *indirect* network effects may result from the fact that the benefit from using a certain ontology increases the more complementary goods and services are available. On the other hand, the need to communicate with partners leads to strong *direct* network effects. The more potential communication partners exist that use an ontology compatible to the own, the greater the chances for an easy and efficient communication (for an analytical analysis of standards in communication networks refer to [Buxmann et al. 99] or [Westarp 00]).

For participants in a communication network it seems likely that the diffusion of ontologies is dependent on the characteristics of their communication partners, the density of the network and the intensity of interaction. For example, recent empirical studies indicate the existence of standardization pressure among business partners and that the direct network effects vary dependent on whether adopters within the individual network of business partners or within the entire market network (installed base) are taken into account [Trunda/Westarp 98], [Westarp 99].

Business networks like in the automobile industry are good examples of this. Generally, we find strong communication in the network around large automobile vendors. This makes it favorable if not necessary to use compatible ontologies. The network structure is dense connecting the vendor with its suppliers who in turn might be connected with suppliers of lower levels. This structure makes it more likely that a certain business vocabulary, probably proposed or enforced by a strong vendor, diffuse faster than they would in other industries with less density and less hierarchical structure. With increasing integration of business partners, common ontologies in this context are not only the basis EDI solutions, but also standard software like enterprise resource planning (ERP) systems (e.g. SAP or Peoplesoft) or new XML-based document management systems.

While the phenomenon of *positive network effects* is analyzed in many economic publications, the existing approaches seem to be insufficient to explain important phenomena like the coexistence of different standards despite strong network effects, small but stable clusters of users of a certain standard or the phenomenon that strong players in communication networks force other participants to use a certain standard.

Searching for appropriate instruments to model the software market, two areas of research seem to be promising.

On the one hand, *theory of positive network effects* analyses the specific characteristics of markets for network effects goods. On the other hand, *diffusion models* focus on explaining and forecasting the adoption process of innovations over time.

4.1 Theory of positive network effects

Existing theoretical models of positive network effects[16] focus on individual buying decisions of goods, marketing strategies of competing vendors, supply and demand equilibria, and welfare implications. The research results offer a good basis for analyzing network effects in a general way. Nevertheless, the models are in many ways not sufficient to model diffusion processes of ontologies, since they abstract from the highly relevant topology of the agents' communication network. In empirical research[17] we find approaches that explicitly analyze software products but no explicit study of network effects of ontologies or languages is

[16] refer to [Rohlsfs 74], [Katz/Shapiro 85] [Farrell/Saloner 86], [Church/Gandal 96] or [Gröhn 99]

[17] refer to [Hartmann/Teece 90], [Gandal 94], [Economides/Himmelberg 95], [Moch 95] or [Gröhn 99]

known to the authors. Most of the research assumes that network effects derive from product characteristics like interfaces for data exchange with other applications. By applying the hedonic pricing approach, the hypothesis that the existence of such attributes increases the willingness to pay for the respective product is statistically proven (for a more comprehensive discussion refer to [Gröhn 99]). The results indicate the existence of network effects in the software market but do not give relevant help for our network diffusion approach.

4.2 Diffusion models

The economic analysis of diffusion forecasts diffusion of innovations in markets. In particular, the question arises which factors influence the speed and specific course of diffusion processes [Weiber 93]. Traditional diffusion models are based on similar assumptions (for a comprehensive overview of the traditional diffusion models refer to [Mahajan et al. 90]). Generally, the number of new adopters in a certain period of time is modeled as the proportion of the group of market participants that did not adopt yet. Three approaches are most common [Lilien/Kotler 83: 706-740], [Mahajan/Peterson 85: 12-26], [Weiber 93]: The *exponential diffusion model* (also *external influence model* or *pure innovative model)* assumes that the number of new adopters is determined by influences from outside the system, e.g. mass communication. The *logistic diffusion model* (also *internal influence model* or *pure imitative model*) assumes that the decision to become a new adopter is only determined by the positive influence of existing adopters (e.g. word of mouth). The *semilogistic diffusion model* (also *mixed influence model*) considers both, internal and external influences. A famous example of the latter is the *Bass model*, which has been used for forecasting innovation diffusion in various areas such as retail service, industrial technology, agricultural, educational, pharmaceutical, and consumer durable goods markets [Bass 69], [Mahajan et al. 90].

Despite the existence of various diffusion models, the approaches are not sufficient to model the diffusion of network effect products. Schoder names three areas of deficits [Schoder 95: 46-50]. First of all, there is a lack of analysis concerning the phenomenon of critical mass. Furthermore, the traditional diffusion models are not able to explain the phenomenological variety of diffusion courses. Third, the models do not sufficiently consider the interaction of potential adopters within their socio-economical environment, how adoption changes relationship to other participants of the system, and how the willingness to pay a certain price changes with an adoption within a certain group.

Help for analyzing the relevance of interaction in diffusion processes comes from research activities in the area of network models of diffusion of innovations [Valente 94]. *Relational models* and *structural models* are two important approaches in this area. Relational models analyze how direct contacts between participants of networks influence the decision to adopt or not adopt an innovation. In contrast, structural models focus on the pattern of all relations and show how the structural characteristics of a social system determine the diffusion process.

Although these models stress the relevance of the agent system's microstructure and use concepts like *opinion leadership*, *group membership*, *personal network density* and *personal network exposure* to explain diffusion processes, they do not account for the abovementioned *positive network externalities*.

4.3 A Diffusion Model with Network Externalities

Basis of our simulation is a simple model of the adoption decision in network effect markets. The terminology is similar to the model of [Katz/Shapiro 85], but we will interpret the terms differently. Let *r* denote the stand-alone utility of an ontology (i.e. the willingness to pay even if no other users in the MAS exists) and *f(x)* denote the additional network effect benefits (i.e. the value of the externality when *x* is the number of other adopters). For reasons of simplification we assume that all network participants have the same function *f(x)*, i.e. their evaluation of network benefits from adopting an ontology is identical. We also assume that the network effects increase linearly, i.e. *f(x)* increases by a certain amount with every new user. The total benefit derived from a given ontology can thus be described by the term $r+f(x)$. Let *p* be the price or the cost to learn or implement the new ontology, then a network agent is inclined to implement an ontology if and only if $r+f(x)-p>0$. In case of *v* competing ontologies, he decides for the ontology yielding the maximum surplus (in case this exceeds 0):

$$\max_{i\in\{1,\ldots,v\}}\{r_i + f(x_i) - p_i\}$$

If the surplus is negative for all *i* then no ontology is adopted.

Unlike most of the existing diffusion models we conducted simulations by modeling *relational* diffusion network in which the decision to adopt is not influenced by the *installed base* within the whole network, but rather by the adoption decisions within the *personal* communication network. The significance of this for the own decision can simply be demonstrated by the following example. Figure 8 shows the communication network environment of agent *A* who wants to choose an ontology that serves her individual needs.

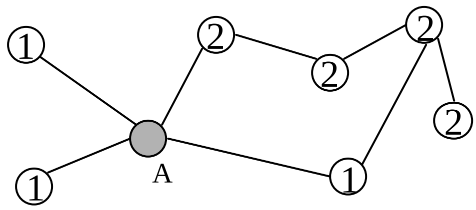

Figure 8. Communication network environment of agent A.

There is a choice of two ontologies (1 and 2) and we assume that both are free of cost and have identical expressiveness, i.e. are perfect substitutes. In traditional models agent *A* would choose ontology 2 since the installed base with 4 existing adopters is larger. In a relational network approach we only focus on the relevant communication partners of *A* who will decide for ontology 1 since the majority of his direct communication partners speak this language.

Our simulations are based on the assumption that network structure, the agents' preferences and the prices of the ontology are constant during the diffusion process. All the results presented below are based on a network size of 1,000 agents, distributed randomly on the unit square, i.e. their *x*- and *y*-coordinates get sampled from a uniform distribution over [0; 1]. In a second step, the network's structure is generated by either choosing the *c* closest neighbors measured by euclidean distance (*close* topology) or selecting *c* neighbors randomly from all *n*-1 possible neighbors (*random* topology). This distinction is made to support our central hypothesis, namely: Ceteris paribus (e.g. for the same network *size* and *connectivity*) the *specific* neighborhood structure of the network strongly influences the diffusion processes.

The graphs in figure 9 show randomly sampled cases of the *close* topology (exemplary for 100 agents and a connectivity *c* of two, five and ten respectively). As we see, a low number of neighbors may lead to a network structure which is not fully connected, i.e. its agents can only experience network externalities within their local cluster. The standardization processes in individual clusters can not diffuse to any agents of a different cluster. These "sub-populations" evolve in total separation and it is therefore rather unlikely, that all the isolated regions evolve to the same global ontology. With increasing connectivity (five or ten neighbors), the chances that a network is not connected become rather small, i.e. every sub-group of agents, agreeing on a specific standard, may "convince" their direct neighbor clusters to join them. The "domino effects" finally might reach every agents even in the most remote area of the network. However, the number of "dominos" that have to fall before a standard which emerged far away in a certain area of the network reaches the local environment of an actor and therefore influences the decision to adopt is typically much higher than in the corresponding graph with *random* topology. Speaking more formally, the average length of the shortest path connecting two arbitrarily chosen vertices of the graph (i.e. the number of neighbors you have to traverse) is smaller for the same connectivity if the graph has a random topology.

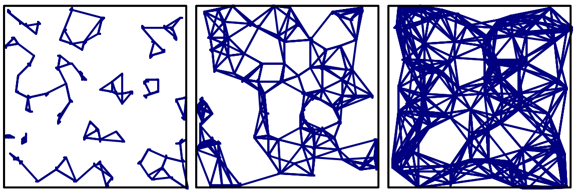

Figure 9. Typical networks with two, five or ten closest neighbors (close topology).

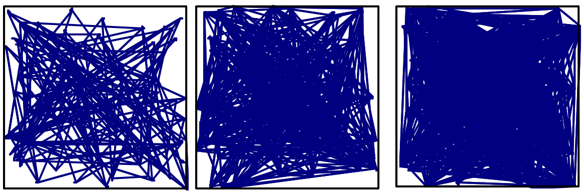

Figure 10. Typical networks with two, five or ten random neighbors (random topology).

Figure 10 shows the graphs with the same connectivity (2, 5, and 10) but *random* topology. The optical impression of a higher connectivity (which is an illusion) results from the fact that we selected "neighbors" to represent an asymmetric relation. That is, when agent x gets positive external effects by a neighbor y, it is unlikely in the *random* topology that vice versa, y also gets positive effects from x. Of course, within the *close* topology symmetric neighborhood is more common meaning that there is a higher probability that if y is the closest neighbor from the perspective of x, at the same time x is also the closest neighbor from the perspective of y. In this case the two links are plotted on top of each other and that is why the close topology graphs look less connected.

Of course most real-world networks represent an intermediate version of these extreme types, but since the costs of bridging geographic distance get less and less important the more information technology evolves, the tendency is clear. Electronic markets will rather resemble the *random* type of structure (since we select our partners by other criteria than geographical distance), while in markets for physical goods (or face to face communication) the physical proximity is still a very important factor for selecting business partners and therefore, the *close* topology will be a good proxy to the real world network structure.

A total number of 3,000 independent simulations were run with 1,000 agents and 10 different ontologies until an equilibrium was reached. The distribution reached in this equilibrium was then condensed into the Herfindahl[18] index used in industrial economics to measure market concentration (e.g. [Tirole 93]). In the following diagrams, every small circle represents one observation.

The left diagram in figure 11 illustrates the strong correlation (0.756) of connectivity and equilibrium concentration for *close topology*. Despite of this strong correlation it can clearly be seen that even in networks with 200 neighbors per agents (i.e. a connectivity of 200) the chances are still very low that one ontologie completely dominates the agnent society. For *random topologies* (figure 11 right) an even stronger correlation (0.781) is obtained. All the correlations illustrated are significant on the 0.01 level.

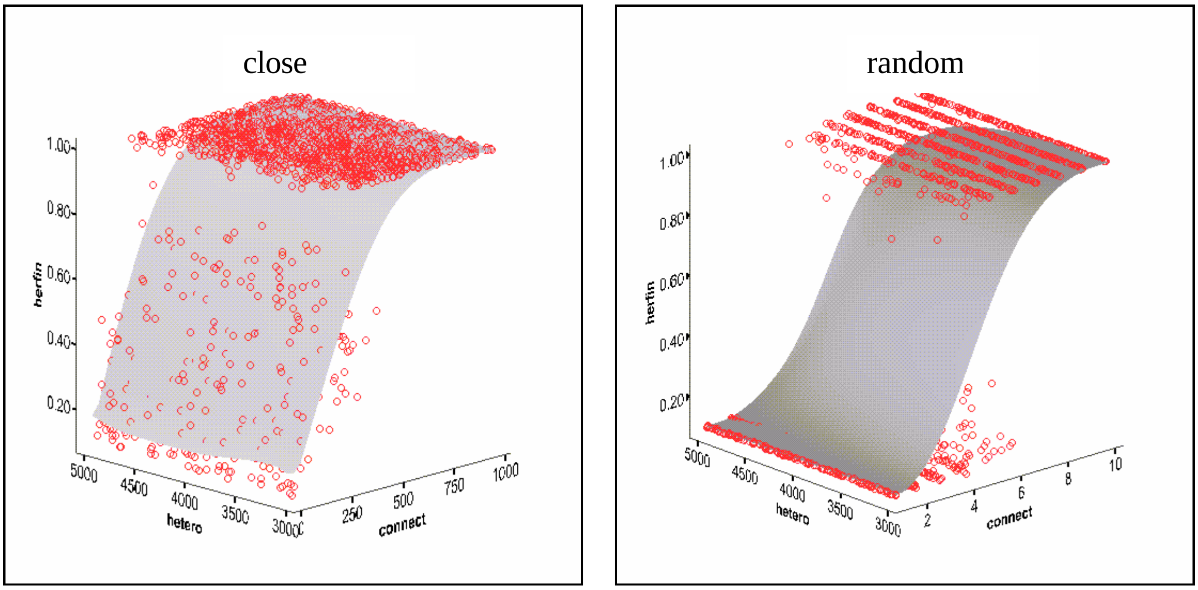

Figure 11. Equilibria in close topology and random topology networks.

Note that the scale of connectivity is extremely different in the two graphs of figure 4. It is obvious that the likelihood for total diffusion of only one ontology is very high in random topology network even for very low connectivity. The heterogenity of the individual benefit does not have a significant impact, as long as the cost of implementing the ontology is low[19] (since it gets dominated by the network effects), this situation changes, however, as soon as cost of implementation increases[20].

[18] The Herfindahl index is calculated by summing up the squared market share for each vendor. If all market shares are evenly distributed among our ten alternative products, we get the minimal concentration index of 10*(0.1)^2 = 0.1 while we get a maximal concentration index of 1*1^2+9*0^2 = 1 if the diffusion process converges to all consumers using one identical software product.

[19] The prices were set to 50 and the benefits from each neighbor sharing the own ontology was arbitrarily set to 10,000, reflecting a situation of very low implementation cost (probably being adequate for the adoption of "free" ontlogies by electronic agents).

[20] When the cost of implementation rises, the inertia to stick with the current ontology also increases, but the topological effect stays the same. This situation is probably more adequate for the adoption of laguages or new knowlede by human agents.

Comparing the two graphs *strongly* supports our hypothesis that for a given connectivity the indirect domino effects are much stronger for *random* topology networks and thus the diffusion process shows much higher tendencies towards standardization. To test this statistically, we ran a Kolmogorov-Smirnov test [Hartung 89: 520-524] rejecting the hypothesis that the concentration indices obtained for close and random topologies follow the same distribution on a significance level better than 0.0005 (KS-Z of 2.261).

5. Conclusions and further research

Although the advent of a semantic web is one of the most relevant areas of standardization, a diversity of knowledge representation languages and not yet interoperable tools prevents a seamless integration of ontologies. As we show by commenting on our own experiences on the way towards an integrated logistics ontology, numerous design decisions could lead to incompatibilities, a system of distributed ontology management should help to prevent or at least identify.

By conducting simulations we demonstrated that the structure (connectivity and topology) of the respective communities significantly influences the adoption processes of new ontologies. For both, random and close topology networks the concentration process correlated strongly with the agents connectivity. Since the globalization of research and business will shift the communication topologies towards the random topology, the pressure for standardization will increase, although topic centered communities may still coexist in isolation, refusing to "surrender" to the competing standard.

Up to now we assumed that only one of the competing ontologies is used at the same time. Although a common assumption in many network effect models, it is not too realistic for human agents, learning a new language and one major goal of the semantic web will certainly be to reduce the cost for software agents to communicate using multiple ontologies.

Speaking in economic terms, one of the major goal of the semantic web initiative is to overcome the disjunction of the incommensurable ontologies, allowing for an "incremental traversal" from one ontology to an other thereby drastically reducing the agents' switching costs.

6. REFERENCES

[Bass 69] Bass, F. M. (1969): A New Product Growth for Model Consumer Durables, in: Management Science, Vol. 15, No. 5 (January 1969), pp. 215-227.

[Buxmann 99] Buxmann, P./Weitzel, T./König, W. (1999): Auswirkung alternativer Koordinationsmechanismen auf die Auswahl von Kommunikationsstandards, in: Zeitschrift für Betriebswirtschaft, Ergänzungsheft 2/99, pp. 133-151.

[Chen 76] Chen, P.P-C. *The Entity-Relationship Model-Toward a Unified View of Data*; ACM TODS vol.1 no.1, pp. 9-36, March 1976.

[Church/Gandal 96] Church, J./Gandal, N. (1996): Strategic entry deterrence: Complementary products as installed base, in European Journal of Political Economy, Vol. 12 (1996), pp. 331-354.

[Connolly 01] Connolly, D.; van Harmelen, F.; Horrocks, I.; McGuinness, D.L.; Patel-Schneider, P.L.; Stein, L.A.: "DAML+OIL (March 2001) Reference Description". World Wide Web Committee (W3C) Note 18 December 2001. Also available from http://www.daml.org/2001/03/daml+oil-index.

[Couloris 01] Couloris, G.; Dollimore, J.; Kindberg, T.: Distributed Systems Concepts and Design (Third Edition) Addison-Wesley, 2001.

[Economides/ Himmelberg 95] Economides, N./Himmelberg, C. (1995): Critical Mass and Network Size with Application to the US FAX Market, Discussion Paper EC-95-11, Stern School of Business, New York University.

[Farrell/Saloner 86] Farrell, J./Saloner, G. (1986): Installed Base and Compatibility: Innovation, Product Preannouncements, and Predation, in: The American Economic Review, Vol. 76, Nr. 5 (December 1986), pp. 940-955.

[Fensel et al. 01] Fensel, D.; Horrocks, I.; van Harmelen, F.; McGuinness, D.L.; Patel-Schneider, P.F.: OIL: An Ontology Infrastructure for the Semantic Web; in IEEE Intelligent Systems, Vol. 16, No. 2, March/April 2001.

[Gandal 94] Gandal, N. (1994): Hedonic price indexes for spreadsheets and empirical test for network-externalities, in: Rand Journal of Economics, Vol. 25 (1994), No. 1, pp. 160-170.

[Genesereth 92] Genesereth, M. R.; Fikes, R. E.: Knowledge Interchange Format, Version 3.0 Reference Manual; Technical Report Logic-92-1, Computer Science Department, Stanford University.

[Gröhn 96] Gröhn, A. (1996): Netzeffekte in der Software-Industrie: Eine Analyse der empirischen Literatur, Ein Modell der Netzeffekte in der Software-Industrie, Kieler Arbeitspapier Nr. 743, The Kiel Institute of World Economics, Kiel.

[Gröhn 99] Gröhn, A. (1999): Netzeffekte und Wettbewerbspolitik. Eine ökonomische Analyse des Softwaremarktes, Kieler Studien 296, Tübingen.

[Gruber 93] Gruber, T.R.: A translation approach to portable ontologies. *Knowledge Acquisition*, 5(2):199-220, 1993.

[Hartmann/Teece 90] Hartmann, R. S./Teece, D. J. (1990): Product emulation strategies in the presence of reputation effects and network externalities: some evidence from the minicomputer industry, in: Economics of innovation and new technology, Vol. 1-2, pp. 157-182.

[Hartung 89] Hartung, J. (1989): Statistik: Lehr- und Handbuch der angewandten Statistik, München.

[Hendler 00] Hendler, J.; McGuinness, D.L.: The DARPA Agent Markup Language; IEEE Intelligent Systems, Vol. 15, No. 6, November/December 2000, pages 67-73.

[Horrocks 00] Horrocks, I.; Sattler, U; Tobies, S.: *Reasoning with Individuals for the Description Logic SHIQ*; in MacAllester, D. (ed): Proceedings of the 17th International Conference on Automated Deduction (CADE-17), Lecture Notes in Computer Science, Germany, 2000.

[Horrocks 02] Horrocks, I.: DAML+OIL: a Description Logic for the Semantic Web; Bulletin of the Technical Committee on Data Engineering 25 (2002) pp.4-9.

[Katz/Shapiro 85] Katz, M. L./Shapiro, C. (1985): Network externalities, competition, and compatibility, in: The American Economic Review, Vol. 75, No. 3 (June 1985), pp. 424-440.

[Lilien/Kotler 83] Lilien, G. L./Kotler, P. (1983): Marketing Decision Making. A Model Building Approach, New York.

[Mahajan et al. 90] Mahajan, V./ Muller, E./Bass, F. M. (1990): New Product Diffusion Models in Marketing: A Review and Directions for Research, in: Journal of Marketing, Vol. 54 (January 1990), pp. 1-26.

[Mahajan/Peterson 85] Mahajan, V./Peterson, A. P. (1985): Models for Innovation Diffusion, Sage Publications.

[Moch 95] Moch, D. (1995): Ein hedonischer Preisindex für PC-Datenbanksoftware: Eine empirische Untersuchung, in: Harhoff, D./Müller, M. (Hrsg.): Preismessung und technischer Fortschritt, Baden Baden.

[OWL 02] Requirements for a Web Ontology Language; W3C Working Draft 07 March 2002; http://www.w3.org/TR/2002/WD-webont-req-20020307/

[RDFS 02] RDF Vocabulary Description Language 1.0: RDF Schema; W3C Working Draft 30 April 2002; http://www.w3.org/TR/rdf-schema/

[Rohlsfs 74] Rohlfs, J. (1974): A theory of interdependent demand for a communications service, in Bell Journal of Economics, 5 (1974), pp. 16-37.

[Schoder 95] Schoder, D. (1995): Erfolg und Mißerfolg telematischer Innovationen, Wiesbaden.

[Sowa 00] Sowa, J.F.: *Knowledge Representation: Logical, Philosophical, and Computational Foundations*, Brooks Cole Publishing Co., Pacific Grove, CA, 2000.

[Tirole 93] Tirole, J (1993): The theory of industrial organization, 6th ed., Cambridge, Mass.

[Trunda 98] Trunda, T./Westarp, F. v. (1998): EDI solutions in large companies - A case study with 3Com, http://www.wiwi.uni-frankfurt.de/IWI/projectb3/eng/ publikat/cases/3Com.htm.

[Uschold 98] Uschold, M.; King, M.; Moralee, S.; Zorgios, Y.:The Enterprise Ontology. Knowledge Engineering Review 13(1) pp. 32-89.

[Valente 94] Valente, T. W. (1994): Network Models of the Diffusion of Innovations, Hampton Press, Cresskill, NJ.

[Weiber 93] Weiber, R. (1993): Chaos: Das Ende der klassischen Diffusionsforschung?, in: Marketing ZFP, H. 1, 1993, pp. 35-46.

[Westarp 00] Westarp, F. v./Weitzel, T./Buxmann, P./König, W. (2000): The Standardization Problem in Networks - A General Framework, in: Jakobs, K. (Ed.): Standards and Standardization: A Global Perspective, Idea Publishing Group.

[Westarp 99] Westarp, F. v./Weitzel, T./Buxmann, P./König, W. (1999): The Status Quo and the Future of EDI, in: Proceedings of the 7th European Conference on Information Systems (ECIS'99), Copenhagen 1999, pp. 719-731.