Robert Meersman and Christophe Debruyne. Hybrid Ontologies and Social Semantics. In DEST 2010, pages 92-97. IEEE Press, 2010

Hybrid Ontologies and Social Semantics

Keynote talk, IEEE DEST 2010 Conference, Dubai, 12-15 April, 2010.

Robert Meersman and Christophe Debruyne
VUBrussel STARLab
Pleinlaan 2
B-1050 Brussels
Belgium

Abstract - **Semantic Web, Social Web, and new economic challenges are causing major shifts in the pervasive fabric that the internet has become, in particular for the business world. The internet's new role as participatory medium and its ubiquity lead to dense tri-sortal communities of humans and businesses mixed with computer systems, and semantically interoperating in a well-defined sense. Many of the challenges and ongoing (r)evolutions appear to produce as yet seemingly contradictory requirements and thus produce potentially very interesting research areas. We argue that linguistics, community-based real world "social" semantics and pragmatics, scalability, the tri-sortal nature of the communities involved, the balance between usability and reusability, and the methodological requirements for non-disruptive adoption by enterprises of the new technologies provide vectors for fundamental computer science research, for interesting new artefacts, and for new valorisations of enterprise interoperability. We posit that one such development will likely result in hybrid ontologies and their supporting social implementation environments –such as semantic wikis– that accommodate the duality and co-existence of formal reasoning requirements inside systems on the one hand and of declarative knowledge manipulation underlying human communication and agreement on the other hand.**

I. Introduction

To state that the Web is turning pervasive and ubiquitous has become something of a platitude. As its density increases, perhaps it will soon be more appropriate to think of it as a tissue rather than a web, with warp and weft provided by numerous –meaningful– services connecting its three main sorts of nodes, computer systems, humans,[1] and businesses. Note that in stating it so we take the physical connectivity for granted; it is the existence of (computer-based) services between nodes that will determine the "real", or if one wants, semantical topology of the "future" Internet of Everything. For this reason we will use the terms Internet and Web below as interchangeable except when referring to the Internet as merely the infrastructure. Note that by having enterprises as first class citizens one becomes motivated to make explicit concepts of group, organisation, community, business ecosystem etc. as modelling primitives leading to richer models than bi-sortal networks where businesses are represented by their human and system "components".

The Web presently is also very rapidly migrating from an information retrieval medium into a participatory one, and in combination with pervasiveness this causes three parallel and interconnected evolutions to take place simultaneously, albeit not synchronized nor occurring at the same pace. These evolutions are respectively driven by

- *technology* (exemplified by the semantic Web);
- *social forces* (newly manifested as the Social Web, or Web 2.0);
- *economical forces*, the internet now being the medium of choice for most content-based interaction of businesses with vendors, customers, and other enterprises. Enterprises as such, or organised groups and communities in general, indeed provide the reason and context for the technology and the application of these social forces.

Not coincidentally, the principal agents in these evolutions are again respectively computer systems, humans and enterprises. In fact, we could envisage for the purpose of our arguments that there are three "parallel" networks in play, each one connecting the nodes of one sort. Interesting research questions emerge for instance at the places where these networks –necessarily– must interact. To see this, we first turn to the semantic Web.

II. Interoperability and the Semantic Web

The formal semantics of a (computer-based) system quite simply is the correspondence between this system and some real world as perceived by humans. It is usually given by a formal mapping of the symbols in the system's description to objects in that real world, such that relationships and logical statements in the specification language can be assigned a truth value depending on whether a certain state of affairs among objects exists in the real world. As the real world is not accessible inside a computer, if we want to store and reason about semantics the world needs to be replaced by an agreed *conceptualization*, often in the shape of a formal (mathematical) construct. A computer-based, shared, agreed formal conceptualisation is known as *an ontology*. Ontologies constitute the key resources for realizing a semantic Web. While theoretically ontologies should be perfect renderings of a real world, in practice they evolve as successive (one assumes, ever better) approximations of it [G95].

The fundamental characterizing principle of the Web from the viewpoint of semantics is the *autonomy* of its nodes. Yet in

[1] this sequence does not imply an order of importance, of course. But note that as *active* agents *within* the WWW, humans *are* in fact relative late-comers…

spite of this in an increasingly pervasive environment web-based systems at these nodes will wish or need to make use of each other's data or processes. Under conditions of respect for autonomy such collaborative arrangement among systems is called *interoperation* and usually implies that some mapping has to occur at runtime, either interpreted or compiled, from the autonomous individual data models and process specifications to some ontology, which in this manner acts as a repository for the semantics of the domain in question. Each mapping remains the responsibility of the system's owner. When on the other hand one is in a position to modify, merge or federate the data and/or process models into a common model, we instead prefer to speak of *integration*.

The notion of interoperability may be transferred to enterprises as is e.g. implicitly done in [EIVP-08]. Businesses inherently are autonomous, making the ability to interoperate critically dependent on shared formal resources such as ontologies, standards, and environments conducive to agreements and which then may be implemented "around" the enterprise's systems to have these interoperate. Assuming asymptotically total internet pervasiveness it is therefore realistic to say that two enterprises are interoperable if and only if their web-based information systems are.

Evidently, system interoperability implies reconciliation or at least balancing of two opposing concerns, the "right" of systems to (technical, business) autonomy and the (social, economical) requirement to collaborate. This in turn leads to a well known trade-off between *usability* (a system's suitability for a particular function or purpose) and *reusability* (the capacity of a solution to be applied in another context). Reusability leads to *leverage* [Z08] which is a fundamental feature of the concept of "generativity", a term coined by that author [ibid.] for the capacity of a technology to produce ("unanticipated") change through ("unfiltered") contributions from a wide audience, and therefore as such induces and promotes innovation. Leverage constitutes a measure for the "general purposeness" of a system and is achieved easier in systems and technology with fewer constraints; examples are IP, the Windows PC, and HTML. It should be obvious that in order for ontologies (touted by some as the "silver bullets" of the semantic Web) to be generative, their specification formalism should maximize leverage. One of the current obstacles to widespread adoption of "real", i.e. formal-reasoning-based semantic Web technology appears to be that a language such as OWL (and even RDF(S)), while elegant, well-founded by theory and robust, easily leads to tightly constrained ontologies that lack reusability. Unfortunately this is not (yet) compensated by an increase in usability; there is claimed to be a proliferation of OWL ontologies (some "tens of thousands" as claimed by [HSHBW08]) but most of these are actually just single-application data models and finding one "usable" for a particular purpose may be more expensive than building one's own —leading to an increase of neither (much) adaptability nor reusability.

This combination may so far have prevented both the leverage and the adaptability needed in order to achieve critical mass, i.e. large scale adoption of semantic Web technology (the fact that SW technology was often presented as disruptive did not help either, esp. in the absence of a proper methodology that involved respect for legacy systems and existing workflows, etc.). On the other hand, RDF(S) (a much simpler formalism and underlying languages such as OWL) seems to be better placed to achieve such generativity status, having been adopted by major database vendors. It is also the mainstay of the Linked Open Data initiative [LOD], which aims to link vast numbers of data elements inside existing resources to others using a trivially simple, and therefore very scalable and standardized URI mechanism. It is conjectured that this initiative may contribute significantly to unlocking large volumes of legacy data, key resources on the "business Web", turning them into semantic interoperability resources. Interestingly, the annotation of existing company data models and database content turns out in many cases to be a "social" activity involving business-level agreement about the contexts for the data to be linked.

III. From the Semantic to the Social Web

There has not been a direct technological or research transition path from the Semantic Web to the emerging Social Web (or Web 2.0 as it is sometimes called). In fact one could claim that this evolution took everybody by surprise –there certainly was little anticipation of it in the SW research community except perhaps for a visionary paragraph in [B99] stating[2], "*Computers can help if we use them to create abstract social machines on the Web: processes in which the people do the creative work and the machine does the administration… The stage is set for an evolutionary growth of new social engines…*". Part of the explanation for this surprise could be that in such social networks there is in fact today no conceptual or technological dependency on semantics –yet.

The Social Web is characterized by large communities of human agents who participate and contribute through the on-line network to a social activity, often just the establishment itself of that community around some more or less agreed themes. Current successes such as Facebook, Flickr, Twitter, MySpace etc. have a distinct one-of-a-kind quality (or at most very few of a kind) and often admit to the absence of an a priori business model. Implementations of Web 2.0 systems do not (again, as yet) require real technological innovation beyond scalability, nor is it clear which shape such innovation should take. Semantically speaking, networks such as Facebook, LinkedIn, Plaxo and others "own" their respective domains in the sense that they are largely in control of the interface semantics (typically plug-ins) of what appears on their networks, and so need only agreements on the technical rather than at the domain level. This situation is due to change as such systems will become more open and involve domain knowledge, i.e. accommodate –for a fee, likely– autonomous plug-ins that exploit the substantial socio-economic resources being built up in such networks. On the other hand however, as widespread use of tagging already indicates, it is not difficult to envisage already now how Web 2.0 communities could *contribute* to the creation and maintenance of semantics resources such as ontologies. A first approach likely will be the exploitation of the tags, and their standardization, or

[2] The authors are grateful to Jim Hendler for pointing this out

agreement on their intended use. As agreement is primordial for systems and enterprise interoperability, on-line communities will indeed increasingly become the environment where these semantics agreements are engendered and negotiated, and we shall call *social semantics* the result of such process of determining meaning by community agreement.

Indeed, it turns out that even though most semantics resources are shared, these often have been developed earlier and autonomously, within a "closed" environment such as an enterprise, committee, research group, etc.. Web 2.0 technology, as has been reported in contexts of Open Source, Wikipedia, Flickr and others[3], may in fact allow to exploit a kind of "inverse" Pareto Principle[4] where the 80% "long tail" of the community, until now responsible for "only" 20% of contributions to such resources, possess the means to become significant.

All this in turn leads to a focal role for natural language (NL) based technology –not so much the usual parsers or even NL understanding systems, but NL in its core function as the carrier of choice for the communication and mediation processes among humans that lead to the necessary agreements. As is borne out by database design methods and techniques such as NIAM and ORM [HM08], the closer the link between this human NL communication and the system and/or business communication that results from it, the more likely such systems will work as intended by their various stakeholders. This is particularly relevant for the interfaces at places where the human, systems and business webs interact, as these constitute the fulcrums where human communication needs to be meaningfully mapped onto systems- and/or enterprise interoperation in order to achieve the leverage so avidly sought. In particular it will be necessary to bridge the "impedance gap" that results from the different level at which humans communicate and agree on concepts, and the level at which knowledge representation is deemed practical for use by system-based, formal reasoners. We argue that this must lead to *hybrid* ontologies[5], where concepts on the one hand are circumscribed linguistically and (mostly) declaratively by agreement within (human) communities, and on the other hand identified formally (and unambiguously) for use in computer-based information systems. Both conceptual schemas will coexist as separate entities; a (likely negotiated) mapping from the linguistic circumscriptions to system identifiers is required.

Certain NL-based techniques also allow scalable solutions to ontology engineering through a *separation of concerns* similar to the one exemplified by databases, which rigorously separate the schema level from the instance level. This separation in turn results in *data independence*, the defining property of databases, allowing application programs to become maximally insensitive to changes in data representation. By analogy, as is implemented in the DOGMA[6] paradigm [M99], one can decompose ontologies into (large) sets of simple, "plausible" and linguistically formulated facts –called *lexons*– and a layer of *ontological commitments* that are used to select lexons, annotate applications, and specify constraints defining the use of the concepts in the ontology. A trivial example of a lexon would be

< γ, Dot, with_color, color_of, Red >

where γ is a *context identifier* pointing to a resource (e.g. a text about laser pointers) assumed to identify unambiguously – to human users– concepts of dot, red, and the roles they play with respect to one another in this context. A commitment (for a specific chosen application linked to this context) could e.g. state that any dot can have at most one color.

Domain-level agreements on lexons are supposed to be easy (they need only be plausible) and may be arrived at –in a scalable fashion– by Web 2.0 methods; lexon bases may come out rife with apparent redundancy, but are intended to be very reusable and so provide semantic leverage. Usability and its associated complexity are delegated to commitment specification and remain the sole responsibility of the application, making the overall approach generative in the sense discussed in Section 2. The NL aspect of lexons specifically leads to semantics in a natural way, as any human language has evolved over millennia by natural selection on the basis of its link with the real world. Note that in this formalism hybrid ontologies could (somewhat naively) be implemented by linguistic context resources on the human side and [e.g. an RDF representation of] lexons on the systems side, and the unique concept disambiguation condition mentioned above as the mapping between them. DOGMA assumes some additional axioms [DMM07] but we omit these here as they are less relevant for our arguments. The DOGMA Studio tools for (non-hybrid) ontology engineering are commercially available from the 2009 Collibra spinoff initiative of VUB STARLab [Collibra]. Collibra also developed the Business Semantics Glossary (BSG), which is a workspace based on wiki technology where both business and technical users collaboratively define and govern the meaning of the business assets in their business context.

Wiki technology has been put forward as a mean to reach agreement and share knowledge about different subjects over the past decade. This technology powers the Wikipedia [Wikipedia] online encyclopaedia containing over 3 millions articles in English. The advantage of wiki technology is that anyone can add content without much technical knowledge. Wiki technology has been adapted in the field of ontology engineering to enable non-technical users to create, visualize and maintain ontologies.

Because of its power, semantics-enhanced wiki technology has been used and explored in a number of projects, such as: the Semantic MediaWiki [KVVHS07] extension to Media-Wiki to add information about the data within pages, and the relationships between pages. OntoWiki [ADR06] facilitates the visual presentation of a knowledge base and involves the

[3] TED talk by Clay Shirky "Institutions vs. Collaboration" July 2005, http://www.ted.com/talks/clay_shirky_on_institutions_versus_collaboration.html

[4] A.k.a. the 80-20 rule stating that for many events, roughly 80% of the effects come from 20% of the causes, http://en.wikipedia.org/wiki/Pareto_principle

[5] concept and terminology coined in recent collaborative work with Tharam Dillon [forthcoming]

[6] Developing Ontology-Grounded Methods and Applications

community; however users are required to have some formal background in knowledge engineering due to OWL's groundings in description logics. The DBPedia [ABKLCI08] initiative extracts structured data from Wikipedia to store and publish it on the Web. The dataset can then be used to perform queries in SPARQL, but currently needs a lot of cleansing. Somewhat related is the work done on WikiDB [WikiDB] that enables the definition and querying of data within a page, much like a database.

GOSPL [DRM10] (Grounding Ontologies with Social Processes and Natural Language) is one example of recent initiatives aiming to enable communities to develop and maintain a representation of their (business) world. This involvement is considered essential for facilitating the uptake of LOD, which annotates and exposes datasets on the World Wide Web. Thus the community as a whole needs to reach an agreement on the meaning of such annotations of legacy data. The LOD initiative however relies on RDF(S) and URI mechanisms to represent these annotations, which are not laymen friendly. The GOSPL prototype for instance is based on DOGMA using Social Web technologies (MediaWiki, DokuWiki) to allow members of a community to express knowledge partly in their own language (See Fig. 1), and to support provenance by tracking changes to the shared agreement on this knowledge. The prototype is currently used within the TAS[3] project[7] to allow end users (e.g., security and privacy experts) to easily develop conceptual models on the domain of security and privacy to provide security policy interoperability within the project.

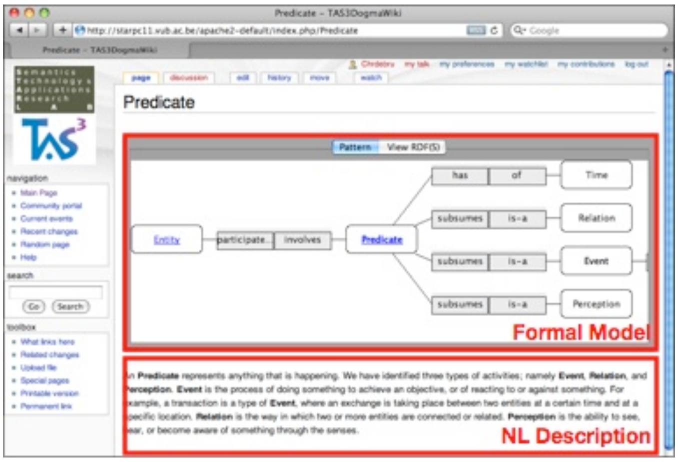

Fig. 1: Screen shot of the GOSPL wiki showing the distinction between the formal and informal definition of a concept.

It is interesting to compare GOSPL with other technologies such as BSG. Both are used for Knowledge Management with, among others, governance through fine grained user roles and permissions and versioning of both the formal and the informal definitions. In GOSPL agreements are reached based on natural language and these shared agreements are then transformed into RDF(S) to annotate datasets, enabling a community to (re-)publish their data on the Web. The aim is to involve every possible stakeholder, relying on the wisdom of the crowds, in reaching a shared understanding. For instance BSG rather focuses on enterprise interoperability within a more controlled group of stakeholders. GOSPL also wants to provide a prototypical hybrid aspect by implementing links between the informal communication in a community and a formal definition and examples for the terms and lexons, individually or as a whole, as they denote something meaningful in the real world, and envisions the publication of data through application commitments written in Ω-RIDL. Ω-RIDL allows mapping application symbols (in XML, Relational Database, etc.) to concepts and lexons in the ontology.

The DOGMA MESS methodology and tools illustrate some of these "social" processes in ontology engineering and deployment in business organisations and communities [DMM06]. To this end De Leenheer in his PhD thesis [D09] cleverly adapts the well-known SECI knowledge externalization/internalization method of Nonaka and Takeuchi [NT95], which in principle was intended for humans-only deployment, to the (bi-sortal) community of humans mixed with information systems. This allows communication and agreement processes between humans to be mapped to detailed system interoperability requirements, and vice versa. The resulting ontologies however are not explicitly hybrid; their concepts are all (uniquely) internalized inside systems albeit originally playing roles in lexons that are assumed to have emerged from NL communication, interview, agreement, etc..

Tools such as GOSPL aim to make the technical and non-technical interface to knowledge co-exist by having communities perform and manage the mappings between the linguistic circumscriptions and system identifiers. The community will play thus a more prominent role in what the knowledge means (semantics) and how it is going to be used (pragmatics), promoting the community as a *first-class citizen* in the ever ongoing mutual evolution of semantics and pragmatics of knowledge.

It will be quite interesting to see if and how the mechanisms of a Social Web, and the processes that lead to its applications, might assist humans to deal with the issues above. Early work e.g. on identity management in the context of privacy and security of transactions, exploiting (as yet simple) semantics of FOAF [S08] may be symptomatic of some of the ground-breaking research and development that is needed.

IV. The Business Web

Apart from the above issues of semantics and natural language, for enterprises (this includes governments and other organised goal-oriented groupings) the adoption of semantic technologies implies a methodology, considerations of legacy systems, scalability of solutions, sustainability of the technology and its cost effectiveness, and the presence of a so-

---

[7] The research leading to these results was partially funded by the EC's FP7 program under grant agreement number 216287 (TAS[3] – Trusted Architecture for Securely Shared Services).

called Virtuous Circle [8] connecting technology vendor(s), enterprise architect, and customer base in a purposeful and productive dialog.

Presently the business Web is the least developed of the three webs, but over time it stands to produce the greatest value from the Semantic-Social-Business combination. One of the main reasons is that in terms of numbers SMEs[9] would constitute the largest and densest on-line communities; however an SME typically is concerned with locally optimizing its place in the value chain (its immediate customers and its suppliers, mostly) and therefore enjoys as yet limited benefit from the high connectivity potential of a business Web, unlike e.g. banks and governments. As "semantic standardization" sets in this situation will change with the correspondingly decreasing risk, and cost, of entry. Social mechanisms may be expected to come into play here in two ways: first by creating business ecosystems sharing common concepts, relationships, events, generic goals and workflows, and secondly to support the very process of standardization itself (note that a standard or reference model may be seen as an object similar to an ontology). The first significantly large networks of SMEs are already emerging along these principles [H08].

The business Web, social Web and semantic Web indeed are bound to have legion points of interaction, likely nearly all of a semantic nature for the reasons discussed in the preceding sections. For instance, existing databases have enormous semantic leverage potential, since annotating a data model in principle provides one with a "semantics" for each tuple described by it; as populations typically are large or very large compared to the data schema, the gain can be substantial. More likely than not, such annotations and their ontologies will have to be re-extracted from the data models by a social process of agreement, and involving players or resources outside of the enterprise. Note incidentally that in spite of many articles in the SW literature failing to make the distinction, data models are *not* ontologies, since they in general lack the element of semantic agreement beyond the boundary of the organisation that developed them and owns them. For more arguments on this issue, see [M01]. As already mentioned in Section 2, the Linked Open Data initiative [LOD] aims to unlock the content of databases as well as other resources in a fashion that prepares it for future semantic exploitation (reasoning). Whether this will happen remains to be seen, as LOD is based on a pointer mechanism (URIs) to connect data as "concepts". Pointers lead to implementation-specific "semantics" as they must be dereferenced, and in turn lead to complex update processes if consistency is to be maintained in a distributed system, a difficulty well known since the bygone days of e.g. CODASYL [O78]… We note in passing that this disadvantage is not present in lexons precisely because of their linguistic NL origin, thereby relating to pointers in much the same way that the Relational Model relates to earlier pointer-based database models.

Finally, for a business Web to reach critical mass and sustainability new methodologies must be developed –the cost of introducing a new technology is high and the process must be made teachable and repeatable. Critical properties of such methodologies are to make the process of technology adoption sustainable (i.e. non-disruptive) and scalable: see the discussion on OWL in Section 2; maybe this should be compared to the –unsuccessful– attempt in the early 90s to introduce the otherwise very elegant Datalog technology.

Authors' note. Earlier versions of sections I, II and parts of section IV were published by the first author in [M10].

V. References and Bibliography

[ABKLCI08] Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak R., Ives, Z., "Dbpedia: A nucleus for a web of open data", in: Proceedings ISWC+ASWC 2007: 722–735 (2008)

[ADR06] Auer, S., Dietzold, S., Riechert, T., "Ontowiki - a tool for social, semantic collaboration", in: Proceedings ISWC 2006, LNCS 4273: 736–749 (2006)

[B99] Berners-Lee, T., "Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web by Its Inventor". Harper San Francisco, 1999.

[Collibra] http://www.collibra.com

[DRM10] Debruyne, C., Reul, Q., and Meersman, R., "GOSPL: Grounding Ontologies with Social Processes and Natural Language". Proceedings ITNG 2010 (forthcoming)

[D09] De Leenheer, P., "On Community-based Ontology Evolution", PhD thesis, VUB STARLab, 2009.

[DMM06] De Leenheer, P., de Moor, A., Meersman, R., "DOGMA-MESS: A Meaning Evolution Support System for Inter-organisational Ontology Engineering", Proceedings ICCS 2006: 189-202 (2006).

[DMM07] De Leenheer, P., de Moor, A., Meersman, R., "Context Dependency Management in Ontology Engineering: a Formal Approach", Journal on Data Semantics VIII, LNCS 4380, Springer-Verlag: 26-56 (2007).

[EIVP-08] Man-Sze Li et al., (eds.), "Unleashing the Potential of the European Knowledge Economy –Value Proposition for Enterprise Interoperability", European Commission, 2008.

[EIRR-08] Charalaibidis, Y. *et al*. (eds.), "Enterprise Interoperability – Research Roadmap", European Commission, 2008.

[G95] Guarino, N., "Formal ontology, conceptual analysis and knowledge representation", Int. J. Hum.-Comput. Stud. 43(5-6): 625-640 (1995).

[HM08] Halpin, T., Morgan, T.: "Information Modeling and Relational Databases", Morgan-Kaufman, 2008.

[H08] Hepp, M., "GoodRelations: An Ontology for Describing Products and Services Offers on the Web", Proceedings EKAW 2008: 329-346 (2008)

[HSHBW08] Hendler, J., Shadbolt, N., Hall, W., Berners-Lee, T., Weitzger, D., "Web science: an interdisciplinary approach to understanding the web", CACM 51(7), 2008.

[KVVHS07] Krötzsch, M., Vrandecic, D., Völkel, M., Haller, H., Studer. R., "Semantic wikipedia". Journal on Web Semantics, 5(4):251–261, 2007.

[LOD] http://linkeddata.org/

[M99] Meersman, R.: "Semantic Ontology Tools in IS Design", Proceedings ISMIS 1999: 30-45 (1999).

[M01] Meersman, R.: "Ontologies and Databases: More than a Fleeting Resemblance", In: Proceedings of OES/SEO Workshop Rome, LUISS Publications, 2001.

[M10] Meersman, R.: "Hybrid Ontologies in the Tri-Sortal Internet of Humans, Systems and Enterprises", in: Proceedings of InterOntology'10, Keio University Press, Japan, 2010.

[8] As opposed to a Vicious Circle. In a Virtuous Circle positive effects are reinforced through dialogue among the indicated stakeholders. Virtuous Circles in adjoining phases of the technology adoption lifecycle interact with each other since adopters (problem owners) in one phase often are the solution providers in the next. [Ongoing work by M. Brodie, R. Meersman et al. to be reported elsewhere]

[9] Small-to-Medium Enterprises

[NT95] Nonaka, I., Takeuchi, H., "The Knowledge-creating Company: How Japanese Companies Create the Dynamics of Innovation", Oxford University Press (1995).
[O78] Olle, T.W., "The Codasyl Approach to Data Base Management", Wiley (1978).
[S08] Story, H., "Building Secure Open & Distributed Social Networks", http://www.slideshare.net/bblfish/building-secure-open-distributed-social-networks-presentation, 2008.
[SMJ02] Spyns, P., Meersman, R., Jarrar, M., "Data Modelling versus Ontology Engineering", SIGMOD Record 31(4): 12-17 (2002).
[WikiDB] http://www.mediawiki.org/wiki/Extension:WikiDB
[WikiPedia] http://www.wikipedia.org/
[Z08] Zittrain, J., "The Future of the Internet and How to Stop it", Yale University Press, 2008.