Agreeing While Disagreeing, a Best Practice for Business Ontology Development

Emanuele Della Valle, Irene Celino, and Dario Cerizza

CEFRIEL – Politecnico of Milano, Via Fucini 2, 20133 Milano, Italy
{emanuele.dellavalle,irene.celino,dario.cerizza}@cefriel.it

Abstract. The agreement is a crucial part of our living together. Important opportunities for sharing resources, integrating systems and collaborating depend on our ability to agree. While we are interested in methods and technologies that support shared agreement, we somehow tend to forget the disagreement; indeed, it is also a part of reality.
In the Semantic Web field, most of research activities investigate ontology agreement and its formalization, rather than the disagreement and the best practices about its extended use.
We claim that the industrial uptake of the Semantic Web is severely limited by the fact that, in the real world, shared agreements are difficult to reach and maintain and that "agreeing while disagreeing" is a common practice. In this paper we illustrate a best practice, which we learned from practical experience, that makes use of an (almost) unexplored potentiality of the Semantic Web to express disagreement, and we describe its use in the e-employment sector.

1 Introduction

Whenever we agree with someone, we feel free to express our consent (*"I think you're right"* or *"I agree with you"*), because the agreement does not imply any conflict. On the other hand, whenever we have to dissent from someone, we feel embarrassed to express our disagreement (*"I'm afraid I can't agree with you"*) and we often try to find the most polite expression in order not to offend our interlocutor (*"I agree with you up to a point, but..."*, *"That's quite true, but..."* or *"I agree with you in principle, but..."*, etc.). As a consequence, every spoken language provides for a lot of expressions to formulate (or to hide) dissent and a fewer expressions, that we daily use without any worry, to state accordance and harmony.

It seems that we fear both the disagreement and the conflicts that may rise by showing it to the others. After all, we can happily work together by committing to a limited agreement, while hiding all the necessary (and unnecessary!) discordances. Such behavior is theoretically founded in the *common knowledge* theory [1], which demonstrates that, when showing the whole agreement, parties cannot "agree to disagree". Therefore, the only way to disagree is to show only part of the agreement. It may appear counter-intuitive, but a good and common practice is to "agree while disagreeing".

What happens when you have to develop a Semantic Web application that deals with such a common behavior? Since an accepted definition of ontology is a "formal, explicit specification of a *shared* conceptualization" [2], we can clearly assume that an ontology is a kind of formal and explicit agreement among the parties that commit to it.

Is there room for disagreement? Being able to share the disagreement may not be a desirable feature, because we tend to agree on many general principles, to commit to a clear subset of the possible implications of this principles and, in the meantime, to keep our disagreement. After all, however, disagreement exists.

Our thesis is that the Semantic Web is elaborating all necessary tools to formalize the agreement and the disagreement, and that their appropriate use can result in a method that helps in developing real-world Semantic Web applications.

The rest of the paper is structured as follows: section 2 states the problem of "agreeing while disagreeing" and presents the conceptual solution we elaborated in several research projects; the notion of ontological commitment is introduced in section 3. Section 4 draws a boundary line between agreement and commitment. Section 5 explains in details our best practice for expressing disagreement and making it partially public; section 6 outlines the method we put in practice in elaborating the SEEMP solution and its advantages; finally section 7 presents conclusions and future work.

2 The Problem and a Conceptual Solution

We have been technically leading several European projects (i.e., COCOON[1] and SEEMP[2]) and Use Case work packages of European and Italian projects (i.e., in TripCom[3] and NeP4B[4]) that aim at fostering the adoption of Semantic Web in health care, government and business. In all those projects, we face the problem of integrating, using a Semantic Web approach, data and applications that belong to parties that are definitely interested in collaborating, but are also competitors. Therefore, they have to agree about the sharing of all the information that will foster collaboration, but they also should not be forced to disclose information that will reduce their competitive advantages.

In those cases, parties agree on general principles, but then they only commit to a subset of the possible implications of these principles and keep for themselves all their disagreements.

For instance, in the domain of e-employment (scenario of the SEEMP project), different Employment Services, both public and private, collect CVs and job offers in order to match demand with supply. Each Employment Service (ES) covers either a region or an occupational sector. As a result, the employment market is severely fragmented and many ESs perceive the need of sharing information in order to provide a better service to their customers. However, they

[1] `http://www.cocoon-health.com/`
[2] `http://www.seemp.org/`
[3] `http://www.tripcom.org/`
[4] `http://www.dbgroup.unimo.it/nep4b/`

would never exchange CVs or job offers, since they contain sensible information (like contact details); instead, the ESs use the "anonymized" versions of CVs and job offers, which we name respectively candidacies and vacancies. Therefore, if an ES exchanges a candidacy/vacancy with another one, it potentially enlarges the possibilities of finding a match, without giving to the other one the chance to by-pass it and directly contact the employee/employer.

In a similar manner, in the logistic domain (scenario of the NeP4B project), different logistic operators may join their efforts to acquire a larger market share by presenting themselves under a common brand. They make available information on their number and type of vehicles, their position on the territory, etc., but, being competitors, they only partially disclose their respective price lists. A joint logistic platform should have enough information to calculate the price of a shipment request, but not enough to let one of the logistic operator to calculate a competitive counter-offer.

Therefore, we draw the conclusion that the common notion of *shared agreement* is not enough to manage the complexity of the industrial scenarios we face. We believe that two important notions have to be made explicit: *commitment* and *disagreement*. It is worth noting that usually when parties provide ontology commitment the intended meaning is that all parties commit to the "entire" ontology. On the contrary, we propose to give a "subjective" meaning to commitment and disagreement which does not presume a common knowledge among all the parties: two parties may commit to (or disagree with) different part of the agreement.

In order to move from the problem statement to the solution conception, we need to find appropriate methods and technologies. We need conceptual elements to capture the notions of agreement, commitment and disagreement, to make them operational and to express the respective relations among them.

Ontologies have been used and are good for formalizing and sharing the agreement. The notion of commitment is usually associated to the notion of ontology, and this is certainly true in the context of agent communication (see also section 3). In agent-based systems, all agents usually share a single ontology. The Semantic Web vision, however, foresees an ecosystem of ontologies, because of the very nature of the Web which is "fractal" [3]. Ontologies can be co-invented, they can partially overlap and, in developing a new ontology, the importing of existing ones is encouraged [4]. We believe that the "practical" meaning of ontological commitment in the Semantic Web is slightly different from the original one (see also section 3). In formal terms, committing to an ontology that imports several other ones is the same as committing to one big ontology obtained by the union of all of them; however, in practical terms, committing to the ontology that includes the import annotations is partially an "unconscious" commitment, in that it means trusting the knowledge engineer who decided which ontologies to import.

Therefore, our best practice is to distinguish between the *reference ontology*, which captures the shared agreement, and the *local ontologies*, which captures the commitment and the disagreement of the various parties. We propose to

build the reference ontology including all the details that are needed to allow for a meaningful communication between each pair of parties, thus including details that most of the parties would consider either inessential or not sharable. Then, each party can develop its local ontology, partially by importing parts of the reference ontology, and partially by developing its own point of view. Every time a local ontology imports a part of the reference ontology[5], the party is considered to commit to the imported parts of the reference ontology. Moreover, every time a part of the local ontology is aligned to a part of the reference ontology (e.g. by the means of ontology-to-ontology mediators [5]), the party is also said to commit to that part of the reference ontology. A particular attention should be given in capturing also the source of disagreement within the local ontology. Finally, each party should make available to all other actors the part of the local ontology that explains its different point of view without causing conflicts.

Sections 3-5 are devoted to explain the details of our conceptual solution.

3 Ontological Commitment

In the philosophy of language and metaphysics, the *ontological commitment* of a proposition is the definition of some of the demands that the sentence's truth imposes on the world, e.g., the existence of one thing presupposed or implied by asserting the existence of another one.

In the context of artificial intelligence, the ontological commitment is a direct consequence of the ontology definition: an ontology is a specification of the concepts and relationships that can exist for a community of agents and the ontological commitment is the agreement to use the shared vocabulary defined by the ontology in a coherent and consistent manner [6–8].

From the previous definition, it follows that an ontological commitment is a guarantee of *consistency*, but not *completeness* with respect to queries and assertions [9]: the agents sharing a vocabulary do not need to share a knowledge base; each of them knows things the others do not, and an agent that commits to an ontology is not required to be able to answer all queries that can be formulated in the shared vocabulary.

Also in the Semantic Web community, we often refer to the ontological commitment in the sense explained above. We are used to build ontologies to express the objects, concepts, and other entities that are assumed to exist in some area of interest and the relationships that hold among them; the resulting ontology represents an abstract, simplified view of the world that we wish to formalize for some purpose. As a consequence, every knowledge-based system built on top of an ontology is committed to its conceptualization. This means that, when two different systems commit to the same ontology, they can assert the existence of different things and individuals (*in-completeness* of the commitment), but

[5] Please note that the behavior of this import feature differs from the behavior of `owl:import` annotation, because it allows for a *partial* importing. However it can be mapped to `owl:import` is the reference ontology is correctly modularized.

they share the same vision of the conceptual model expressed by the ontology (*consistency* of the commitment).

Several different methodologies developed to build ontologies [10] are based on this notion of ontological commitment, in that they assume the ontology engineering process to be aimed at formalizing the obligation to a specific conceptualization of the world (or of a part of it). As we mentioned in section 2, the only difference on the notion of ontological commitment between the Artificial Intelligence and the Semantic Web world seems to be that, in the latter case, the engagement to the ontology is explicit and conscious with regards to the concepts, relations, axioms and individuals defined within the ontology itself, while, in the Semantic Web, it is somehow unconscious with respect to the ontologies imported via the `owl:import` mechanism.

4 Agreement vs. Commitment

In proposing our best practice, we would like to stress the difference between the achievement of an agreement and the commitment to the obligations it entails.

To better explain this distinction, we take as an example the situation in which there are three different actors wishing to pave the way for a fruitful collaboration between them. To start their interaction, they need to formalize the common domain of discourse. If two of them need the specification of some concepts on which the third actor disagrees, they have two possibilities:

1. either ignoring the disputed concepts; this results in a poor and weak agreement that prevents the two concordant actors from fully collaborate;
2. or including the controversial concepts; this results in a stronger agreement that however can heavily hinder the relationship with the third actor.

We believe that the *shared agreement* must be large enough to allow for a *meaningful communication* between all the three parties to be integrated. It can include details that are needed by two parties and inessential for the remaining one. However, with this extended meaning of shared agreement, the third actor will refuse to commit directly to the *entire* agreement, because it includes details that are useless from its point of view.

Our best practice proposes to solve the problem by allowing for a *partial commitment* to the shared agreement, which means that each party engages itself in a durable obligation of complying with just a part of the agreement. Parties that need details will commit to the part of the agreement that contains them, whereas those that do not need them will not.

Therefore, to preserve the common negotiation and to enable, as much as possible, a win-win situation between all the parties of a business deal, we suggest to *keep the separation between the agreement and the commitment to it*. In other words:

– by *agreement*, we mean the description of a common field of action, potentially going in all the details which enable a meaningful communication between at least two parties of the arrangement; while,

– by *commitment*, we mean the specification of (the part of) the agreement on which an actor engages itself by complying with it; the commitment therefore implies an active and durable obligation to a part of the agreement, while preserving the party's reservation and disagreement on the rest of the agreement itself.

We distinguish from the definition of ontological commitment given in [9] in that we extend the notion of in-completeness of the commitment from the factual knowledge to the conceptual knowledge: each party (or agent) is free to discard parts of the ontological agreement and to undertake to the sharable portion only.

As mentioned in the section 2, the realization of our definition of agreement relies in the explicit conceptual formalization of the *reference ontology*. This means that, whenever three or more parties settle an arrangement and reach a shared agreement, they formalize it in the reference ontology. The ways to (partially) commit to the reference ontology are detailed in section 5.

5 Preserving the Disagreement

The commitment to the agreement is the way to express the accordance and the dissent with regards to the mutual consent. Therefore, we need a way to explicitly express the disagreement of a party on the object of the common action field.

We believe that, in order to preserve the disagreement, it is useful to consider two orthogonal dimensions of the commitment: positive vs. negative commitment and public vs. private commitment. A final summary of these notions is displayed in table 1 on page 8.

5.1 Positive vs. Negative Commitment

We say that a party expresses a *positive commitment* whenever it imports a part of the reference ontology or aligns a part of its local ontology with the reference one. Positive commitment provide a means for meaningful communication on the basis of the parts of the reference ontology that received commitment by other parties. Such parties, thus, can ask queries to the positive-committed party and they may expect answers whenever the knowledge base of the party contains useful facts.

On the contrary, a party expresses a *negative commitment* in two ways. The easiest one is by avoiding to express a positive commitment on a part of the reference ontology. However, this way may not be enough: the second way consists in formalizing parties' specific points of view (i.e., the source of disagreement).

The party that *refrains from expressing a positive commitment* on a part of the reference ontology prevents other parties to establish meaningful communication that involves that part of the reference ontology. When other parties ask queries about such part of the reference ontology, the party will not answer, because it does not understand the query.

However, the same party can do more than refusing to commit, it can *formalize its point of view* as a part of its local ontology. As a result, one party

can refine its positive commitment, which is supposed to include leeways for disagreement, by expressing its specific point of view.

For instance, in the case of employment, each European country has its own regulation in terms of job legislation and skills/certifications required to apply for a job. This is especially relevant for physician, lawyers, teachers, etc. Those regulations are mandatory for each country, but, being "local", they cannot fall within the shared agreement (i.e., the reference ontology).

As a concrete example, let's consider two Employment Services (ES), one Swiss and one Italian. Both ESs express a positive commitment on the concepts related to University occupations. However, the legislation about the prerequisites to apply for a University professor position is different between Switzerland and Italy: the two countries disagree on the necessity of holding a Ph.D. title. Therefore, the Swiss ES also makes explicit in its local ontology that each candidate for a professor position should hold a Ph.D. title (whereas in Italy this is not mandatory).

The negative commitment, under the form of an explicit formalization of the local peculiarities, can be used by the discordant party, in order to actively filter the queries it receives.

In the example above, the Swiss ES can write a rule that filters candidates that do not hold a Ph.D title. So, if the Italian ES asks for available University professor positions for an Italian candidate that did not get a Ph.D. title, the Swiss ES will not provide any matching vacancy, even if one is available.

5.2 Public vs. Private Commitment

We say that a party makes its commitment *public* if it discloses to the other parties involved in the agreement its positive/negative commitment. It is worth noting that making the entire commitment public is not mandatory.

We believe that the positive commitment should be always made public, because it lets other parties know in advance if it is worth asking a query (this can be done by checking if the query involves the part of the reference ontology which the receiver commits to). However, a party may decide to only *partially* revealing its positive commitment to all the participants, because it prefers answering only to a subset of them; in this situation, the party discloses its ability to answer certain queries only to the "friendly" parties and by using a different communication channel.

In a similar manner, a party may decide to disclose only part of its negative commitment. Making public a negative commitment is particularly meaningful if it communicates a well known disagreement (e.g., different legislations among EU countries). Such negative commitment can be evaluated by other parties before establishing a collaboration, in order to prevent formally invalid business processes to be activated. However, keeping private large part of the party's specific point of view is quite usual, because by revealing it the party can come into conflict with others.

Such *private* part of the negative commitment can also be used to express filters that prevent "illegal" information from being exchanged. Let's consider

for instance an Italian logistic operator (LO) which is required to transfer a load of soft drugs. Being such request illegal in Italy (but legal in some other countries where the request could have been started), the LO should not only refuse to process the shipment order, but also prevent the request from entering into the elaboration system.

Our Best Practice in Preserving the Disagreement

- *Public Positive Commitment* is the space of "traditional" commitment (as proposed by T. Gruber [9]), which allows for a meaningful communication with other parties.
- *Public Negative Commitment* is the space for disclosing knowledge about well known disagreement that will not cause conflicts among the partners, but that, on the contrary, allows for checking in advance if asking queries to a party is worth a try.
- *Private Positive Commitment* avoids a party to let every other parties know that it committed to answer certain queries; however, it leaves the possibility to ask queries to a subset of partners that are informed about the positive commitment (by using a different channel).
- *Private Negative Commitment* enables the possibility to write filters that prevent "illegal" information to be exchanged, without disclosing knowledge that may cause a party to come into conflict with other ones.

Table 1. The "traditional" notion of commitment can be refined in four different notions, to deploy real Semantic Web application that enable meaningful communication among parties while leaving crucial leeways of disagreement.

5.3 Language and Tool Support

In putting into practice what we discuss above (see also section 6), we were severely limited by the lack of support in standard languages (RDF, RDFS and OWL) and tools.

Import annotations in OWL are not flexible enough; a partial workaround would be a careful modularization of the reference ontology. However, a full implementation of our approach would require an importing clause that enables to select a part of the ontology.

Moreover, in order to express disagreement, we need both a rich ontological language for expressing inequality and a rule language for describing ontology alignments and formalizing filters on data exchange. For what regards the former, we notice that the current support for inequality expression in OWL is limited to `owl:differentFrom` and `owl:disjointWith`; the latter requisite is partially satisfied by different proposals, but no standard Semantic Web rule language exists yet.

For those reasons, we look forward to following the standardization process of OWL 1.1 [11], which includes several new features that we consider interesting for our purposes (e.g., the application of disjointness not only to classes but also

to properties and the possibility to predicate negative assertions). In a similar manner, we expect RIF Core [12] to be of help; however, the tool support will probably be limited until the standardization process is in progress.

Using best-of-breed solutions, in the SEEMP project we adopted the WSML language [13] and WebODE [14] and WSMT [15] as tools. WSML offers a layering of semantic languages including a rule language (WSML-rule), which can be used to describe ontology alignment and to state complex axioms (such as those required to describe filters), on top of an ontological language (WSML-light), which can be used to model the declarative part of the ontologies. WebODE and WSMT were chosen because they provide complete support for WSML, respectively, to model ontologies and describe alignments and filters.

6 The SEEMP Solution

In the following, we present our approach in developing the core of a Semantic Web application that allows parties to partially disagree while doing business together. As explained in details below, our approach mainly comes from the experience we gained during the SEEMP project. The lesson learned and the best practice we would like to share is outlined in table 2.

The Way We Applied Our Best Practice in the SEEMP Project

To formalize agreement and disagreement, we use:

- *Ontologies* as way to express:
  – the general shared agreement among all the parties, which we name *reference ontology*; and
  – the specific points of view of the parties that are only partially sharable with others, which we name *local ontologies*.
- *Ontology-to-Ontology Mediators* as a way to express:
  – the *commitment* toward the reference ontology by (partially) aligning the local and the reference ontologies; and
  – the formal expression of the *disagreement* of the party, by defining filters that exclude the possibility for facts expressed in the local ontology to be translated in the reference form or vice-versa.

Moreover, with regards to the respect of each party's competitive advantage or private data, we suggest:

- to *make publicly available* to all parties the reference ontology and the parts of the local ontologies and the respective mediators which express the local point of view and which is not in contrast with the reference ontology;
- to *keep private* to each party the disagreement expressed in the rest of the local ontologies and the respective mediators.

Table 2. Our recipe to express agreement and disagreement.

We successfully experiment the previously described approach in the SEEMP project, in which concretely we have:

- the *SEEMP Reference Ontology*, as the common agreement between the various ESs that take part in the employment marketplace;
- the various *SEEMP Local Ontologies*, which are the specifications of the local agreements, i.e. the formalizations of the peculiarities of the local environments;
- the *SEEMP Connectors*, as the mediators between the real-world information systems and the shared agreement; they represent both the "positive" commitment onto the Reference Ontology and the "negative" commitment, i.e. the formulation of the specific content or needs of each local ES.

The *Reference Ontology* is a core component of the SEEMP system: it is the common "language" to describe the details of the employment sector. It is rich enough to support the semantic needs of all the ESs, currently involved or subsequently integrated in the SEEMP marketplace. The Reference Ontology, as well as the Local Ontologies, have been developed by following the Methontology approach [14]; this serves to have scalable, adaptable and maintainable ontologies.

For what regards the *Local Ontologies*, in the SEEMP architecture we had two possible options to build them:

1. taking as a seed the Reference Ontology: in this case, the concepts in the Local Ontologies are extension in depth of the concepts already present in the Reference Ontology; the consequence is that the data exchange between different ESs is easy, while the complexity of the Connectors between the Local Ontologies and the local schemata is higher;
2. operating a reverse engineering process from ES schema sources: it is the easiest way for ontologizing ESs, since each concept in a Local Ontology is the semantic expression of a relevant concept in the respective ES; the consequence is that the Connectors between Local Ontologies and schema sources are not complex, while the mappings between the Reference and Local Ontologies can be difficult and cause delay in the data exchange.

The SEEMP way adopts the first option at the beginning, when few ESs are present in the marketplace, while when more ESs join the marketplace, the solution will move toward the second option. The balance between the two options is related to the need for a "minimal shared commitment".

Since each ES talks in its own "language" (the Local Ontology, which represents its view on the employment domain), its respective connector is responsible for *resolving* these *heterogeneity issues*, by translating the local content in terms of the Reference Ontology. In this way, all the ESs in the marketplace speak the same language, and heterogeneity issues are solved, because, rather than managing mappings between every possible ontology pair (which becomes unmanageable as the number of ESs grows), each ES need only maintain mappings to and from the Reference Ontology.

The SEEMP approach offers several advantages, which can be analyzed from the point of view of both CEOs (the decision makers) and CTOs (the IT experts).

The SEEMP solution offers CEOs a way to enforce the subsidiarity principle [16], by valuing each ES contribution in the marketplace. This increases the number of interconnections, which in turn results in a larger number of faithful users and business transactions.

From the point of view of CTOs, SEEMP enables an easier maintenance of the integration with other ESs and lower integration costs. Semantics makes the mappings between different terminologies easier, because tools like WSMT [17] can analyze Reference and Local Ontologies (e.g., by comparing sub-structures and by searching for synonymies) and can guide the IT Administrator in drawing the mappings. Thank to this support, the mapping definition process requires less time or, alternatively, it provides more precise mappings in the same amount of time. This results in a lower integration cost.

There is also an added value from the point of view of the final users of the employment marketplace, i.e. job seekers and employers. Many job offers, that today could only be found by inserting multiple times the CV in each ES and by merging manually the results of different ESs, becomes seamlessly available through the interface of each ES.

In conclusion, to achieve all these results, what has to be built is a comprehensive Reference Ontology that encompasses several employment domains. Development and maintenance of this reference part of SEEMP is more then an ICT problem: it is a matter of reaching agreement at organizational level. As already discussed, the goal of SEEMP is reaching a "minimal shared commitment" in which ESs *agree* on high-level aspects, allowing for collaboration among them, *while disagreeing* on minor details that differentiate each ES from the others.

7 Conclusions and Future Work

During these years, the interest in interoperability issues shifted from a merely technical problem (how to develop Internet applications) to a business problem (how to create value by offering cross-enterprise services to the market).

As a consequence, new constraints arose. When a service is offered by a constellation of actors, a central authority is often difficult to set up and all parties aim at minimizing its role. The central authority should perform only those tasks which cannot be performed effectively by the parties. This principle is known as *subsidiarity* [16]: it states that matters ought to be handled by the smallest (or, the lowest) competent authority. Each party acts independently and a central system (if any) performs only operations that exceed the capacity of individual parties. Subsidiarity assumes that each party is interested in fostering collaboration to create a common value, but it saves room for competing with other parties for market share. Therefore, the principle of subsidiarity emphasizes the importance of an interoperability infrastructures, which supports and empowers each party's actions and which strengthens links among the parties in the market.

Contemporary solutions to interoperability issues need to protect different positions and ideas. They should take into consideration that looking for a comprehensive agreement sometimes can be useless or counter-productive, because the practice to exclude possible causes of conflict from the agreement can produce agreements that are so limited to be almost empty, thus useless. Last but not least, the explicit formalization of disagreement, if well-known and conflict-safe, is needed in most business domains, in order to prevent invalid business transactions from taking place.

In this paper, we proposed a best practice that could help in guiding the deployment of interoperability solutions based on Semantic Web which, in our opinion, has a better chance for industrial uptake. It proposes the formalization of the notion of agreement, commitment and disagreement. It recommends the development of a reference ontology, which is the common source of agreement, and several local ontologies, which express commitment and make the disagreement explicit. The adoption of our best practice in the SEEMP project to the employment market gave us several positive feedbacks and we are currently working in applying it in the NeP4B project to the logistic market.

We believe that we will be able to develop in the near future a full-fledged methodology and a comprehensive tool set, that support several of the tasks that now we can perform only manually.

Part of the method that we are formalizing consists in the iterative construction of both reference and local ontologies, based on the analysis of the public positive feedbacks of each partner. Starting from a basic reference ontology, parties provide rough public positive feedbacks. By formally evaluating the part of the agreement that received the commitment of every partner, it is possible to understand the level of *minimal* meaningful communication that can be established among all the parties. At the same time, each partner can understand which part of the agreement it has to commit to, if it wants to achieve a higher level of communication with one or more parties. After this analysis, the un-committed part of the reference ontology can be eliminated and parties are asked to provide again public positive commitment. After each iteration, the sources of conflict in the reference ontology are reduced and a broader positive commitment is reached. In the end, each party is invited to make public the negative commitment that cause no conflict.

We need to make the method formal, to develop supports tools, especially visualization ones, and to asses its usefulness in several projects. The refinement of the SEEMP solution and the development of NeP4B prototypes will be of great importance in the next years.

Acknowledgments

This research has been partially supported by the SEEMP EU-funded project (IST-4-027347-STREP) and by the NeP4B Italian-funded FIRB project (MIUR-2005-RBNE05XYPW). We thank all the colleagues, and in particular Asunción Gómez-Pérez and Gang Zhao, for the fruitful discussion.

References

1. Robert J. Aumann: Agreeing to Disagree. The Annals of Statistics **4**(6) (1976) 1236–1239
2. Rudi Studer, V. Richard Benjamins, Dieter Fensel: Knowledge engineering: Principles and methods. Data & Knowledge Engineering **25**(1-2) (1998) 161–197
3. Tim Berners-Lee: The Fractal Nature of the Web, working draft. http://www.w3.org/DesignIssues/Fractal.html (1998-2007)
4. Nigel Shadbolt, Tim Berners-Lee, Wendy Hall: The Semantic Web Revisited. IEEE Intelligent Systems **21**(3) (2006) 96–101
5. Jos de Bruijn, Marc Ehrig, Cristina Feier, Francisco Martíns-Recuerda, François Scharffe, Moritz Weiten: Ontology Mediation, Merging, and Aligning. In: Semantic Web Technologies: Trends and Research in Ontology-based Systems. John Wiley & Sons, Ltd. (2006) 95–113
6. Allen Newell: The knowledge level. Artificial Intelligence **18**(1) (1982) 87–127
7. Nicola Guarino: Formal ontology and information systems. In: Proceedings of the 1st International Conference on Formal Ontologies in Information Systems. FOIS'98, IOS Press (1998) 3–15
8. Mallya, A.U., Huhns, M.N.: Commitments among agents. Internet Computing **7**(4) (2003) 90–93
9. Tom Gruber: A translation approach to portable ontology specifications. Knowledge Acquisition **5** (1993) 199–220
10. Gábor Nagypál: Ontology development. In Rudi Studer, Stephan Grimm, Andreas Abecker, eds.: Semantic Web Services, Springer Verlag (2007)
11. Peter F. Patel-Schneider, Ian Horrocks, Boris Motik: OWL 1.1 Web Ontology Language: Structural Specification and Functional-Style Syntax. Technical report, W3C (2006)
12. Harold Boley, Michael Kifer: RIF Core Design. Technical report, W3C (2006)
13. Jos de Bruijn, Holger Lausen, Axel Polleres, Dieter Fensel: The web service modeling language: An overview. In: Proceedings of the 3rd European Semantic Web Conference (ESWC2006), Budva, Montenegro, Springer-Verlag (2006)
14. Gómez-Pérez, A., Fernández-López, M., Corcho, O.: Ontological Engineering. Springer Verlag (2003)
15. Kerrigan, M.: D9.1v0.1 Web Service Modeling Toolkit. Technical report, University of Innsbruck (2005)
16. 2004/387/EC: Decision of the European Parliament and of the Council on Interoperable Delivery of pan-European Services to Public Administrations, 2004 (2004)
17. Adrian Mocan, Emilia Cimpian: Mappings creation using a view based approach. In: MEDIATE-2005. Volume 168 of CEUR Workshop Proceedings., CEUR-WS.org (2005) 97–112