



UNIVERSITY OF TRENTO

DEPARTMENT OF INFORMATION AND COMMUNICATION TECHNOLOGY

38050 Povo – Trento (Italy), Via Sommarive 14
<http://www.dit.unitn.it>

OPEN KNOWLEDGE SEMANTIC WEBS THROUGH PEER-TO-PEER INTERACTION

Dave Robertson, Fausto Giunchiglia, Frank van Harmelen, Maurizio Marchese, Marta Sabou, Marco Schorlemmer, Nigel Shadbolt, Ronnie Siebes, Carles Sierra, Chris Walton, Srinandan Dasmahapatra, Dave Dupplaw, Paul Lewis, Mikalai Yatskevich, Spyros Kotoulas, Adrian Perreau de Pinninck and Antonis Loizou

May 2006

Technical Report # DIT-06-034

Open Knowledge

Semantic Webs Through Peer-to-Peer Interaction

Dave Robertson¹, Fausto Giunchiglia², Frank van Harmelen³, Maurizio Marchese², Marta Sabou⁴, Marco Schorlemmer⁵, Nigel Shadbolt⁶, Ronnie Siebes³, Carles Sierra⁵, Chris Walton¹, Srinandan Dasmahapatra⁶, Dave Dupplaw⁶, Paul Lewis⁶, Mikalai Yatskevich², Spyros Kotoulas³, Adrian Perreau de Pinninck⁵, and Antonis Loizou⁶

¹ Informatics, University of Edinburgh, UK

² Information and Communication Technology, University of Trento, Italy

³ AI Department, Vrije Universiteit Amsterdam, Netherlands

⁴ Knowledge Media Institute, Open University, UK

⁵ Institut d'Investigació en Intel·ligència Artificial, Barcelona, Spain

⁶ Electronics and Computer Science, University of Southampton, UK

Abstract. We present a manifesto for a new form of knowledge sharing that is based not on direct sharing of “true” statements about the world but, instead, is based on sharing descriptions of interactions. By making interaction specifications the currency of knowledge sharing we gain a context to interpreting knowledge that can be transmitted between peers. The narrower notion of semantic commitment we thus obtain requires peers only to commit to meanings of terms for the purposes and duration of the interactions in which they appear. This lightweight semantics allows networks of interaction to be formed between peers using comparatively simple means of tackling the perennial issues of query routing, service composition and ontology matching. Although the entire system described in this manifesto has not been built, all its components use established methods; many of these have been deployed in substantial applications; and we summarise a simple means of integration.

1 Introduction

The Semantic Web is to traditional knowledge representation as the Web is to traditional hypertext. What fundamentally distinguishes the Web from hypertext is not the way pages are represented but the way they are used, via a simple but effective form of automated processing: the Web search engine. From most people’s point of view the workings of this mechanism are mysterious, but mystery is fine when in a short time anyone can find information dispersed across the globe. The interaction required of everyone is always the same - Web search via keywords - and the components sought are either textual pages or programs (generating pages or forms) that individual designers can configure freely.

Semantic Web efforts encourage designers to specify pages and programs in a knowledge representation language. What fundamentally distinguishes the

Semantic Web from traditional knowledge representation is not the formal representations of pages and programs but the way these representations are used in automated processing. The use we have in mind is to improve the accuracy of Web searches and to allow components to be connected automatically rather than simply be discovered individually. It simplifies our explanation, without loss of generality, if we think of all components (pages and programs) as peers capable of supplying information in some formal language: a page supplies information about that page; a program generates information, perhaps requiring input information in order to do so (we return to the data-centric versus process-centric view in Section 6.1). We use the word “peer” at this point simply to underline that no component is given special authority over others, other than via the manner in which components interact with humans and with each other. This way of attempting to share knowledge between peers encounters elementary problems, familiar from traditional software and knowledge engineering:

- When two different components supply exactly the same formal expression this does not imply that they mean the same thing. Conversely, when they supply different formal expressions they may mean the same thing. In traditional knowledge (and software) engineering we avoid this problem by reaching a consensus amongst component designers on an ontology. To obtain an ontology one needs an oracle to decide which formal terms to use. There are only two sources of oracle: human or mechanical. Neither source scales well. Human oracles can only agree in small groups for narrow tasks or domains. Mechanical oracles need to acquire domain knowledge via automated learning methods, none of which has scaled beyond narrow tasks or domains.
- Only some combinations of components work well together, even when the meanings of their appropriate inputs and outputs correspond, because of assumptions made deep within the design of components. In traditional knowledge (and software) engineering this problem is detected through integration testing: by running assemblies of components comprising the system and checking for defects such as range errors (where the type of an input is correct but its value is outside the range anticipated by the component's designers). It is infeasible to test all ranges for all variables for all components so integration testing only covers some small subset of the possible interactions.
- Components may fail without warning, either because they are removed or as a consequence of their environment. Traditionally this problem is solved by ensuring that the components are in a shared, controlled environment. No such controls are available in an open environment, in which we may not always know who, what or where are our peers.
- Even without the problems above, theoretical results (*e.g.* [6]) show that in general it is impossible to guarantee consistently shared common knowledge in asynchronous distributed systems. The engineering solution to this problem is to provide mechanisms for establishing synchronisation between appropriate components, but for this we must have a frame of reference to bound the synchronisation - otherwise synchronisation itself falls foul of the general consistency problem.

As an example, suppose we want to find the addresses of expert researchers on wave energy in Scotland. A conventional Web search using Google doesn't find this information easily because the information about who is an expert is hard to infer from conventional Web pages and the large number of pages discussing wave energy tend to swamp out the few home pages of expert researchers and their groups. A semantic web approach might do better if appropriate information services could be combined. Suppose that an expert finding service does exist and it offers to supply as output a set, S , of names of specialists if given as input a request for experts in some country, X , and discipline, D . A simple specification of this service (made much simpler than, say, an OWL-S, specification since we need only a simple example here) might be this:

service : <i>expert_finder</i> input : <i>request_experts(country(X), discipline(D))</i> output : <i>specialists(X, D, S)</i> operation : <i>get_experts(X, D, S)</i>
--

Suppose also that two address finding services exist in Scotland, allowing one to send a request for the address of a person with a given name, P , and receive an address, A , for him or her. One of these services is run by a UK address company (*uk_address*); the other by the University of Edinburgh (*univ_ed*). Both, by coincidence, have identical specifications:

service : <i>address_finder</i> input : <i>request_address(person(P))</i> output : <i>address(P, A)</i> operation : <i>get_address(P, A)</i>

A semantic web search engine could not obtain the set of addresses we require using only these service specifications (assuming that the specifications connect to actual services via some grounding mechanism not discussed here) because it is missing some vital functionality:

Functionality 1 Interaction specification: *Wherever there is distributed computation there is an issue of control over message passing between services. The expert_finder service supplies a set, S, of names of experts but the address_finder service deals with individual names. Something must infer that we should take each element of S and pass it to the address_finder, collecting the results. We cannot always infer this by examining only the types and values of variables - just like normal programming, someone or something must supply the necessary control structure.*

Functionality 2 Interaction coordination: *We need to know that some form of interaction involving expert_finder and address_finder would be useful for the task. Notice that our task is not, and could not be, specified with the services because the engineers of those services could not be assumed to predict all the possible tasks of which they might be a part. To provide any automation we*

must separately obtain a formal description of the interaction we might need. And this description will most likely be provided using language terms that are different from those used by the services which want to interact with a given service. Some mechanism must make sure that the services are fed appropriate information. The variable, D , in the expert_finder service will need to be bound to the name of some scientific discipline that the service recognises and it is unlikely to recognise all conceivable discipline descriptions. Will it accept “wave energy”? We don’t know simply by looking at the type of the variable.

We have described major problems associated with the aspiration of semantic web efforts. These are traditional engineering problems but their traditional solutions do not apply in open environments. The key to solving these, we shall argue, is in making specifications of interactions an integral part of knowledge exchange between peers. In Section 2 we summarise a compact but expressive formal language for specifying and computing interactions. We then describe the basic mechanisms needed to support this computation in an open environment: ontology alignment (Section 3.1), discovery (Section 3.2) and visualisation (Section 4). Section 5 then combines these elements within a minimal model of interaction, shareable between peers. Finally, since we view this as an evolutionary manifesto, in Section 6 we compare our approach to the major existing paradigms.

2 Interaction Specification

The functional requirement we must satisfy here is: given that a peer has been given a model of interaction, use this computationally to control its own behaviour and communicate to other relevant peers the behaviours expected of them for this interaction.

From previous implementations by the authors using the system described below, there are several different ways in which a peer might exploit an interaction model with a precise computational behaviour:

- It might simply run this model locally, sending messages to other peers but not allowing them to be privy to the broader picture of interaction contained in the model.
- It might distribute appropriate components of the model to those peers with which it interacts, on the assumption that each of them will separately run that component - producing a distributed computation.
- It might use the model as a form of script which it passes entire to the next peer with which it needs to interact - producing complex interactions amongst many peers via pairwise interactions between peers.

None of these computational models is optimal. Server or script based computation keeps the interaction model in one piece, giving advantages in imposing constraints on interaction that apply across groups of peers. On the other hand, a server based computation means that peers other than the one acting as server

have little control over the course of the interaction; while a script based computation only works for interactions that can be deconstructed into a series of pairwise interchanges. Since no winning execution strategy is known, it makes sense to use an interaction modelling language that has a semantics independent of any one of these strategies but that is capable of being executed by any of them.

The need to supply this sort of information has been recognised by many in the semantic web community. The most direct solution is to specify the process of service combination, and the roles undertaken by the services in that process, in an executable language. We give below a specification for our running example in one such language: LCC. Each of clauses 1 to 4 defines the message passing behaviour of a role in the interaction. Clause 1 defines the role we wish some peer (identified by name C) to perform: that of a locator for experts which, given a country, X and a discipline, D , identifies a list of addresses, L . To undertake this role, C must send a message to an expert finder, E , requesting names of experts and then receive a set of names, S from E before taking the role of an address collector. Clause 2 defines the address collector role which involves recursing through the set of names, requesting, and then receiving, an address for each from an address finder, F . Clauses 3 and 4 define our earlier *expert_finder* and *address_finder* services in terms of the message passing required by each of them. Note that these make specific commitments to the temporal behaviours of the services in our interaction (for instance we prescribe that an *expert_finder* is contacted only once in this interaction while an *address_finder* may be contacted many times).

$$\begin{aligned}
& a(\text{expert_locator}(X, D, L), C) :: \\
& \quad \text{request_experts}(X, D) \Rightarrow a(\text{expert_finder}, E) \text{ then} \\
& \quad \text{specialists}(X, D, S) \Leftarrow a(\text{expert_finder}, E) \text{ then} \\
& \quad a(\text{address_collector}(S, L), C) \tag{1} \\
& a(\text{address_collector}(S, L), C) :: \\
& \quad \left(\begin{array}{l} \text{request_address}(P) \Rightarrow a(\text{address_finder}, F) \Leftarrow S = [P|Sr] \wedge L = [A|Lr] \text{ then} \\ \text{address}(P, A) \Leftarrow a(\text{address_finder}, F) \text{ then} \\ a(\text{address_collector}(Sr, Lr), C) \end{array} \right) \text{ or} \\
& \quad \text{null} \Leftarrow S = [] \wedge L = [] \tag{2} \\
& a(\text{expert_finder}, E) :: \\
& \quad \text{request_experts}(X, D) \Leftarrow a(\text{expert_locator}(X, D, L), C) \text{ then} \\
& \quad \text{specialists}(X, D, S) \Rightarrow a(\text{expert_locator}(X, D, L), C) \Leftarrow \text{get_experts}(X, D, S) \tag{3} \\
& a(\text{address_finder}, F) :: \\
& \quad \text{request_address}(P) \Leftarrow a(\text{address_collector}(S), C) \text{ then} \\
& \quad \text{address}(P, A) \Rightarrow a(\text{address_collector}(S), C) \Leftarrow \text{get_address}(P, A) \text{ then} \\
& \quad a(\text{address_finder}, F) \tag{4}
\end{aligned}$$

Fig. 1. Example interaction model

It is not the task of this paper to explain LCC in detail or to justify its use rather than some other process specification language (for that argument see [10]). For our current purposes, the salient features of LCC are:

- It is an executable specification language, providing the features common to such languages (such as variables, data structures and recursion). All of these were needed to deal with even our simple running example.
- Despite its specificity in terms of data and control structure, so we know precisely how we want our services to interact, there remain obstacles to achieving that interaction:
 - Which specific terms will work when communicating between services? This is discussed in Section 3.1.
 - How do we know which actual services to use? In our LCC specification the variables E and F are unbound but messages must be sent to specific services. This is discussed in Section 3.2.
 - What happens if our peers are human operated? For example, the operation $get_experts(X, D, S)$ in Clause 3 might involve asking a human who the experts are. At that point the human needs to know enough of the interaction’s context to give an informative reply (see Section 4).

Principle 1 *Interaction models are declarative, executable specifications that may be understood independently of any particular peer or execution model. The conditions under which they are run, however, requires more run-time support than a traditional executable specification language.*

3 Interaction Coordination

The functional requirement we must satisfy here is: given a peer with no knowledge of how to interact with others to perform some task, obtain for it a description of an appropriate interaction.

One way to do this might be through synthesis, either fully automated (*e.g.* [2]) or interactive (*e.g.* [8]), so that peers could compose appropriate service clusters for whatever task arises. Although an important means of initiating some kinds of interaction specification, synthesis is unlikely to be the main source of this functionality for three reasons. First, as in Section 1, the specifications we need are non-trivial to synthesise automatically and would require specialist expertise to synthesise interactively. Second, synthesis can be time consuming and peers in a semantic web are likely to need fast responses. Third, and perhaps most importantly, sharing information about useful interactions is a way to propagate experience about effective knowledge sharing - so obtaining a model of interaction that has been widely used and is popular with one’s peers may in many cases be better (and much easier) than building one from scratch. We now consider two key enablers for this form of sharing: dynamic ontology matching and peer-to-peer interaction model sharing.

3.1 Dynamic Ontology Matching

In our introduction we argued that traditional methods of ontology matching do not scale to open knowledge sharing systems. Our focus on interaction has not, however, eradicated the issue. Instead, we have a different form of ontology matching problem to consider: matching terms that appear dynamically in the course of an interaction. Being autonomously and independently defined inside each peer, most of these terms will be semantically heterogeneous. Thus, while one peer could have *expert_finder* as its service role, the others could have *person_finder*, *expertICT_finder*, *expert_broker*, and so on. Notice that while the first and the fourth role denote services which are essentially equivalent, the second is more general than *expert_finder*, while the third is less general. It is a fact that these terms are always used in the context of some local, a priori defined, often left implicit, ontological description of the world. And this influences not only the specific equivalent terms used to describe a concept but also the level of generality of the concept itself.

The solution we propose is to construct semantic mappings (*e.g.* more or less general, equivalent to, disjoint from) existing between the terms used during interaction. One such example is the mapping $\langle \textit{expert_finder}, \textit{person_finder}, LG \rangle$, stating that *expert_finder* is less general (*LG*) than *person_finder*. These mappings are those defined and used in C-OWL [9]. Their main advantage over “syntactic mappings”, namely mappings which return an affinity measure in the interval [0,1] (see, *e.g.*, the mappings constructed by the state of the art system COMA [7]), is that the information carried by the semantic relation can then be exploited in many ways, for instance when fixing mismatches (see “Term matching” below).

We discover semantic mappings, using the method implemented in the S-Match system [5]. This is applied in at least three different phases:

- *Role matching*: aligns the different ways in which roles are described when initiating or joining an interaction. An example is the mapping $\langle \textit{expert_finder}, \textit{person_finder}, LG \rangle$.
- *Term matching*: aligns (structured) terms within the clause defining a role in order to undertake an interaction. An example is matching *get_address*, which in one peer could take two arguments (*e.g.* name of a person and his/her address) and in another three arguments, where the third argument (*e.g.* *Type_of_Comm*) could discriminate whether we need an address for personal or work communication.
- *Query / Answer matching*: takes place when running an interaction model and deals with the semantic heterogeneity arising from the statement of a query and in the values returned in its answers. For example, an interaction model specifying that the *address_finder* needs to look up the address for Stephen Salter by invoking the *get_address('Stephen Salter', A)* operation is not guaranteed (as we are with the ontology of messages) to match perfectly to the operation the peer actually can perform. Perhaps the operation used by the peer is *find_address* and the surname is expected first (as in '*Salter, Stephen*').

The kind of (semantic) matching that we need here differs from the previous approaches in that it is not done once and for all, at design time on statically defined ontologies, but, rather, it is performed at run-time. This moves the problem of *ontology* (and data) *integration*, widely studied in the literature, to the problem of *ontology coordination*. As discussed in [4], the problem of data and ontology coordination is characterised by various new difficulties (beyond the obvious requirement of very fast response times). Coordination is dynamic so exactly what needs to be coordinated depends on what is interacting. Peers have partial knowledge of their network so cannot always predict which other peers will interact with them. Ontology matches made in this way without prior consensus are therefore intended to support query answer sets that are *good enough* for the interaction in hand but not necessarily complete or always correct.

This form of dynamic semantic matching might at first sight seem a harder problem than conventional static matching. We can, however, utilise our coordination framework to turn it into a more limited, easier problem. From Principle 1, interaction models are shareable resources. So too are mappings used with them. This cuts down the number of new mappings that are needed as models are reused. Furthermore, the context in which mappings are constructed is much more limited than with traditional ontology mapping. Since the purpose of mappings is to ensure that a specific interaction model functions correctly then the only issue is whether the meaning associated with the particular use of an operation required of a peer corresponds to an operation it can do. This is a much less demanding task than the general task of mapping entire ontologies between peers, for two reasons. First, we need consider only the fragment of the ontology that applies to a specific interaction. Second, the commitments we make in a mapping can be weaker and more easily judged: for instance, mapping “visit to Italy” to “holiday trip” may make perfect sense for an interaction with a holiday service even though the mapping does not hold in general.

Principle 2 *Models of interactions are developed locally to peers. However, they must be shared in order to achieve interaction coordination. This is achieved by dynamically matching, at run time, terms in the interaction models. This happens both when synthesising them and when running them to answer specific queries. Dynamic semantic matching does this by considering the terms in the context (defined as a local ontology) of the involved peers.*

3.2 Interaction Model Sharing and Discovery

The obstacle to a peer wishing to acquire an interaction model in an open system is knowing who to ask and how to ask. We cannot expect a peer to know all other peers or to know the details of an interaction. We can expect, however, that each peer knows about some distributed discovery service(s) and that each peer knows some features of the interaction in which it is willing to participate. These features need not be complex - keyword matching may suffice. In our running example, a peer with no knowledge of how to find addresses of experts might ask the discovery service for interaction models described by the keywords

{*expert, address*}. The task of the semantic discovery system is then to locate interaction models with these features and peers that want to participate in these interactions across the peer network.

Principle 2 separates interaction models from services, with the advantage that interactions can be shared. To make them work, however, they must connect to specific services. This requires choice. In our running example we must choose, when binding *F* (the *address_finder*) whether we use *uk_address* or *univ_ed*? That may depend on the names suggested as wave energy experts by the *expert_finder* service. If 'Stephen Salter' were a name suggested then *univ_ed* might be the best choice (since Salter is a professor at Edinburgh). If not, then *uk_address* might be a better bet because it is more general. Hence the context set by earlier interactions conditions subsequent choices of services.

There are three (non-exclusive) ways to tackle this issue:

Sharing of experience between peers : The most direct way to overcome problems like those above is (like traditional Web search engines) to refer to records of past experience. If we know, for example, that someone else had used the LCC specification in our example successfully with the variable bindings $E = e_find, F = univ_ed, X = 'Scotland'$ and $D = 'wave\ energy'$ then if we were to ask a similar query we might follow a similar binding pattern.

Using the semantic discovery service : To facilitate ranking of peers or interaction models, the semantic discovery service can calculate and maintain statistical information about keywords and contexts from all interaction models and all peers in the system.

Collaborative recommendation of interaction models and services : By identifying emergent groups amongst peers, based on sharing of interaction models between peers, groups of users that are likely to use services under similar contexts can be inferred. In addition we can require that the local ontologies of peers in the same cluster (or at least those segments relevant to the candidate interaction protocols) can be automatically mapped to each other. In this setting the problem of choosing the best candidate component is reduced to collaborative filtering. The appropriateness of interaction models retrieved by the system and specific services can be assessed based on the frequency of their use within the community, while the added mapping requirement can ensure that the input parameters will be provided in the format expected by the service.

Principle 3 *The choice of interaction models and peers to occupy roles in them is determined by a distributed discovery service. Evidence of role performance in interactions may be routed to this service. Interaction models in need of role performers and role performers in need of related interaction models consult this service.*

In recent years, a wide variety of resource discovery systems have been proposed and developed. Most of them, like UDDI and Napster, support attribute-based retrieval. Their major disadvantage is that they are centralised approaches.

To alleviate these problems, many peer-to-peer indexing systems have been proposed, and basic methods such as distributed hash tables, latent semantic indexing and semantic overlays have been developed.

We assume that data or services are described through sets of terms which we call *descriptions* and that the system contains a large number of such descriptions. Our routing methods extend the work described in [12] to provide statistical information about terms. Initially, peers use a distributed hash table to muster data and group descriptions by term. Since we make no assumptions of shared ontologies, different terms may be used for similar concepts. In [11] a method for automatic word sense discrimination is proposed. Words (or in our case terms) and their senses are represented in a real-valued high-dimensional space where closeness in the space corresponds to semantic similarity. As input for this method we use the descriptions for each term. From the representations of terms in the high-dimensional space, we can then extract information about term generality, term popularity, related terms and homonyms.

4 Visualisation for User Interaction and Interaction Monitoring

There are two constraints in the interaction model for our running example: *get_experts(X, D, S)* in clause 3 and *get_address(P, A)* in clause 3. We have assumed that each of these would be satisfied automatically via a service call. Sometimes, however, constraints may need to be satisfied by human interaction if all automated means of constraint satisfaction fail. A variant of our example might involve the satisfaction of *get_experts(X, D, S)* by a human expert who expresses an opinion on the set, *S*, of experts in a given country, *X*, in domain of expertise, *D*. When these interactive constraints must be satisfied then it is necessary to link interaction model specifications to visualisation specifications (which then can be interpreted via a Web browser or similar mechanism). There are several (non-exclusive) ways to do this:

- By carrying the visualisation specification with the interaction model. This allows interaction model designers to customise visualisations to interactions.
- By providing alternative visualisation methods on peers. This allows limited local customisation to account for style choices.
- By building customised visualisers for very heavily used interaction models. This is appropriate for tasks where the visualisation is the inspiration for the interaction model - for example if we wished to have a complex geospatial visualisation on peers but maintain consistency across peers of information viewed within that visualisation framework.
- By generating visualisation directly from the structure of an interaction model. This is appropriate for tasks such as monitoring the state of an interaction or investigating a failure - a facility not essential to all users but essential for some.

Principle 4 *Interaction models must permit versatility in visualisation: providing default visualisations for common structures but also allowing customisation of visualisation by both peers and interaction model designers. This can be achieved by adopting a standard, declarative markup language for visualisation that each peer may interpret without needing to understand the deeper semantics of the constraints themselves.*

5 A Minimal, Most General Interaction Model

We have used the interaction model of Figure 1 as an illustrative example of the various reasoning components needed for peer-to-peer discovery, sharing and collaboration. This example is domain specific (as are most interactions) but we can also specify more general forms of interaction. The most general of these describes how a peer manages other interaction models.

$$\begin{aligned}
 & a(\text{peer}, X) :: \\
 & \quad (a(\text{discoverer}, X) \text{ or } a(\text{collaborator}, X) \text{ or } a(\text{sharer}, X)) \text{ then} \quad (5) \\
 & \quad a(\text{peer}, X) \\
 & a(\text{discoverer}, X) :: \\
 & \quad \left(\begin{array}{l} \text{descriptors}(K) \Leftarrow a(\text{discoverer}, Y) \text{ then} \\ \text{discovered}(\mathcal{P}) \Rightarrow a(\text{discoverer}, Y) \Leftarrow \text{locate}(K, \mathcal{P}) \end{array} \right) \text{ or} \quad (6) \\
 & \quad \left(\begin{array}{l} \text{descriptors}(K) \Rightarrow a(\text{discoverer}, Y) \text{ then} \\ \text{add_to_interaction_cache}(\mathcal{P}) \Leftarrow \text{discovered}(\mathcal{P}) \Leftarrow a(\text{discoverer}, Y) \end{array} \right) \\
 & a(\text{collaborator}, X) :: \\
 & \quad m(M, \mathcal{P}, \mathcal{S}, \mathcal{O}) \Leftarrow a(\text{collaborator}, Y) \text{ then} \\
 & \quad \left(\begin{array}{l} m(M', \mathcal{P}'', \mathcal{S}'', \mathcal{O}') \Rightarrow a(\text{collaborator}, Z) \Leftarrow \left(\begin{array}{l} \text{match}(M, \mathcal{P}, \mathcal{S}, \mathcal{O}, M', \mathcal{P}', \mathcal{S}', \mathcal{O}') \wedge \\ \text{expand}(M', \mathcal{P}', \mathcal{S}', \mathcal{P}'', \mathcal{S}'', Z) \end{array} \right) \text{ or} \\ m(M, \mathcal{P}, \mathcal{S}, \mathcal{O}) \Rightarrow a(\text{collaborator}, Z) \Leftarrow \text{routing}(M, \mathcal{P}, \mathcal{S}, \mathcal{O}, Z) \text{ or} \\ \text{null} \Leftarrow \text{completed}(M', \mathcal{P}', \mathcal{S}') \end{array} \right) \quad (7) \\
 & a(\text{sharer}, X) :: \\
 & \quad \text{share}(\mathcal{P}) \Rightarrow a(\text{sharer}, Y) \Leftarrow \text{in_interaction_cache}(\mathcal{P}) \text{ or} \quad (8) \\
 & \quad \text{add_to_interaction_cache}(\mathcal{P}) \Leftarrow \text{share}(\mathcal{P}) \Leftarrow a(\text{sharer}, Y)
 \end{aligned}$$

Fig. 2. Interaction model for a basic peer

Figure 2 gives a minimal specification of peer interaction (reduced to its essentials to save space). It describes the key roles of a peer: discovery and sharing of interaction models; and collaboration driven from interaction models communicated between peers. In the definition: \mathcal{P} is an interaction model; \mathcal{S} is the state of the interaction in which the peer currently is engaged; \mathcal{O} is the set of ontology mapping applying to \mathcal{S} ; M is a message; X , is the unique

identifier of a peer; *locate*(K, \mathcal{P}) means the peer can find a \mathcal{P} , described by keyword list, K ; *add_to_interaction_cache*(\mathcal{P}) adds \mathcal{P} , to the cache known to the peer; *in_interaction_cache*(\mathcal{P}) select \mathcal{P} , from the cache known to the peer; *match*($M, \mathcal{P}, \mathcal{S}, \mathcal{O}, M', \mathcal{P}', \mathcal{S}', \mathcal{O}'$) extends \mathcal{O} , adapting $(\mathcal{P}, \mathcal{S})$, to peer giving $(\mathcal{P}', \mathcal{S}', \mathcal{O}')$; *expand*($M, \mathcal{P}, \mathcal{S}, \mathcal{S}', M', Z$) expands \mathcal{S} given M , yielding M' sent to per Z ; and *completed*($M, \mathcal{P}', \mathcal{S}'$) means that M completes this interaction for this peer.

Figure 2, of course, is not complete because it does not define important aspects of routing, ontology matching, *etc.* that we discussed in earlier sections. For these the reader is referred to the papers cited in appropriate earlier sections. Our point here is that even the basic knowledge sharing operations of a peer can be understood in terms of shareable interaction models.

6 Comparison to Current Paradigms

6.1 The Data-centric and Process-centric Semantic Web

One view of the Semantic Web is data-centric, where nodes are data sources with which we associate formal specifications of content. The onus is on curators of data (or engineers of services supplying data) to author their specifications appropriately so that generic systems can discover and use data, guided by this additional information. The intention in doing this is to describe key aspects of the semantics of content - so called, semantic annotations. The difficulty in practise with using only semantic annotations is that to gain consensus on what the annotations should be it is necessary for them to be used for practical purposes by the peer group to which they are relevant. From this it follows that the data-centric paradigm needs to be supported by a way of sharing patterns of usage and knitting them into semantic annotations. The interaction models described in this paper are a means of expressing such patterns. Peer-to-peer routing makes it possible to share these on a large scale. Various forms of ontological alignment (including manual alignment) can then be applied to allow peers to select (and collectively reinforce) specific patterns of usage that work when combining data.

The need to represent potential interactions has long been recognised, hence the process specification elements of OWL-S. In OWL-S, however, an interaction process is associated with a service (and with its data) rather than being separately defined. By separating interaction specifications from data annotations we obtain three crucial advantages:

- We no longer have to define generic processes for services. Instead we expect to have many, perhaps very domain specific, interaction models that we then align more narrowly with services.
- We no longer have to aim for broad ontological consensus across services because data is used via narrow alignments.
- By losing the above two restrictions we are able to knit together services with less sophisticated formal languages.

6.2 Web Service Architecture

Our approach is intended to complement and extend a traditional Web service architecture by addressing a number of restrictions. The key extensions that we are proposing, and the restrictions that they address, are summarised below:

- The Web Service Architecture, while distributed, is not inherently peer-to-peer. In particular, there is no support for efficient routing of messages between services, service discovery is performed in a centralised manner using registries, and there is the assumption that services will always be available at a fixed location. In our Peer-based architecture we relax these restrictions. We provide efficient query routing between components to prevent bottlenecks, we support component discovery using distributed search techniques, and we can cope with components that are not always available through dynamic substitution.
- The lightweight interaction models that we define avoid problems associated with dynamic service composition. Our models define precisely how interaction should be performed with individual components, and also how composition of components should be performed. We do not rely upon complex planning operations or require the construction of detailed workflows by users, although our methods do not exclude these methods where appropriate.
- The basic Web services architecture does not contain any support for assessing trust across services when conducting interactions. Because our methods maintain explicit models of interaction to coordinate services we can apply a repertoire of trust assessment methods to these: from evidence based or provenance-based methods through to methods based on statistical analysis (on groups of interactions) or social analysis (on groups of peers with shared interactions). Importantly, we can associate different measures of trust with appropriate interactions, since one measure will not fit all situations.

6.3 Grids

In [1] three generations of grids are identified: first generation, where proprietary solutions are aimed mainly at sharing computational resources of large super-computing centres for high-performance applications; second generation, where (through middleware and standardisation) grids have become more ubiquitous; and third generation, which shifts to a service-oriented architecture, and the use of meta-data to describe these services. In this third generation, services are given well-defined, machine-processable meaning so as to enable autonomic configuration of the grid and assembly of services - the so-called “semantic grid” elaborated in [3].

Our approach is consistent with the semantic grid vision but without our methods being predicated on sophisticated underlying Grid infrastructure. Traditional grid systems connect together specific services (or types of service) in stable networks - the aim being to do as much as possible to make these networks robust, secure and resistant to failure. We concentrate on specifying interactions,

which may take place with different combinations of services in an open, peer-to-peer environment where the only essential infrastructure requirement is the ability to pass messages between peers. In this sense, our aim is an “everyman’s grid” in the sense that we aim to maintain integrity of interaction (a key grid objective) without requiring specialist (centralised) infrastructure or computing resources to do so and at a very low entry cost.

Acknowledgements

This research is funded by the OpenKnowledge project (www.openk.org).

References

1. F. Berman, G. Fox, and A. Hey, editors. *Grid Computing: Making the Global Infrastructure a Reality*. John Wiley & Sons, Inc., New York, NY, USA, 2003.
2. J. Blyth, E. Deelman, and Y. Gil. Automatically composed workflows for grid environments. *IEEE Intelligent Systems*, July/August, 2004.
3. D. De Roure, N.R. Jennings, and N.R. Shadbolt. The semantic grid: A future e-science infrastructure. In F. Berman, G. Fox, and A. J. G. Hey, editors, *Grid Computing - Making the Global Infrastructure a Reality*, pages 437–470. John Wiley and Sons Ltd., 2003.
4. Giunchiglia F. and Zaihrayeu I. Making peer databases interact: A vision for an architecture supporting data coordination. In Klusch M., Ossowski S., and Shehory O., editors, *6th international workshop on cooperative information agents (CIA 2002)*, Madrid, 18-20 September 2002.
5. Giunchiglia F., Shvaiko P., and Yatskevich M. S-match: an algorithm and an implementation of semantic matching. In Bussler C., Davies J., Fensel D., and Studer R., editors, *The semantic web: research and applications, 1st european semantic web symposium (ESWS'04)*, volume 3053 of *LNCS*, Heraklion, 10-12 May 2004.
6. J. Y. Halpen and Y. Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3):549–587, 1990.
7. H.H.Do and E. Rahm. COMA - a system for flexible combination of schema matching approaches. In *Proceedings of Very Large Data Bases Conference (VLDB)*, pages 610–621, 2001.
8. T. Oinn, M. Addis, J. Ferris, D. Marvin, M. Senger, M. Greenwood, T. Carver, K. Glover, M. Pocock, A. Wipat, and P. Li. Taverna: a tool for the composition and enactment of bioinformatics workflows. *Bioinformatics*, 20(17):3045–3054, 2004.
9. Bouquet P., Giunchiglia F., Van Harmelen F., Serafini L., and Stuckenschmidt H. C-OWL: contextualizing ontologies. In Fensel D., Sycara K. P., and Mylopoulos J., editors, *2nd international semantic web conference (ISWC 2003)*, Sanibel Island (Fla.), 20-23 October 2003.
10. D. Robertson. Multi-agent coordination as distributed logic programming. In *International Conference on Logic Programming*, Sant-Malo, France, 2004.
11. Hinrich Schutze. Automatic word sense discrimination. *Computational Linguistics*, 24(1):97–123, 1998.
12. R. Siebes. pnear - combining content clustering and distributed hash tables. In *Proceedings of the IEEE'05 Workshop on Peer-to-Peer Knowledge Management*, San Diego, CA, USA, July 2005.