

# Ontology Translation for Interoperability Among Semantic Web Services

Mark H. Burstein and Drew V. McDermott

■ Research on semantic web services promises greater interoperability among software agents and web services by enabling content-based automated service discovery and interaction and by utilizing . Although this is to be based on use of shared ontologies published on the semantic web, services produced and described by different developers may well use different, perhaps partly overlapping, sets of ontologies. Interoperability will depend on ontology mappings and architectures supporting the associated translation processes. The question we ask is, does the traditional approach of introducing mediator agents to translate messages between requestors and services work in such an open environment? This article reviews some of the processing assumptions that were made in the development of the semantic web service modeling ontology OWL-S and argues that, as a practical matter, the translation function cannot always be isolated in mediators. Ontology mappings need to be published on the semantic web just as ontologies themselves are. The translation for service discovery, service process model interpretation, task negotiation, service invocation, and response interpretation may then be distributed to various places in the architecture so that translation can be done in the specific goal-oriented informational contexts of the agents performing these processes. We present arguments for assigning translation responsibility to particular agents in the cases of service invocation, response translation, and match-making.

Semantic web services are to the semantic web what current-day web services are to the web, as we know it now. Tim Berners-Lee, Jim Hendler, and Ora Lassila's vision of the semantic web (Berners-Lee, Hendler, and Lassi-

la 2001) was of a future web populated by pages enriched by their association with sharable semantic representations, which describe both content and, in the case of services, functionality. Semantic web services will publish machine-interpretable descriptions of their capabilities and interaction models so other software agents can find and use them without prior built-in knowledge about how to call their application program interfaces (APIs). They will soon support the development of personal software agents or semantic web clients for such things as comparative shopping, information discovery and travel planning, and compositions of those services. Less glamorously, but perhaps more importantly, these techniques may soon enable business-to-business interactions that are more dynamic, support semiautomated service composition on the scientific computing grid, and enable mobile, wireless devices to interact seamlessly with the services discovered as they move about.

To fulfill these promises, published semantic service descriptions must be used in a variety of ways. Services will be *discovered* by agents matching a client's service requirements against service capabilities. Clients will *invoke* services by deducing from their descriptions the content of the messages required to request those services and interpret their responses, which may range from straightforward acknowledgements to indications of failure to requests for additional information. Finally, by using these descriptions of each service's effects and usage constraints, agents may compose multiple services, roughly the way classical AI

planners use planning operators. We use the word *agent* for these clients to emphasize the goal of giving them the ability to reason about the services they deal with.

Internally, semantic web service clients must be able to determine when to outsource an internal goal or function to a remote service, select among some suggested candidate services, and reason about *how* to interact with the selected service based on the service's published description and the clients' own internal goals and knowledge. This includes decisions about how to provide ancillary information the services may require (such as credit information, access certificates, and so on). Since services may require extended interactions, service descriptions may include interaction protocols that these clients must be able to follow. A noteworthy example is failure-recovery procedures.

A number of auxiliary semantic web agents can help in achieving these aims. Semantic matchmakers (Paolucci et al. 2002a) that act like web search engines or intelligent universal description, discovery, and integration (UDDI)<sup>1</sup> web service registries may assist with automated service discovery by cataloging and recommending services to clients. Authentication and policy authorization services may assist both clients and servers to know who they are, and what kinds of interactions they can have. Ontology and mapping registries may help to ensure that agents have consistent and complete sets of concepts, relations, and rules for service-related reasoning.

Semantic service descriptions are developed using a mix of domain ontologies and shared, general-purpose ontologies, such as those defining the structures for representing service capabilities. OWL-S,<sup>2</sup> formerly DAML-S (Ankolenkar et al. 2002), is a semantic web ontology developed by a group of researchers in the Defense Advanced Research Projects Agency's (DARPA) DARPA Agent Markup Language (DAML) program to address these latter, structural aspects of service descriptions. The World Wide Web Consortium (W3C) has formally recommended the Resource Description Framework (RDF) (Lassila and Swick 1999) for building web-compatible structured semantic descriptions and the Ontology Web Language (OWL) (Dean et al. 2003) for publishing ontologies (web documents) defining structured classes and relations. OWL is a semantic description language with a formal semantics and an extensible markup language (XML) / RDF syntax. By leveraging features of RDF, OWL ontologies can be shared, combined, and extended simply by publishing web documents. The

Ontology Web Language for Services (OWL-S) is a set of ontologies, developed in OWL, for writing machine-readable semantic descriptions of web services such as those whose APIs are described using the Web Service Description Language (WSDL) (Christensen et al. 2002).

In this article, we review the basic elements of OWL-S and its intended use model and then discuss the role that ontology mapping and translation must play in interactions between clients and services that use different domain ontologies. We will argue why, in practice, we expect particular agents will need to be responsible for translating the content of messages produced at different stages of their interaction and why it may at times be difficult for mediators to relieve the functional agents (services and clients) of this responsibility.

## Communities Sharing Ontologies and Ontology Mappings

Although this article is about semantic translation of service requests, OWL-S is useful even when clients and services use the same ontologies. Clearly, service interactions are simpler when one doesn't have to translate the meaning of messages sent between agents, so one might ask why not just share a single consistent set of ontologies? When ontologies are shared within a community of people and software services, then major barriers to efficient and timely information sharing are removed. This is why businesses and governments have spent millions of dollars in recent years trying to unify ontologies and database schemas to the largest extent possible. Unfortunately, these efforts have also demonstrated that there are also costs associated with ontology or schema merging that grow with the number of disparate systems and modeling perspectives being combined. At some point, the level of detail in the unified representational model is greater than that needed by any of the individual applications, increasing the complexity of maintaining all of them. An alternative suggested by Uschold and Gruninger (2002), among others, is to share ontologies within tightly integrated communities, while allowing for mediated interaction with other communities. Each community can develop some of its own ontologies and share or extend other widely shared concepts, maintaining locally more detailed models within communities with special responsibilities.

When two communities are going to interact, their ontology developers must define at least partial mappings between the ontologies

each uses. They may also need to extend the original ontologies so that messages that previously had no translation can be translated into the recipients native ontology. Although there is a sizable literature on automated and semiautomated techniques for developing mappings (see Rahm and Bernstein [2001], Noy and Musen [2002], and Kalfoglou and Schorlemmer [2003] for surveys), we anticipate it will require at least some involvement of people familiar with each of the ontologies involved for some time to come. RDF and OWL are designed to enable semantic web communities to share a syntactic language for defining ontological terms and communicating semantic descriptions. We submit, though, that ontology mappings<sup>3</sup> will also need to be published using a similarly standardized language, and agents will also need to interpret these mappings when communicating with “foreign” agents. OWL, as a terminological description language, is not powerful enough to describe all of the kinds of mappings between the concepts that it can be used to define. In our own recent work, we have used first-order logical rules (Dou, McDermott, and Qi 2002, 2003) to define these mappings and have previously explored the controlled application of second-order rules to generate translation programs (Burstein et al. 2003). In this article, we focus on the question of how to use mappings, once published, to interact with services that use different ontologies.

## The OWL-S Ontologies for Semantic Web Services

OWL-S is a collection of “upper ontologies” for describing web services from two primary perspectives, advertising and discovery, and planning and execution. It does not address (or need to address) specific domain issues (such as product taxonomies, taxonomies of types of inputs, and so on) as these ontologies can be developed, shared, and adapted by communities of service providers. OWL-S service models consist of three abstract components. The service profile ontology is used to describe service “advertisements” suitable for use by a semantic web service matchmaker. It includes an extensible set of properties for describing the purpose of a service (as defined in a service hierarchy), its input types, output types, and other features that might be relevant discriminators between services from the perspective of potential users of that service.

An OWL-S matchmaker (Paolucci et al. 2002a) accepts advertisements from services and queries from clients looking for services. It

finds candidate services for clients by matching queries described by partial service profiles against the profiles registered by services.<sup>4</sup> Several OWL-S matchmakers have been developed, including one that extends the web service registry system UDDI (Paolucci et al. 2002b).

The OWL-S service process model includes concepts for describing in detail the interfaces and functional properties of atomic services, namely, their inputs, preconditions, conditional outputs, and effects. It also includes a vocabulary for process composition so service providers can describe how clients should execute sequences of services and their associated data flow. The most recent release of OWL-S treats process specifications as named individuals with complex descriptions, in a manner consistent with the activity model in the Process Specification Language (PSL) (Schlenoff et al. 2000). Preconditions, postconditions, and effects are represented using encapsulated logical expressions.

The OWL-S process model ontology is also based in part on traditional representations of classical AI planning operators, as exemplified by the PDDL language used in the annual AI planning competition (Fox and Long 2003). The premise here is that the reasoning required to identify and invoke an appropriate service is essentially the same as selecting from among a set of planning operators to achieve a goal and then executing it. Input parameters represent the information required to invoke the service, and output parameters represent the information that is returned by the service. Preconditions, and conditional service effects, whose ranges are logical formulas, are to be interpreted by the client reading the service description as additional constraints on the inputs it must provide, with the effects also describing how the service changes the state of the world in a way consistent with the client agent’s goals. Many planning systems have been used to implement OWL-S service clients, including those described in McIlraith and Son (2002), Paolucci et al. (2003), and Wu et al. (2003).

There is good news and bad news about applying traditional planning algorithms to the semantic web service problem. The bad news is that the assumption that the state of the world is completely known at all times is obviously wrong. In fact, a large fraction of the actions taken on the web are to gain information (McDermott, Burstein, and Smith 2002). The good news is that there are a number of planners that can handle simple contingent planning with partial state knowledge and information acquisition goals (see the survey published in

this magazine by Weld [1999]), and the standard notation for action definitions actually fits the web pretty well. Actions are invoked and monitored by messages, and the effect of a message often does consist of changing a few things discretely. Our own work has been based on extending the estimated regression planner, Optop (McDermott 2002, 2003), to do contingent planning with OWL-S process models.

The last part of the OWL-S model, the service grounding, describes the relationships between the inputs and outputs of atomic processes and the elements of a particular message transport model. Specifically, the grounding ontology represents mappings of process parameters onto portions of WSDL message specifications. The OWL-S notion of an atomic process can be viewed as an abstract representation of a WSDL operation that provides additional semantics for it. Although WSDL is rapidly becoming the standard XML language for describing web services, it has no internal means for specifying the semantics of a service's message patterns. WSDL message parameters have XML datatypes, whose semantics must be interpreted by the client program, by virtue of its programmer having read the accompanying documentation.

The OWL-S ontology thus adds the semantic elements necessary for agents to reason dynamically about the relationship between their internal goals and the types of information required to formulate service requests and interpret service responses. Once the information required to create a request has been identified, OWL-S provides a mapping to WSDL (the grounding) to enable the request to be turned into a Simple Object Access Protocol (SOAP) or hypertext markup language (HTML) format.

## OWL-S Usage Model

The basic use model envisioned for OWL-S, illustrated in figure 1, consists of the following steps, not all of which are required in all cases. Services first advertise themselves by sending their service profiles to a matchmaker and by publishing service process models on the web so clients can read them (1). Agents requiring a service to achieve goals pose queries to the matchmaker consisting of a partial or abstract description of the service it requires, including nonfunctional or quality of service requirements (2). The matchmaker compares these abstract descriptions to its library of service profiles and returns URI references for candidates that could be used (3). The querying agent then reads each candidates' published service descriptions and selects one to use. The agent de-

termines how to make a request to the selected service by unifying its goals and preferences with the effects specified in the service process model, in order to determine a semantically valid set of inputs to the server. These inputs are then mapped onto a WSDL input message pattern using the OWL-S service grounding, and the resulting message is sent (6).

The service receives the request and determines whether it can perform the request. It may acknowledge the request, send an error, request additional information, or (generally, on completion) send a reply stating the service results (7). Whatever reply is sent is parsed using a WSDL error or output message template. This message is then mapped (using the process output grounding) into a semantic description of the output parameters of the OWL-S process, from which the agent determines which of the published processes' effects occurred and whether the agent's goal has been accomplished (8).

Viewed abstractly, this process bears some resemblance to the web service use model envisioned for UDDI-WSDL-SOAP, except that it does not involve a programmer querying UDDI, reading the WSDL models found there, and implementing those interfaces. The client software agent is responsible for the interaction with the matchmaker, the interpretation of which candidate services are most appropriate, the determination of the information required to invoke each service, and the interpretation and response to messages returned by the service.

Variations on this execution model for OWL-S have been implemented by a number of academic and industrial researchers (see, for example, Gaio, Lopes, and Botelho [2003]; Paolucci et al. [2003]; and Sirin, Parsia, and Hendler [2004]). These implementations have finessed the ontology translation issue by assuming a shared set of ontologies. However, as suggested above, we cannot assume that agents always share ontologies. Translation will be required to achieve the broad interoperability envisioned for semantic web services, and so we need an architecture that supports it.

## Service Invocation Reasoning across Different Ontologies

The central questions we discuss here are Where does translation fit into an architecture supporting the OWL-S use model? and How do the goals and knowledge of specific agents influence the decision about which agents can or should do translation? The traditional approach to inserting translators into distributed

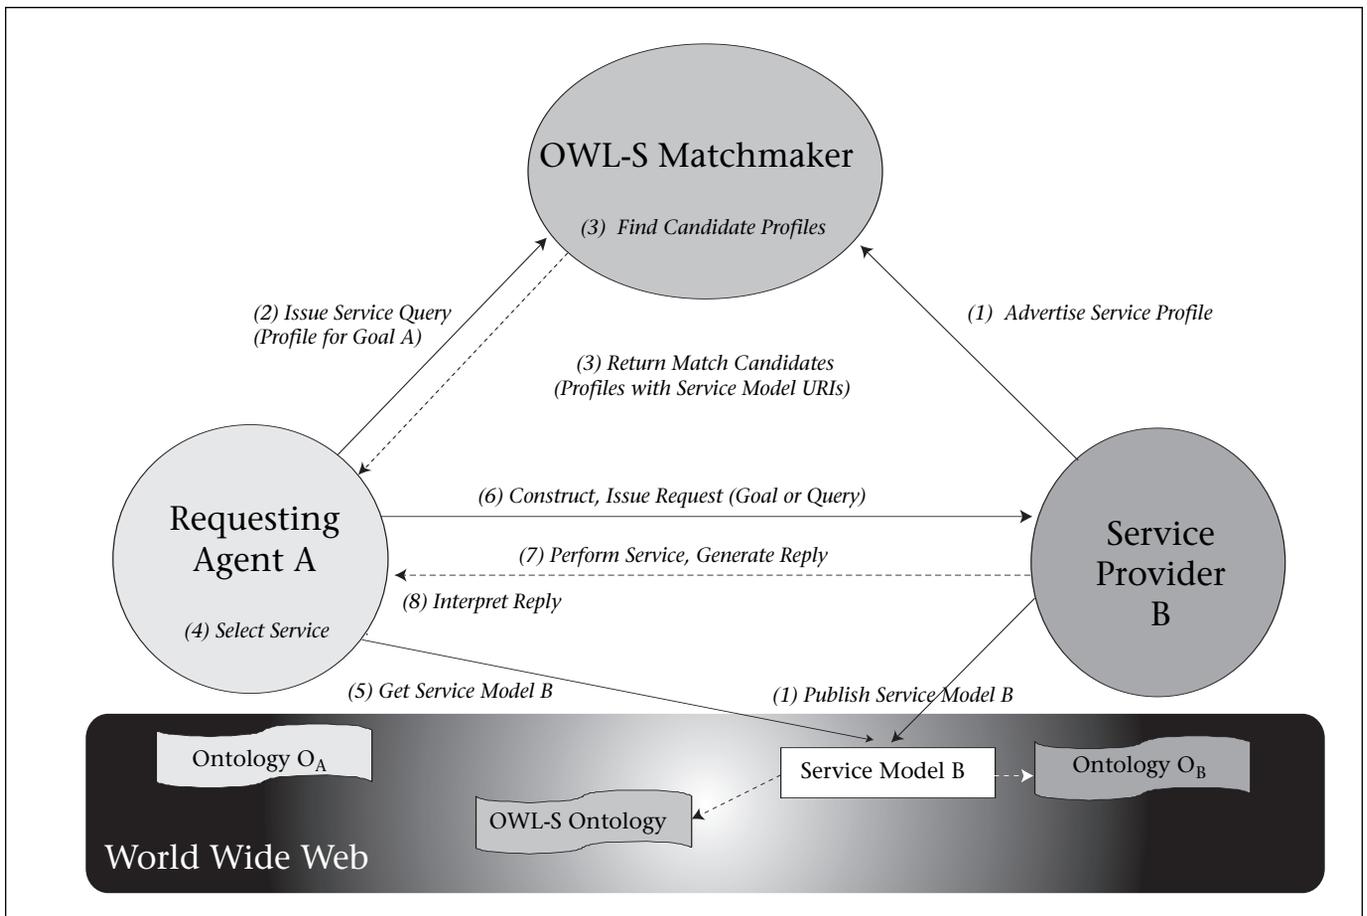


Figure 1: Basic OWL-S Semantic Web Service Use Model.

architectures is to insert a middle agent between client  $C$  and server  $S$  to translate messages from language  $LC$  into language  $LS$  and vice versa. We will argue that this works only when the sender can construct a well-formed message and the target ontologies and ontology mappings are known. It doesn't work, or work well, when a translator needs information that is local to a client or service provider, rather than just the ontologies used by the agent receiving the translation.

Consider the difference between a client's internal goal and a well-formed request to a service. For example, to request a purchase, a service may require specific information about the requestor (such as login identification, valid means of payment) and the manner of accomplishing the goal (for example, shipping method). The request message must be constructed dynamically by reasoning about the relationship between internal client goals and the required input parameters specified in a published service description. The first translation task is therefore the translation of domain elements of the service description from the

ontologies used in the published service description into ontologies used by the client. We have developed a model showing how the client can handle this translation task as it plans how to construct a service request, by reading the service description in its original form along with the ontology mappings needed to interpret the constraints on request input parameters. In an environment with all of these axioms, the client's plan reasoning will effectively translate the parts of the service description it needs to find the needed information in its own knowledge base.

Figure 2 sketches a typical situation in which a personal semantic web agent, which we will call *MyAgent*, is tasked by user Mark to buy a book about XML from an unspecified book sales service on the web. *MyAgent* knows a number of things about Mark, and it knows how to contact commercial semantic web services. *MyAgent* first queries a matchmaker to find a suitable service and discovers there is one called *Books4Sale*, with an OWL-S process model as shown. The *Books4Sale* process model states that, when correctly invoked, it has

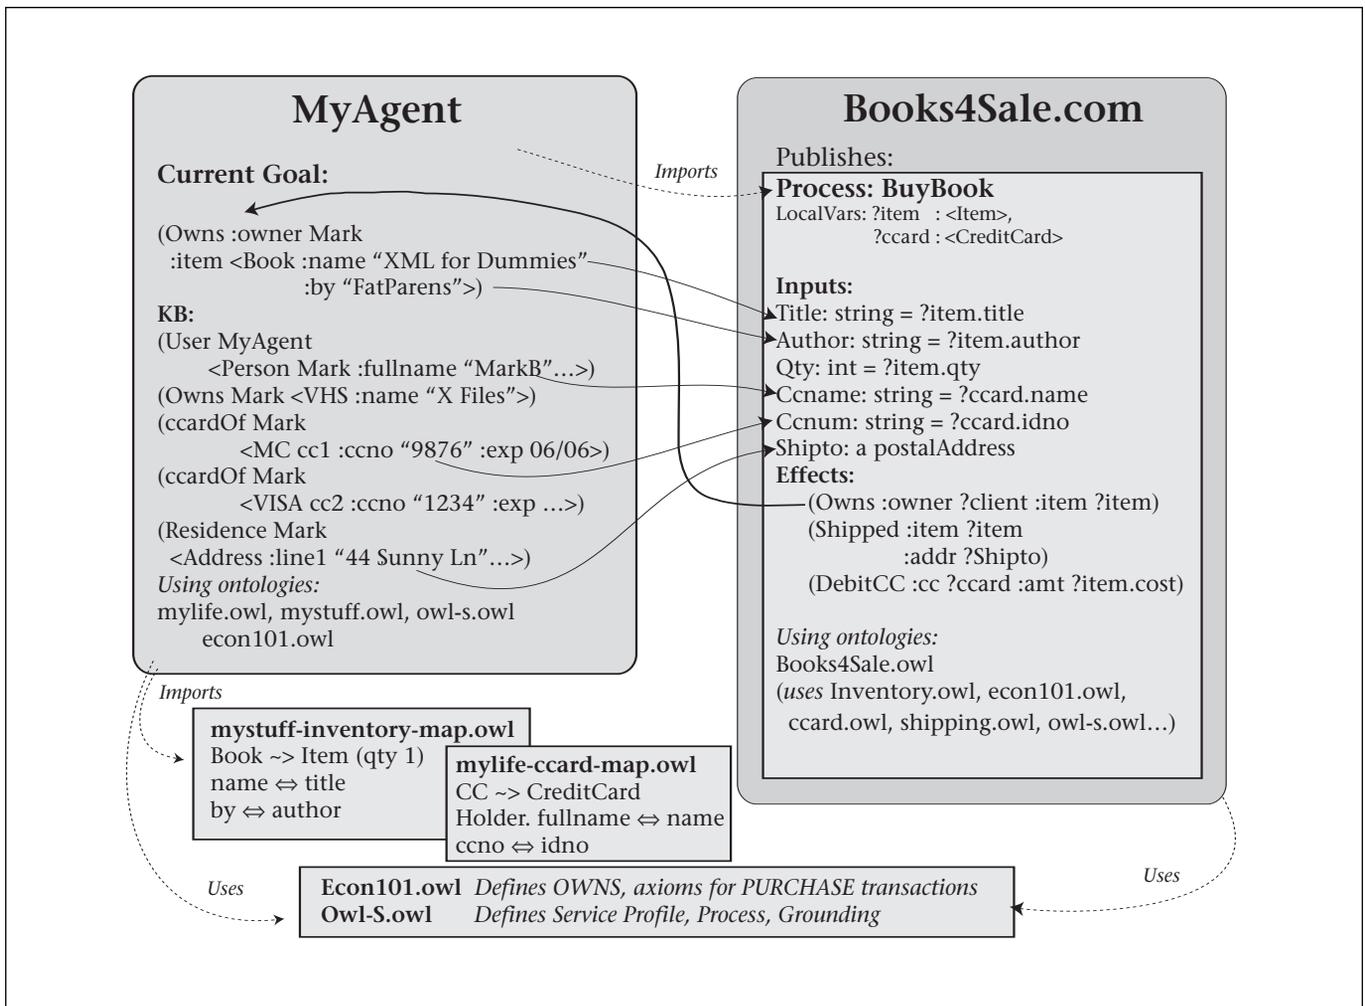


Figure 2. Book Buying Example.

the effect that the client will own an Item with the specified title and author. MyAgent can plan to achieve its goal by determining a set of inputs to this service that will produce an effect matching its goal and by then using the accompanying OWL-S service grounding to execute the action by formulating a WSDL message that can be passed to the Books4Sale service.

The Books4Sale process model specifies as required inputs an item title and author, a credit card number and expiration date, and a shipping address, described using terms from its domain ontologies. It produces as output an order confirmation number and shipping tracking number (not shown). The problem is that MyAgent uses an ontology (MyStuff) different than that used by Books4Sale. MyStuff includes a Book class with a *by* property for the author and a *name* property for the title that is used to represent the object of its goal. Furthermore, instances of the Item class in the Books4Sale service ontology represent different sets of books (identified internally by international

standard book numbers [ISBNs]), each with an associated quantity in stock, while instances of the MyStuff Book class represent individual books, each assigned an inventory number by the owner when purchased. These differences illustrate the kinds of problems that arise when various developers design ontologies for related but different purposes.

Table 1 summarizes the partial mappings between terms in the two ontologies. The two concepts can be related for the purposes of the transaction as long as the key descriptive properties of the Book class (needed as service inputs) have mappings to the corresponding descriptive properties of the Item class and the bridging axioms capture the condition that an Inventory:Item description with qty = 1 can refer to the same entity as a description of a MyStuff:Book.

By expressing these mapping rules as additional axioms in the planner's knowledge base, the required inputs to the BuyBook process associated with the client's goal can be identified

by the planner during goal regression. The identification of these inputs hinges on two things: (1) the client can match its goals to a translation of some of the proposed process's effects, unifying items referenced as arguments to these goals with the variables in those effects, and (2) the type restrictions and preconditions specified for input variables referenced in these effects translate to conditions in the client's ontology consistent with the objects specified in the client's goals.

In short, as in classical planning, the process's effects must unify with the client's goal, and the process's preconditions must be able to be satisfied by the client. The catch is that this goal regression reasoning will only succeed if the constraints are translated into the client's ontology, either in advance by translating the whole service description or incrementally by the planner by backward chaining in the presence of the necessary ontology-mapping rules.

On the semantic web, translating a service process model like BuyBook need not be done by a middle agent if the necessary ontology mappings are available directly. Each client can find and read any ontologies that it does not have that are referenced in messages. If the semantic web architecture were to support mechanisms for finding and loading ontology mappings, then the client could relate descriptions in unfamiliar ontologies to the concepts it natively uses while planning. Taken together with its own axioms, these ontology mapping rules enable it to complete its reasoning whenever the mappings cover all terms used in the target service's parameter type constraints, preconditions, and effects.

Figure 3 shows schematically the situation within the client after loading the service description, the published ontologies it references, and mappings of terms in those ontologies to those of the client. There are sets of ontologies used internally by the client ( $O_{Ci}$ ) and by the service ( $O_{Si}$ ) and some shared ones ( $O_{Shared}$ ), like OWL-S. There are also published mappings ( $M$ ) that relate terms in some of the ontologies of the two partners. In this context the client is able to extend its reasoning to elements of the service ontology  $O_{Si}$  for which mappings are known, reasoning about those parts of the merged ontology as if the client shared them with the server.

Once the client is able to interpret the information requirements of the service and determine how to satisfy those requirements, it can generate the service request. It may even be possible to construct a valid request with an incomplete set of mappings, if the missing map-

MyStuff	Bridging relation	Inventory
Book	InstanceClass-SetClass	Item
Book.name	Equivalent Property	Item.title
Book.by	Equivalent Property	Item.author
Book.pubDate	Equivalent Property	Item.PubDate
Book.purchDate	- N/A -	
Book.itemno	- N/A -	Item.ISBN
	Cardinality(inst) = 1 >>	Item.qty

Table 1. Mappings between MyStuff.owl and Books4Sale.owl.

pings cover optional service inputs. For example, if Books4Sale also used the Item class to represent other goods besides books that it sold, but required different service parameters to specify those items, then the fact that there were no mappings to the MyStuff ontology to translate those additional service parameters could be ignored. In contrast, a separate translation agent, called to translate the full service description for the client, might fail because no complete mapping between the ontologies was available. Furthermore, when the client is responsible for translation, there is the possibility of getting assistance from the user to identify whether the missing mapping information is relevant or perhaps even to supply the mapping.

Elsewhere (Dou et al. 2002, 2003) we described an implemented approach to translation based on this model that can be used to translate datasets (OWL/RDF files). This system (OntoMerge) works by reading the source and target ontologies into a first-order inference engine (OntoEngine), along with a set of bridging axioms or articulation rules (Mitra, Wiederhold, and Kersten 2000) that act as translation rules between terms of the two ontologies. The ontologies can be loaded together without clashing because of the distinct URIs used to identify terms in each model, and these URIs are then linked through references in bridging axioms. If the bridging axioms are bidirectional, translation can be performed in either a push or pull fashion. With the push approach, a dataset is translated by asserting its content and forward chaining to find implications expressed in the target ontology. Similar techniques based on applying sets of rewrite rules have been described in Mitra, Wiederhold, and Kersten (2000) and Chalupsky (2000), although the latter system uses more syntactic transformational rewrite rules. These techniques are useful for translating fully specified messages, such as query or service responses.

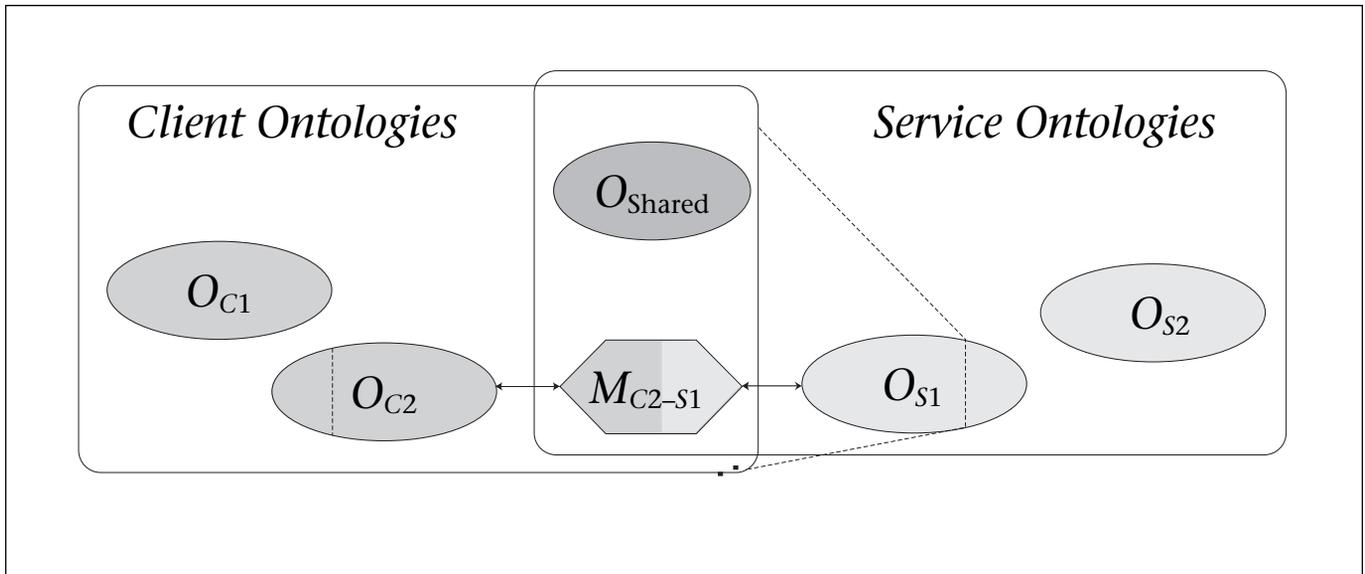


Figure 3. Ontology Merging by Loading Published Ontologies and Mappings

With OntoMerge, however, you can also use a “pull” approach, in which queries are expressed in the target ontology, and OntoMerge backchains to find the answers given the assertions made in the source ontology. We have used this pull approach in the reasoner of the Optop planner to demonstrate the planning of service invocation requests even when the services are described using OWL-S and “foreign” domain ontologies. As the planner reasons about how to satisfy its own goals (by unifying them with service effects) and infer whether the service’s preconditions can be met, its backward chaining naturally utilizes the bridging axioms to locate the information required for service inputs. For example, after loading the bridging axioms of table 1, Optop determines that the title and author inputs to BuyBook can be unified with the *name* and *by* fields of the book specified as its internal goal.

### Identifying Values for Input Parameters Unrelated to Goals

When services publish request message formats, they typically include information that is not represented directly in corresponding client goals. Many other pieces of information, such as a login ID and password may be required. If such required inputs are simple facts known to the client, then these facts can be found by the client during precondition evaluation using the technique previously outlined, because required inputs can be treated like knowledge preconditions of the process. But consider the credit card and shipping information required in our sample book-buying

process. The knowledge precondition for a shipping address is to know the address to which the purchase should be shipped (a service effect), and this may depend on such things as whether it is intended as a gift or not. Decisions of this kind are common when dealing with unfamiliar web services, although when dealing with commercial services there are a number of common types of required information that can be anticipated even when the specific service to be used is unfamiliar. For example, since purchasing anything requires some form of payment, one can anticipate the need to decide among the client’s available credit cards and implement decision criteria based on the cost and purpose of the items and the available balances of the cards. Whether to buy a warranty is another common question. Less common, more item-specific requirements would include special wrapping preferences, nonstandard shipping methods, or whether to buy item-specific optional accessories, and so on.

One way to handle the common options on web service requests is to give the client agent default policies for deciding such questions. For example, there could be a default rule that purchases should be shipped to the address of the buyer (user). It is a default rule since it would not apply when a specific location was specified as part of the user’s goal. Most current AI planning systems deal with nongoal related effects only to the extent that they try to avoid effects that negatively affect their plans. But many services’ inputs effectively ask their client to refine the goal or its object (for example, new or used book?) or require the client to

further specify the method by which it is achieved (for example, selection of a shipping method). Planning systems that can use background goals and preference policies to make these additional distinctions will be needed to more fully automate web service invocation.

These examples highlight why request translation is primarily about making the service description useful to the client rather than translating the request message the client produces once it knows what is required. Since one cannot anticipate all of these input requirements from unfamiliar services, one must translate the constraints on input parameter's implied by service preconditions and effects into knowledge requirements so that the client's decision policies can be applied correctly. Assuming that the basic structure of each service description is provided by the shared OWL-S ontology, then what remains are the domain-specific elements of the model. Either the client translates individual input requirements as it plans its request, as we have suggested, or the service description is entirely translated for the client by a translation agent beforehand.

There are several problems with this. One problem is that ultimately the client must apply the service grounding to send the message, and groundings are published by the service provider using the service provider's ontology. Thus, after determining a valid set of request inputs, the client must translate these parameter values back into the service's own ontologies. If done remotely, this would have to be a separate process from translating the service description for use by the client. Another problem with the external translator approach is that it may fail if there are missing mappings, whereas, if the client does the translation, it could use a partial set of mappings to reason directly with the published description and then ask for user help to decide whether and how to translate any optional parameters for which mappings were unavailable. The specific distinctions made in a service provider's ontology may lead to the real time development of new mappings because the service ontology may make distinctions not known to be relevant to the client agent's ontology developer, and mappings cannot exist where one community or the other does not represent the concepts or relations involved. For example, the MyStuff ontology may not represent the distinction between hardcover and softcover books. If that is a required input to a book-buying service, the requirement for it would fail to be translated, and user advice would be needed. Ultimately, we will need to develop mixed-initiative approaches to automated service interaction in

such open domains that enable our client agents to get additional user assistance and to learn incrementally about the ontologies of unfamiliar services and their mappings into local ontologies.

## Translating Responses to Service Requests

Thus far, we have suggested why service invocation may involve translation within a client, rather than using a middle agent. We now ask whether there are circumstances in which knowledge local to a particular agent is needed to translate service responses. Translating responses to requests or queries is perhaps the most well studied kind of translation. Here, the message to be translated is fully formed by one agent (the service provider) independent of the recipient's knowledge state, and the translator must simply identify the mappings needed to transform the description into one in the recipient's ontology. Most of the work on heterogeneous information retrieval is focused on the problem of developing the necessary ontology mappings for this (Rahm and Bernstein 2001).

In most cases, response translation can be done by any agent that has access to the source and target ontologies and the necessary articulation or mapping rules. However, there are some circumstances in which the mappings cannot be represented precisely, such as when the classification of particular items depends not on the structure of the relations in the ontologies but on their specific extent. Gio Wiederhold has frequently illustrated this by explaining that different administrative offices at Stanford University use slightly different models for who qualifies for membership in the Employee class. One includes only salaried employees, while the other is broader and includes nonsalaried staff, such as emeritus professors like himself. The distinction between these classes is not modeled by systems in either office because neither database represents any attributes that capture the distinction. Thus, no translation rule can be written that succinctly captures the exception to the rule associating the two classes. The only way to translate data about particular individuals being communicated between the two offices is to access the set of all known employees in the target environment and test whether each person whose information is being transmitted is an employee. Effectively, in such situations, the translation needs to be done by the recipient, as a middle agent would not have all the data necessary.

## Translation Issues Associated with Dynamic Service Discovery

The final question we consider is the architectural placement of translation functionality for queries to semantic web service matchmakers. Recall that matchmakers receive advertisements called service profiles from services that wish to be utilized, and clients looking for services find them by querying matchmakers with general descriptions of their requirements represented in the form of a partial or abstract service profiles. Pointers to candidate services are returned for consideration by clients.

So what happens when different services use distinct ontologies to advertise their services? If a matchmaker is to maintain a catalog of service advertisements, then a range of possibilities must be considered. At one end of the spectrum, the matchmaker might use its own set of ontologies to organize the services in a single framework. At the other end of the spectrum, it would maintain every service profile in its original form. Let's call the first approach the *Yahoo approach* and the second the *Google approach*, since Yahoo uses a uniform taxonomy constructed by its staff, and Google indexes web pages based on the content of the original sources.

With the Yahoo approach, each service advertisement must utilize concepts in the matchmaker's ontology to describe the represented service. Clients must then translate their queries into abstract service descriptions using the matchmaker's ontology. With OWL-S profiles, this would mean finding concepts for the class of service desired (for example, a restaurant), the classes of acceptable outputs (Italian food), and the desired effect (for example, food served in-house versus packaged for take-out) and quality of service (for example, maximum price). You might recognize the problem with this from using an online yellow pages. The client needs to be mindful of the expected number of answers. For example, if one wants to find a Chinese restaurant in Rome, is it better to ask for a service that provides listings of Chinese restaurants in Rome? or just restaurants in Rome? or Chinese restaurants in Italy? Finding a useful, targeted answer depends on the number of candidates in each of those somewhat different conjunctive classes. So an effective translation of a service query could depend on both the specific mappings of terms into the matchmaker's ontology and on the selection of appropriate abstractions of those terms to find a small, targeted set of answers. If such knowledge is routinely needed to get the best results, then an architecture where the

matchmaker does the translation is preferable, since the matchmaker can reformulate the query in a context where knowledge about what parts of the registry space are sparsely populated is available.

With the Google approach, services are allowed to describe themselves using their native ontologies. As a result we see a different set of motivations for having the matchmaker translate queries. Here, it is possible in the extreme case that every new service advertisement is described using a different set of ontologies. How can a client agent translate its own query, or have a middle agent do it, when each query must be compared to many profiles, all represented using different ontologies? The only possible strategy requires matching and translation to be performed effectively simultaneously within the matchmaker.

For each kind of semantic service matchmaking, then, there seem to be good reasons to argue that the matchmakers should do the translation of queries they receive. Of course, these are extreme characterizations of the ontology mismatch problems to be faced in connection with service discovery. It is still an open, empirical question how best to design matchmakers that handle a wide diversity of service description ontologies.

## The Impact of Knowledge Locality on Translation Processes

An important motivation for the development of semantic web services is their promise of greater and more dynamic interoperability among agents and services. If the semantic web is going to be based on the principle that a thousand ontologies can bloom, then it is important to look at how the translation processes will be used to enable services to interoperate. These questions have not been a major focus in the semantic web services research community to date.

In this article, we have reviewed the assumptions and motivations for the OWL-S approach to dynamic discovery and utilization of semantic web services in order to frame a discussion of the multiple roles that ontology translation processes must play in facilitating semantic interoperability. Along the way we suggested why published ontology mappings will be important and concluded that it is often better to support translation as a function within agents, rather than separate mediators, in semantic web service architectures.

The question of how and where translation processing happens is critical to the design of

architectures for semantic web services because of the many related, partially overlapping ontologies in use by different communities on the semantic web. For good reasons, each community wants to control its own ontologies. Our concern is with how to enable interoperability of services beyond tight-knit communities, so remote clients can discover and use these services to fill one-time or occasional needs. Semantic web service ontologies make it possible for clients to read published semantic service descriptions and reason with them. Translation of the domain terms in these descriptions will also be needed to support message exchanges.

We have argued that translation functionality can critically depend on knowledge local to particular agents. In particular, we demonstrated how translating a client's goal into a request is best done by the client using published service descriptions and ontology mappings, while translating a client's query to a matchmaker is best done by the matchmaker, because the translation depends on the set of ontologies and profiles known to that agent. The take-away message is that the stronger the need is for a translator to access a localized knowledge context during translation, the more difficult it will be to locate that translation process in an independent middle agent.

### Acknowledgements

This writing of this article was funded under a contract with Yale University to the DARPA DAML program, contract number F30602-00-0600.

### Notes

1. See the UDDI Technical White Paper, [www.uddi.org/](http://www.uddi.org/).
2. See the DAML Services Coalition, OWL-S, [www.daml.org/services/owl-s/](http://www.daml.org/services/owl-s/).
3. See Madhavan et al. (2002) for a discussion of ontology mapping languages.
4. This matching process is not a straightforward subsumption test, as one might expect. As Paolucci et al. (2002a) show, it is asymmetrical. A good match is one in which the classes of inputs the client can provide subsume those required, while the classes of outputs the service claims it can return subsumes those required by the client.

### References

- Ankolenkar, A.; Burstein, M.; Hobbs, J. R.; Lassila, O.; Martin, D. L.; McDermott, D.; McIlraith, S. A.; Narayanan, S.; Paolucci, M.; Payne T. R.; and Sycara, K. 2002. DAML-S: Web Service Description for the Semantic Web. In *The Semantic Web: Proceedings of the First International Semantic Web Conference (ISWC)*, Lecture Notes in Computer Science. Berlin: Springer-Verlag.
- Berners-Lee, T.; Hendler, J.; and Lassila, O. 2001. The Semantic Web. *Scientific American* 284(5): 34–43, 2001.
- Burstein, M.; McDermott, D.; Smith, D.; and Westfold, S. 2003. Derivation of Glue Code for Agent Interoperation. *Journal of Autonomous Agents and Multi-Agent Systems* 6(3): 265–286.
- Chalupsky, H. 2000. OntoMorph: A Translation System for Symbolic Knowledge. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Seventh International Conference (KR2000)*, ed. A. G. Cohn, F. Giunchiglia, and B. Selman. San Francisco: Morgan Kaufmann Publishers.
- Christensen, E.; Cubera, F.; Meredith, G.; and Weerawarana, S. 2002. Web Services Description Language (WSDL). Report, Massachusetts Institute of Technology, Cambridge, MA ([www.w3.org/TR/](http://www.w3.org/TR/)).
- Dean, M.; Connolly, D.; van Harmelen, F.; Hendler, J.; Horrocks, I.; McGuinness, D.; Patel-Schneider, P.; and Stein, L. 2003. OWL Web Ontology Language 1.0 Reference. Cambridge, MA: World Wide Web Consortium (W3C) ([www.w3.org/TR/owl-features/](http://www.w3.org/TR/owl-features/)).
- Dou, D.; McDermott, D.; and Qi, P. 2003. Ontology Translation on the Semantic Web. In *Proceedings of the International Conference on Ontologies, Databases, and Applications of Semantics*. Lecture Notes in Computer Science. Berlin: Springer-Verlag.
- Dou, D.; McDermott, D.; and Qi, P. 2002. Ontology Translation by Ontology Merging and Automated Reasoning. Paper presented at the European Knowledge Acquisition Workshop on Ontologies for Multi-Agent Systems, Bologna, Italy, July 6.
- Fox, M.; and Long, D. 2003. PDDL2.1: An Extension of PDDL for Expressing Temporal Planning Domains. *Journal of Artificial Intelligence Research* 20: 61–124.
- Gaio, S.; Lopes, A.; and Botelho, L. 2003. From DAML-S to Executable Code. In *Agentcities: Challenges in Open Agent Environments*, ed. B. Burg, J. Dale, T. Finin, H. Nakashima, L. Padgham, C. Sierra, and S. Willmott, 25–31. Berlin: Springer-Verlag.
- Kalfoglou, Y., and Schorlemmer, M. 2003. Ontology Mapping: The State of the Art. *Knowledge Engineering Review* 18(1): 1–31.
- Lassila, O.; and Swick, R. 1999. Resource Description Framework (RDF) Model and Syntax Specification. Cambridge, MA: World Wide Web Consortium (W3C) ([www.w3.org/TR/REC-rdf-syntax/](http://www.w3.org/TR/REC-rdf-syntax/)).
- Madhavan, J.; Bernstein, P.; Domingos, P.; and Halevy, A. 2002. Representing and Reasoning about Mappings between Domain Models. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence*, 80–86. Menlo Park, CA: AAAI Press.
- McDermott, D. 2003. Reasoning about Autonomous Processes in an Estimated-Regression Planner. In *Proceedings of the International Conference on Automated Planning and Scheduling*. Menlo Park, CA: AAAI Press.
- McDermott, D. 2002. Estimated-Regression Planning for Interactions with Web Services. In *Proceedings of the Sixth International Conference on AI Planning and Scheduling*. Menlo Park, CA: AAAI Press.
- McDermott, D.; Burstein, M.; and Smith, D. 2002. Overcoming Ontology Mismatches in Transactions with Self-Describing Agents. In *The Emerging Semantic Web: Selected Papers from the First Semantic Web Work-*

ing Symposium, ed. Isabel Cruz, Stefan Decker, Jérôme Euzenat, and Deborah McGuinness, 228–244. Amsterdam: IOW Press.

McIlraith, S.; and Son, T. 2002. Adapting Golog for Composition of Semantic Web Services. In *Proceedings of the Eighth International Conference on Knowledge Representation and Reasoning (KR2002)*. San Francisco: Morgan Kaufmann Publishers.

Mitra, P.; Wiederhold, G.; and Kersten, M. A. 2000. Graph-Oriented Model for Articulation of Ontology Interdependencies. In *Advances in Database Technology (EDBT'2000): Seventh International Conference on Extending Database Technology*. Berlin: Springer-Verlag.

Noy, N. F. and Musen, M. A. 2002. Evaluating Ontology Mapping Tools: Requirements and Experience. In Workshop on Evaluation of Ontology Tools at EKAW'02 (EON2002). SMI Report 2002-0936, Stanford Univ., Stanford, CA.

Paolucci, M.; Ankolekar, A.; Srinivasan N.; and Sycara, K. 2003. The DAML-S Virtual Machine. In *The Semantic Web: Proceedings of the Second International Semantic Web Conference (ISWC)*, 290–305. Berlin: Springer-Verlag.

Paolucci, M.; Kawamura, T.; Payne, T.; and Sycara, K. 2002a. Semantic Matching of Web Services Capabilities. In *Proceedings of the First International Semantic Web Conference (ISWC)*, Lecture Notes in Computer Science. Berlin: Springer-Verlag.

Paolucci, M.; Kawamura, T.; Payne, T. R.; and Sycara, K. 2002b. Importing the Semantic Web in UDDI. In *Proceedings of Web Services, E-Business, and the Semantic Web: CAISE 2002 International Workshop*. Lecture Notes in Computer Science, 225–236. Berlin: Springer-Verlag.

Rahm, E.; and Bernstein, P. A. 2001. A Survey on Approaches to Automatic Schema Matching. *The International Journal on Very Large Data Bases* 10(4): 334–350.

Sabou, M.; Richards, D.; and van Splunter, S. 2003. An Experience Report on Using DAML-S. In *Proceedings of the Twelfth International World Wide Web Conference Workshop on E-Services and the Semantic Web (ESSW '03)*. New York: Association for Computing Machinery.

Schlenoff, C.; Gruninger M.; Tissot, F.; Valois, J.; Lubell, J.; and Lee, J. 2000. The Process Specification Language (PSL): Overview and Version 1.0 Specification. Technical Report NISTIR 6459, National Institute of Standards and Technology, Gaithersburg, MD.

Sheshagiri, M.; desJardins, M.; and Finin, T. A Planner for Composing Services Described in DAML-S. In *Proceedings of the AAMAS Workshop on Web Services and Agent-Based Engineering*. New York: Association for Computing Machinery.

Sirin, E.; Parsia, B.; and Hendler, J. 2004. Composition-Driven Filtering and Selection of Semantic Web Services. In *Semantic Web Services: Papers from the AAAI Spring Symposium*. Technical Report SS-04-06. Menlo Park, CA: American Association for Artificial Intelligence.

Uschold, M.; and Gruninger, M. 2002. Creating Semantically Integrated Communities on the World

Wide Web. Paper presented at the Semantic Web Workshop, Honolulu, 7 May 2002. ([semanticweb2002.aifb.uni-karlsruhe.de/USCHOLD-Hawaii-InvitedTalk2002.pdf](http://semanticweb2002.aifb.uni-karlsruhe.de/USCHOLD-Hawaii-InvitedTalk2002.pdf))

Weld, D. Recent Advances in AI Planning. *AI Magazine* 20(2): 93–123.

Wu, D.; Parsia, B.; Sirin, E.; Hendler, J.; and Nau, D. 2003. Automating DAML-S Web Services Composition Using SHOP2. In *Proceedings of the Second International Semantic Web Conference (ISWC2003)*. Lecture Notes in Computer Science. Berlin: Springer-Verlag.



**Mark H. Burstein** is a division scientist and the director of the Human Centered Systems Group at BBN Technologies. He is cochair of the Semantic Web Services Initiative's Architecture Committee. He holds a B.S. in mathematics from the Massachusetts Institute of Technology and a Ph.D. in computer science (artificial intelligence) from Yale University.

He is author or coauthor of more than 60 papers on topics including semantic web services, ontology translation, multiagent systems, mixed-initiative planning and scheduling, plausible reasoning, knowledge acquisition, and machine learning. He can be contacted at [burstein@bbn.com](mailto:burstein@bbn.com).



**Drew V. McDermott** is professor of computer science at Yale University. He was educated at the Massachusetts Institute of Technology, where he received a Ph.D. in 1976. His research is in planning, knowledge representation, and interagent communication, with side excursions into philosophy.

He helped create the AI Planning Systems conference, which, as a result of a merger, became the International Conference on Planning and Scheduling. He chaired the Rules Committee for the first International Planning Competition, which created PDDL, the standard notation for defining planning domains. He is the coauthor of two textbooks in artificial intelligence and the sole author of *Mind and Mechanism*, a book on AI and consciousness. He is on the editorial board of *Artificial Intelligence*, and is a Fellow of the American Association for Artificial Intelligence.