FIPA Ontology TC: *Use Cases for Ontologies*

ontology@fipa.org
19th December, 2001

1. Use Cases for Application of Ontologies in Service Clustering

A consumer would like to make arrangements for a trip using a B2C 'trip arranger' service. By using such a service, the consumer is relieved of the burden of having to acquire reservations for one or more modes of transportation, hotel, tickets for special events, etc., in a piecemeal fashion. The customer provides a customer profile that indicates what general arrangements are desired. This includes but is not limited to information such as times of travel, the cost restrictions, and entertainment preferences.

- The 'trip arranger' service is a cluster of services. The 'trip arranger' would sequence through the various services in a manner consistent with its business process design. A business process ontology in the travel service domain could be used to help interpret the customer request. This is not to say that some of the services in the cluster could not be 'in process' concurrently. The actual sequence need not be the same for different customer requests.
    - If a customer wants to visit New York solely to see a theater production, the 'trip arranger' process could use the business process ontology in the travel service domain to attempt to arrange for theater tickets before travel and lodging arrangements are attempted.
    - A customer service request form design could be informed by the business process ontology, travel service domain.

- The 'trip arranger' service utilizes multiple ontologies in interpreting the customer request, including any constraints the customer has. For example, a Transportation Ontology could be used in determining what transport modes are viable for fulfilling the customer need. Ontological information would be used in the following instances:
    - A Washington DC to San Francisco trip is planned to take place over 4 days. Train transportation is not a viable option because of time constraints. Reasoning on temporal features would have to be done, and this activity could make use of a Time Ontology. (It is possible that the 'trip arranger' service uses a Transportation Ontology that makes references to elements of a Time Ontology.)
    - A customer may have requested 'jet service' as a mode of transportation. The 'trip arranger' service makes use a Transportation Ontology in selecting among various airline flights. A flight that is to be made on a propeller plane and fits time and cost constraints is rejected in favor of a flight whose aircraft equipment is a jet.

- This overlaps with Process Ontology: Each service in the 'trip arranger' cluster of services will have some functional task. For example, consider the 'hotel reservation' service. At some point, this service will want to 'make reservation'. In executing the 'make reservation' task, the 'hotel reservation' service will have to utilize the process ontology that encompasses the task of making reservations. At stages of the 'make reservation' task, i.e., in a sub-process, relevant ontologies will be needed. For example, a sub-process could be 'check credit card information', and an ontology or ontologies that encompassed different currencies and credit cards would be needed.

2. Semantic Framework Mapping

2.1 Upper Ontologies

There exist several types of upper ontologies that may be of use in these areas. Higher-level ontologies may be used for:

- Policies such as security policies, conversation policies, etc.
- Conventions such as legal, societal, or ethical conventions.
- Contracts that can be formed between agents.

These ontologies provide knowledge about the structure and content of policies, contracts, conventions; also the meaning of each of the components, at a general level. This includes:

- The terms and basic building blocks used within the policies and contracts. For instance, an upper ontology for policies may contain structural terms, legal terms, societal terms, and other generic terms germane to policies as a whole.
- Axioms defining how these basic building blocks can be composed within a given convention, policy or contract.
- The semantics to handle these terms and reason over them.
- Interrelationships among different upper ontologies, the nature and direction of those relationships, and the relationships of the upper ontologies to the domain-specific ontologies. For instance, policies exist in the context of some conventions and legal system.

Note that one property of these upper ontologies is that a given, say, policy, will not be fully instantiated from the upper ontologies, but rather be instantiated generally as the policy is defined and its use unfolds.

2.2 Domain-Specific Ontologies

These lower-level ontologies are more domain-specific, for example, defining policies in a specific area. A given domain-specific ontology exists in relationship to some specific upper-level ontologies, and those relationships must be represented. Also, a policy represented using some set of upper ontologies may be instantiated differently in different domains and different uses.

These lower levels can have multiple models and representations. The ontologies themselves must allow the ability to reason over terms from a set of ontologies, which is the collection of ontologies that the given instance happened to use. Note that the definition of the domain ontologies should not directly represent the upper ontology terms and axioms, as this hinders reusability and adaptability.

Also, instantiation of policies, conventions and contracts may occur over a period of time. As with the issue of multiple ontologies, reasoning over such policies, etc, must be able to adapt to the fact that they may not yet be fully instantiated.

Use Case

1. Agents A, B and C know about ontologies O1 and O2.
2. Data set D1 provides information about the topography of a given region, R. It has information about the slope and aspect of the region by location. The ontology for D1 is O1, and it associates location with x,y points that refer to coordinate system C1.
3. Data set D2 provides information about the soil types by location in region R. The ontology for D2 is O2, which associates location with x,y points that refer to coordinate system C2.

4. A wants to know about areas in Region R that are good for wine cultivation. A asks B for all the places in R that have a particular aspect and soil type that are known to be suitable for growing grapes of the suitable type.

5. B can satisfy this query by accessing data sets D1 and D2, but must match locations that are described according to two different ontologies, O1 and O2.

6. B repeatedly asks C to translate locations in the O1 ontology into corresponding locations in the O2 ontology in order to match locations that have the appropriate soil type and aspect.

3. Manufacturing

In a manufacturing environment, there are several ontologies within a given process that address safety, content-use, device functionality, restrictions, and constraint theories to what the different levels of workers can perform. All of these have different aspects of ontologies based on equipment status and the type of process being executed.

From these aspects, one understands what the tasks are and who they communicate with in order to accomplish a given task.

There is overlap. Much of the resolution is solved by humans. Agent systems will be used to enable the tracking and communication of results for streamlining tasks.

4. Modeling Externally-Defined References to Objects

FIPA agents may need to converse about or interact with entities having an identifier or reference defined according to some naming scheme that is outside the scope of FIPA. For example, CORBA objects may be referenced using interoperable object references (IORs), World Wide Web resources may be named (and usually accessed) using URIs, and books may be referenced using ISBNs. In some cases these references are unique, but in other cases (such as CORBA IORs) uniqueness may not be required.

FIPA ACL provides the query-ref communicative act for asking another agent for a "reference" to an object corresponding to a given description. In the current FIPA ACL and SL semantics, an "object" is identified by a logical constant considered to be its standard and unique name. This notion of object identification is not powerful enough to account for the range of object reference types that an agent may need to model. Therefore, such "external" references for an object must be modeled explicitly within an agent and ontologies must indicate when a particular class corresponds to a type of "resource" (to use the Semantic Web term) having references of a particular type. For any ontology representation language to be used with FIPA agents, a standard mechanism should be defined for representing this information about classes, and for encoding the associations of instances with references. This would assist in the portability of this information between ontology representation languages.

Use Case

1. Ontology: A Person has a name and can be referenced by a social security number. A Person also optionally has a Web page which is referenced by a URL.

2. Instance data: There is a Person named John Smith referenced by the social security number 1234567. His Web page is referenced by the URL:
   http://www.freewebpages.com/JohnSmith.html

3. An agent sends a query-ref message with the description meaning "the Web page of the Person named John Smith". The response is the URL:
   http://www.freewebpages.com/JohnSmith.html

5. Information Agents

5.1 Information Integration and Fusion

In an information agent system, one of the ways that ontologies are used is to represent knowledge about the data that is being brought in from the various information sources. Information agents may use the ontology like a schema to manipulate the information in a database-like manner, for example, via SQL. This in turn puts some conditions on what might be in an ontology, for instance:

- The ontology must cover the elements of each information source schema that the agent system needs to take advantage of, though it does not imply that all concepts from all information sources be represented in the ontology.

- The portion of the ontology that concerns the data directly needs to have a specification that is amenable to manipulation via relational operations, to wit, that the concepts be represented as entities, relationships and their attributes.

- When related information from different sources needs to be joined using a relational join operation, then the attributes specified in the join must be represented in the same manner. (Note that this is not a requirement in any other situation; there is nothing a priori in information manipulation that forces other attributes to have the same representation.). Typically these are attributes that are viewed from a query processing perspective in a manner similar to keys. In this situation, it may be desirable to have the ontology include knowledge of a canonical representation for those attributes and/or knowledge in the form of axioms as to how alternative representations relate to the canonical representation.

5.2 Information Processes and Information Analysis

Information agents not only access and integrate information, but also implement processes required to analyze, abstract, and extract knowledge from this integrated information. These processes involve composing services into flows; there needs to be some form of Process Ontology that expresses the knowledge about the flows and about how to access the services.

Invoking each service may involve expressing what is desired in terms of some specific Service Ontology that represents knowledge about the type of service and the knowledge needed to invoke the service. A process should be able to refer to service ontologies describing the services accessed by the process, without the process ontology necessarily understanding a priori all the services that any process can invoke. Process Ontologies need to incorporate the relationships necessary to enable the late binding of service steps to ontologies at process definition time.

Invocation of generic information analysis services may require that the service step in the process also be given the (location of the) data and the necessary knowledge about the concepts and organization of the data, expressed in terms of some Information Ontology (as described in the previous section). Again, this means that process ontologies need to incorporate the relationships necessary to enable the late binding of service steps to information ontologies at process definition time.

Use Case 1: Conversation

Conversation is the notion of multi-party, sequential communication between two or more agents. The unit component of a conversation is the message, which contains some propositional cont ent (performative) expressed in the context of one or more ontologies. This implies:

1. Ontology can provide a common context for a conversation.

2. More than one ontology can be reference within a conversation.

3. The scope of ontology within a conversation can be:
   - Individual elements (terms, predicates, etc.) of content expressions
   - Complete content performative of a message
   - A sequence of messages within a conversation
   - The complete conversation

4. Agents may require a policy guarantees that an ontology bound to a conversation at initiation, remain consistent and available for the duration of that conversation of a specified fragment thereof.

5. The ontology will affect the observable semantics of a conversation.

Use Case 2: Shared Plans

A shared plan implies the formation of a group (or team) of agents to achieve a specific goal. Typically the plan will break down the goal into a set of sub-goals that are distributed amongst the agent team. Each sub-goal consists of set of tasks, defined by the plan. This implies:

1. Ontology can provide a common context for all participating agents.

2. More than one ontology may be relevant to aspects of the plan.

3. The scope of an ontology in a shared plan can be:
   - Individual tasks assigned to an agent
   - Sub-goals assigned to one or more agents
   - The entire shared plan (goal) and all agents participating in it

4. Agents may require a policy guarantees that an ontology used remain consistent and available for the duration of the shared context.

Use Case 2: Mutable Ontologies

The stability of ontologies with respect to change over time is a critical issue for agent populations. For instance, the validity of policies and contracts established between agents may rely on consistency in referred ontologies; if an ontology were to change, even in the smallest sense, this could easily invalidate any dependent agreements.

This implies that ontology providers should ideally publish any guarantee of ontology persistency, whether permanent or leased. Without any such guarantee, agents should account for potential instabilities when establishing agreements.

6. A Use Case for a New Member Joining an Organization or Domain

- Service examples (provider view):
  - Credential creation
  - Credential registration
  - Credential exchange
  - Credential revocation

- Service examples (user view):
  - Identifying and authenticating a person
  - Selecting and configuring use of an end-to-end secure channel (policy driven)
  - Protecting personal information
  - Gaining the use of and registering credentials

- Drivers

  Currently the use of security services by applications is hard-coded, predefined, tied to a particular system infrastructure and tied to a particular domain. Example security ontologies:
  - Credential types: X.509 certificates, Driver's License, etc.
  - Credential distribution: PKI such as XKMS, SPKI, Web of Trust
  - Support services: security policies, profiles, negotiation dialogues
  - Process driven: identifying when credential are challenged and when they must be shown or used during the process ontologies, checking credentials
  - Domain specific: SET, etc.

  This creates issues in setting up end-to-end or agreeing security between parties across multiple domains, across heterogeneous service infrastructure and within an open (open in sense that new parties can be added or removed at any stage service) service infrastructure.

Use Cases: A Person Joins a New Organization

When a user joins a new organization, the user needs to gain different types of credentials to access resources within the organization. For example, a type of digital certificate may be used to send secure email, a library card may be needed to access the library service, an access card may be used to access different buildings, network login credentials etc. Each of these credentials will have its own ontology to create, define, register, exchange and revoke credentials. Generally, the identities and addresses of different "first-contact" representatives within parts of the organization will be given to the new user and the user will then contact them to gain the various credentials to operate within the new organization. Some people within the organization need to have knowledge of how to acquire different sets of credentials templates and how to assist new-comers to complete the registration for the credentials. The organization member will need to know how to use the credentials, e.g., when a challenge may be issued and the user must be prepared to present or use the credentials. Parts of this usage process for certain credentials may be explicit, pre-defined by policies within the organization whilst other processes may be more dynamic and react to changes in the wider environment of the organization, for example, if attacks occur in the wider environment, the processes for using the ontologies may vary from time to time.

There are several cross-domain issues here such as maintaining links and credentials with multiple organizations: links (and credentials) may be maintained with previous organizations.