Browsing Schedules - An Agent-based approach to navigating the Semantic Web

Terry R. Payne, Rahul Singh and Katia Sycara

Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213
{terryp, kingtiny, katia}@cs.cmu.edu

Abstract. The Semantic Web promises to change the way agents navigate, harvest and utilize information on the internet. By providing a structured, distributed representation for expressing concepts and relationships defined by multiple ontologies, it is now possible for agents to *read* and *reason* about published knowledge, without the need for scrapers, information agents, and centralized ontologies. Agents can utilize this knowledge to seek and invoke other agents and web services, thus supporting navigation across the Semantic Web. We demonstrate how agents support enhanced navigation on the Semantic Web within a conference-schedule domain, and present three agent-based services: the RETSINA Calendar Agent, which reasons about schedules marked up on the Semantic Web; the DMA2ICal markup tranaslation agent which provides translation services between schedules grounded in different ontologies, and a Conference Agent, that invokes the Calendar Agent.

1 Introduction

The World Wide Web was originally designed as a distributed information space that seamlessly supported human navigation through related, linked documents. Since this medium was originally designed to support human-to-human communication [2], machine or agent mediated assistance has been hindered by the type of markup used (HTML). An emphasis by content providers on presentation and physical design has resulted in a lack of structure, at the content level, and rendered documents opaque to machine comprehension.

The Semantic Web [2] goes beyond the World Wide Web by using a structured, logically connected representation to encode knowledge. It also provides sets of inference rules that can be used to reason over this knowledge. Since the Semantic Web does not require content providers to use a single centralized ontology; but instead supports the use of many different ontologies, it is unrealistic to assume that agents and services will understand all possible markup. Hence many different agents providing specific services will translate/convert information from one ontology to another as required, thus encapsulating and providing specific functionality to the user. By communicating and exchanging information, agents will present the user with a new agent-oriented approach to navigating the semantic web. This transition from using a homogeneous markup for layout on the Web to hetrogenous ontologies for semantic markup will allow agents to seamlessly interoperate, and provide easier, faster, and more flexible access to relevant data.

In this paper, we present three services that share information about conference schedules, and thus provide functionality to the user beyond that available with current web browsers. Section 2 describes a simple *Conference Agent* that presents different conference sessions to the user. Once the user selects the session or sessions of interest, the *Conference Agent* invokes the *Calendar Agent*, described in Section 3, which offers schedule browsing and download functionality. If the *Calendar Agent* fails to understand the markup used, it will utilise a service discovery mechanism to locate a markup translation service, such as the *DMA2ICal* service described in Section 4. This process is discussed before concluding the paper in Section 5.

2 A Simple Conference Agent

The *Conference Agent* is a service designed to provide basic information about a conference, and provide different selection criteria for browsing the conference agenda. It organises the presentation of annotated conference details, such as location, date, hotels, etc, and publishes this content as HTML with additional controls to invoke other Semantic Web agents. One such functionality is schedule browsing, which is encapsulated by the *RETSINA Calendar Agent*.

Once the user has selected the subject areas of interest, the *conference agent* locates the users *calendar agent* and sends it a browse request containing resource references to events corresponding to the subject areas of interest. These events are then presented to the user for viewing. By tasking the Conference agent in this way, users do not need to visually inspcet all the individual events, but can rapidly compare those of interest with other events stored in the user's calendar.

3 RCal - RETSINA Calendar Agent

The *RETSINA Calendar Agent (RCal)* is a distributed meeting scheduling agent that can navigate semantic web content to gather and reason about events and schedules. It works synergistically with a commercial Personal Information Manager (PIM) to store schedules and contact details with those in the user's calendar and address book. Schedules found on the Semantic Web can be browsed via the *schedule browser* (Fig. 1), which lists the events, times, location and attendees of each event. However, the use of semantics can be used to facilitate further browsing or to invoke additional services.

Schedules may be viewed in one of several ways; via the Semantic Web schedule browser, or automatically based on instructions received from trusted agents via an Agent Communication Language (ACL) message. The user can also specify the URI of a schedule (e.g. a conference program) in the semantic web schedule browser. Alternatively, the agent may receive a browse request from

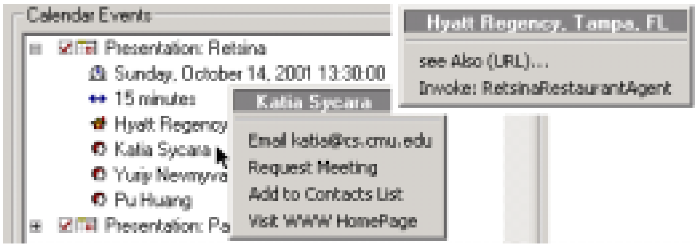

Fig. 1. Browsing schedules & invoking context-based services/agents.

another agent, in which case the schedule browser would appear automatically. The schedule and relevant resources (e.g. information about locations, attendees etc) are then gathered from the Web and presented to the user via the browser.

As concepts are often referenced by a resource URI, they may also include other information that may be relevant to the user. For example, an *ATTENDEE* resource may contain information such as contact details including an email address, the name of their *RCal* agent (used for distributed meeting scheduling [5]), and a WWW Homepage. If these properties are known, then additional services can be offered to the user when a concept is selected (e.g. the user right-clicks the *ATTENDEE* concept *"Katia Sycara"* in Fig. 1). These properties can also be used to query service providers (i.e. other agents) via a discovery infrastructure such as a DAML-S Matchmaker [1]. This form of serendipitous service discovery (as opposed to goal-directed service discovery) attempts to find services that might be of use to the user. *RCal* constructs requests for services based on the properties of the selected concept, and returns a URI of a web page with information that can then be presented to the user. For example, if the location of an event is selected, such as the *Hyatt Regency Hotel* in Fig. 1, then the properties of the location resource will be examined. If they include details about the location's address or latitude and longitude, a request will be automatically constructed and submitted to a middle-agent, which returns the advertisements for the matching services, such as the *RetsinaRestaurantAgent* (Fig. 1). These services can also be offered to the user; if the user selects a service, a query can then be sent to the selected service and the results displayed.

4 DMA2ICal - Agenda Markup Translation Service

Whilst *RCal* can provide browsing and download functionality for schedules marked up using the ICal ontology, it is unable to understand markup using other ontologies, such as that used by *ITTalks*[1] or the DAML Markup Agenda (DMA) Ontology[2]. If the ontology used to mark up a new schedule can be mapped to, or merged with the ICal ontology, then the agent will be able to

[1] See *http://daml.umbc.edu/ontologies/calendar-ont*
[2] See *http://www.daml.org/2001/10/agenda/*

reason about the new markup and understand what concepts are equivalent to the known events, schedules etc. The DMA2ICal agent is a simple translation service that translates agendas that are marked up using the DMA ontology into markup using the ICal ontology. It was designed not only to translate markup between two services, but to demonstrate how such services could be dynamically located and tasked when unknown markup is detected.

```
<dma:Meeting rdf:ID="TAC01">
  <dma:name>Trading Agent Competition 2001 Workshop</dma:name>
  <dma:location resource="#HRTampa" />
  <dma:day>
    <dma:Day>
      <dma:start>2001-10-14T13:45:00</dma:start>
      <items rdf:parseType="daml:collection">
        <dma:Talk resource="http://www.tac.org/2001event.rdf#PainInNEC"/>
        <dma:Talk ID="RetsinaTrading">
          <dma:title>Presentation: Retsina</dma:title>
          <dma:speaker resource="http://www.daml.ri.cmu.edu/people.rdf#ks" />
          <dma:duration>PT15M</dma:duration>
        </dma:Talk>
```

```
<ical:VCALENDAR ID="TAC01">
  <dc:title>Trading Agent Competition 2001 Workshop</dc:title>
  <ical:VEVENT-PROP resource="http://www.tac.org/2001event.rdf#PainInNEC"/>
  <ical:VEVENT-PROP>
    <ical:VEVENT ID="RetsinaTrading">
      <ical:DTSTART>
        <ical:DATE-TIME><value>20011014T134500</value></ical:DATE-TIME>
      </ical:DTSTART>
      <!-- end not included in this example -->
      <ical:LOCATION resource="#HRTampa" />
      <ical:ATTENDEE resource="http://www.daml.ri.cmu.edu/people.rdf#ks" />
     <ical:DESCRIPTION>Presentation: Retsina</ical:DESCRIPTION>
    </ical:VEVENT>
```

Fig. 2. Translating markup from DMA to ICal.

The DMA2ICal service builds a model based on the DMA ontology and the markup to be translated, and then transforms this model into markup using the ICal ontology. When active, the translation agent advertises its capabilities using the DAML-S capability description language with a DAML-S Discovery Service [1]. Other agents, such as *RCal*, can then request translation services by submitting a query with the URI of the agenda to be translated. This agenda is parsed, and an RDF graph based on the markup is constructed. This graph is then traversed to identify the values of the properties within the agenda (such as time, location, speaker) and these values are annotated with ICal markup. This new markup is then returned to the requesting agent.

Whilst this simple approach can be effective if the markup to be translated conforms to a fixed number of ontologies, it lacks the ability to convert specific resources within the markup, such as different ontologies used to annotate the speaker. Also, it only provides a one-to-one mapping between the markup using two ontologies, but not a more generic solution that can be more easily applied to markup using other ontologies. Various schemes (such as Onion [4]) identify the correspondence between different ontologies, and can be used to define domain dependent rules for markup translation. However, the simple approach used by DAM2Cal was sufficient for an initial investigation of autonomous discovery and invocation of translation services for agenda markup.

5 Discussion

This paper demonstrates how service discovery and information sharing can allow agent communities to locate and present relevant services to a user, based on the information that is being browsed. Although several different ontologies may be used to markup content, translation services can transform unknown markup into that which can be understood by the agent. *RCal* makes use of a serendipitous search to look for services that may be of use to the user, based on selected resources. However, in a service rich environment, many irrelevant services may be presented to the user. Thus, work is currently underway to develop profiles of the user's interest, and to infer context (such as locating restaurants in favor of hardware stores when examining the location of a conference site). The *ITTalks* Agent system [6] is an existing web-based system that provides automated, intelligent notification of information techology seminars. Profiles of user preferences, annotated in DAML+OIL [3], are used to suggest those seminars that might be of interest to the user. These examples demonstrate how, by combining the semantics now available through the Semantic Web, communities of agents can interoperate and work synergistically to provide better access to information and functionality than was previously available on the World Wide Web.

6 Acknowledgments

The research was funded by the Defense Advanced Research Projects Agency as part of the DARPA Agent Markup Language (DAML) program under Air Force Research Laboratory contract F30601-00-2-0592 to Carnegie Mellon University. Special thanks goes to L. Miller & S. Decker for proposing the original Semantic Web Calendar Challenge, and to M.Dean for developing the original gen2html routine for generating HTML from DAML Meeting Agenda markup.

References

1. A. Ankolekar et. al. DAML-S: Semantic markup for web services. In *International Semantic Web Working Symposium*, pages 411–430, 2001.
2. T. Berners-Lee, J. Hendler, and O. Lassila. The Semantic Web. *Scientific American*, 284(5):34–43, 2001.
3. J. Hendler and D. L. McGuinness. Darpa agent markup language. *IEEE Intelligent Systems*, 15(6):72–73, 2001.
4. P. Mitra, G. Wiederhold, and S. Decker. A scalable framework for the interoperation of information sources. In *International Semantic Web Working Symposium*, pages 317–329, 2001.
5. T. R. Payne, R. Singh, and K. Sycara. Rcal: A case study on semantic web agents. In *First International Conference on Autonomous Agents and Multi-Agent Systems*, 2002.
6. R.Scott Cost et. al. ITTalks: A case student in how the semantic web helps. In *International Semantic Web Working Symposium*, pages 477–494, 2001.