

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 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 Translation 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. Although this medium was originally designed to do more than simply support human-to-human communication [2], an emphasis 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 supports the use of many different ontologies, one cannot assume that agents will understand all possible markup. However, translation services that convert markup from one ontology to another may be solicited by an agent when it encounters unknown concepts. By communicating and exchanging information, agents can 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 heterogeneous 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 describe agents and services that share information about conference schedules, and thus provide functionality to the user beyond that available with current web browsers. Section 2 describes the *Retsina Calendar Agent* which provides browsing functionality for Semantic Web schedules. This agent can be invoked remotely from the *Conference Agent* when the user indicates that they would like to browse events pertaining to a specific topic. However, if the ontologies used by the schedule are unknown to the *Calendar Agent*, it contacts a service discovery mechanism to locate a domain-specific markup translation service, such as the *DMA2ICal* service. This process is discussed before concluding the paper in Section 3.

2 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 calendar events and contact details. It works synergistically with a commercial Personal Information Manager (PIM) to contrast and store schedules such as conference programs, or recurring appointments with existing events within the user's calendar. Published Semantic Web events that are stored can be monitored by the agent, and updated if necessary. In addition, notifications can be associated with events and sent via email or via another agent to a mobile device to remind the user when each talk is about to start.

RCal understands several different ontologies that may be used to generate schedule markup. The schedules and events themselves are described using the Hybrid RDF Calendar Ontology (iCal)¹. User resources and contact details are described using the Friend-of-a-Friend ontology (<http://xmlns.com/foaf/0.1/>), whereas the Dublin Core ontology (<http://dublincore.org>) is used to provide meta-data about the schedule.

Schedules found on the Semantic Web can be browsed via the *Schedule Browser* (Fig. 1), and compared with existing appointments to identify conflicting events. The browser lists the events, times, location and attendees of each event. The user can then select all, or a subset of the events and import them into the PIM for perusal at a later date. Notifications can also be associated with the imported events to remind the user of impending meetings. The contact details of the people attending the meetings can also be imported into the user's contact list. As the URI referring to each event is stored with the information about the event, *RCal* can periodically check to see if the event has changed, and thus notify the user of the change. In addition, further information can be explored, such as visiting relevant web pages, or invoking services that provide driving instructions or local restaurant recommendations.

The ability to refer to resources defined in different ontologies facilitates the navigation of information not directly related to a schedule. For example, when listing the events within the Semantic Web schedule, *RCal* locates the

¹ See <http://ilrt.org/discovery/2001/06/schemas/ical-full/hybrid.rdf>




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

Name property of an `<ical:ATTENDEE>` resource to list the name of each meeting attendee. Other information may be defined by this resource, such as email or webpage properties, that could facilitate additional browsing or offer additional services (e.g. the user right-clicks the 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, 5]. This form of serendipitous service discovery attempts to find services that might be of use 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. A request for service will be constructed containing details about the location’s address or latitude and longitude, and submitted to a Matchmaker. Advertised service descriptions that match the service request are then returned and offered to the user as additional services. In the *Hyatt* example in Fig. 1, *RCal* discovers the existence of the *RetsinaRestaurantAgent*, and presents this via a drop-down menu option.

2.1 Agent-based Browsing - Invoking the RCal Schedule Browser

As agents reason about Semantic Web markup, they will encounter concepts unrelated to their domain of expertise. By identifying these concepts, and searching for other local agents that provide browsing capabilities, an agent can share the responsibility for browsing this data. This mechanism can be demonstrated by using *RCal*’s schedule browser via a *Conference Agent*, which was designed for browsing conference information marked up on the Semantic Web. This agent organizes 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. A user who is browsing a conference site can select the subject areas of interest and browse a schedule containing only presentations on the selected subject areas. The *Conference Agent* locates the user’s *Calendar Agent* by querying a service registry for agents that: (1) supports schedule browsing; (2) belong to the current user; and (2) are capable of providing browsing facilities on the same display as that used by the *Conference Agent*. Once an agent is found, a KQML-based browse request is sent to this agent (i.e. *RCal*) containing resource references to events corresponding to the subject areas of interest. The *RCal* schedule browser then appears on the

user's display, listing the subset of events. By tasking the *Conference Agent* in this way, users do not need to visually inspect all the individual events, but can utilize *RCal* to rapidly compare events of interest with others stored in the user's calendar.

2.2 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* or the DAML Markup Agenda (DMA) Ontology. To comprehend this markup, an agent would need additional rules (e.g. articulation rules [4]) or the use of translation services. *RCal* overcomes this by locating and soliciting the help of translation services that convert markup from one ontology to another. These services are located using a discovery service based upon the DAML Services Ontology (DAML-S) [1], which describes a service in terms of its capabilities (through an advertised *Service-Profile*), how the service works (through a *Service-Model*) and how it can be invoked (through a *Service-Grounding*).

The *DMA2iCal* agent is a simple domain-specific translation service that translates agendas that are marked up using the DMA ontology² into markup using the iCal ontology. It was designed not only to translate markup between two ontologies, but to demonstrate how such services could be dynamically located and tasked when unknown markup is detected. When this service is made available on the Semantic Web, it advertises its capabilities with a service registry, such as the DAML-S Matchmaker [5]. When *RCal* detects an unknown concept within the markup (for example, the `<dma:Meeting>` resource defined in Fig. 2), it constructs a DAML-S request for the service based on this concept. This request simply consists of the unknown concept, and a desired `<ical:VCALENDAR>` concept. This is then submitted to the DAML-S Matchmaker, which performs a semantic capability match between the request and advertised Service Profiles, before returning a list of possible matching services, including the profile for the *DMA2iCal* translation service. *RCal* then selects and invokes one of these returned services by sending it the URI of the unknown concept. The service then constructs new markup using the iCal ontology, and returns this markup to *RCal* which then presents the schedule to the user. This is the same mechanism used by *RCal* to discover browsing and other services when the user selects concepts from the schedule browser (Fig. 1).

3 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

² See <http://www.daml.org/2001/10/agenda/>

```

<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.ni.cmu.edu/people.rdf#ks" />
          <dma:duration>PT15M</dma:duration>
        </dma:Talk>
      </items>
    </dma:Day>
  </dma:day>
</dma:Meeting>

<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.ni.cmu.edu/people.rdf#ks" />
      <ical:DESCRIPTION>Presentation: Retsina</ical:DESCRIPTION>
    </ical:VEVENT>
  </ical:VEVENT-PROP>
</ical:VCALENDAR>

```

Fig. 2. Translating markup from DMA to ICal.

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 technology 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 WWW.

4 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.

References

1. A. Ankolekar et. al. DAML-S: Semantic markup for web services. In *Int. 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 *Semantic Web Working Symposium*, pages 317–329, 2001.
5. M. Paolucci, T. Kawamura, T. Payne, and K. Sycara. Semantic Matching of Web Services Capabilities. In *First International Semantic Web Conference*, 2002.
6. R.Scott Cost et. al. ITTalks: A case student in how the semantic web helps. In *Int. Semantic Web Working Symposium*, pages 477–494, 2001.