

RCal: A Case Study on Semantic Web Agents

Terry R. Payne
Carnegie Mellon University
The Robotics Institute
5000 Forbes Avenue
Pittsburgh PA 15213, USA
terryp@cs.cmu.edu

Rahul Singh
Carnegie Mellon University
The Robotics Institute
5000 Forbes Avenue
Pittsburgh PA 15213, USA
kingtiny@cs.cmu.edu

Katia Sycara
Carnegie Mellon University
The Robotics Institute
5000 Forbes Avenue
Pittsburgh PA 15213, USA
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. We present the *RETSINA Calendar Agent*, a distributed meeting scheduler, that reads schedules (such as conference programs, events, etc) marked up in RDF on the Semantic Web, and imports these into the user's Personal Information Manager. The embedded Semantic Web Browsing tool allows the user to explore related concepts within the schedule, and to query other agents and service providers for more information.

Categories and Subject Descriptors

I.2.1 [Artificial Intelligence]: Applications and Expert Systems—*Office Automation*

General Terms

Human Factors

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], machine or agent mediated assistance has been hindered by the type of markup used within the documents. An emphasis by content providers on presentation and physical design has resulted in a lack of structure, both at the layout and content levels, and rendered most documents opaque to machine comprehension. The Semantic Web [2] goes beyond the World Wide Web by encoding knowledge using a structured, logically connected representation, and providing sets of inference rules that can be used to conduct automated

reasoning. Whilst the idea of knowledge representation languages is not new, existing languages generally use their own set of ontologies and inference rules to identify and eliminate logical contradictions and inconsistencies. It would be impossible to enforce the use of managed, centralized ontologies and rules within a dynamic and distributed knowledge base such as the World Wide Web without sacrificing expressivity. For this reason, the Semantic Web accepts these contradictions, and does not enforce the use of a single centralized ontology. Instead, many different ontologies are used that encode knowledge, such as the Dublin Core and RSS¹ ontologies. In addition, new ontologies can be defined, existing ontologies can be extended, and links can be added to connect different ontologies and reuse concepts. In this paper, we describe the *RETSINA Calendar Agent*, a distributed meeting scheduling agent, that has been augmented to understand and reason about scheduling information on the Semantic Web, and we demonstrate how additional related agents and services can be discovered and presented to the user, thus enhancing the agent's assistance.

2. RDF & THE SEMANTIC WEB

The Semantic Web expresses concepts and their properties using the Resource Description Framework (RDF) [3], which is in turn expressed in XML. RDF encodes concepts and their relationships as sets of triples, where each triple represents a subject, predicate and an object. For example, the statement "*Harrisburg* is the *State-Capital* of *Pennsylvania*", is a triple containing the subject *Harrisburg*, object *Pennsylvania*, and the predicate *State-Capital* which relates the two concepts. Both subject and object can, in turn, be related to other concepts using additional properties (or predicates), forming a directed graph. RDF extends this natural structure by allowing concepts to be represented by URIs to other concepts. Thus, concepts are no longer terms bound to a single node, but unique definitions that can be shared by multiple documents. For example, one could refer to the author of a paper by their name, but as that name is unlikely to be unique, it becomes problematic to reason about what other papers have been written by the same author. However, the author could be represented by a URI referring to the concept that uniquely defines that author, and thus avoid ambiguities when reasoning. In addition, the author concept may contain properties to other concepts, thus providing a richer environment for reasoning.

Concepts are defined by class-property descriptions within one or more distributed, extendible ontologies. Concept def-

¹See <http://dublincore.org/> & <http://purl.org/rss/>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'02, July 15-19, 2002, Bologna, Italy.

Copyright 2002 ACM 1-58113-480-0/02/0007 ...\$5.00.

initions themselves may be related to other concept descriptions, through equivalence and class/ sub-class hierarchies which facilitates semantic reasoning through subsumption inferencing. Thus, it is possible for two agents to inter-operate semantically, by reasoning about the relationship between two concepts in an ontology.

3. RCAL - RETSINA CALENDAR AGENT

Representing knowledge within RDF marks a shift from unstructured data on the web, to structured knowledge that can be harvested and utilized by agents. 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. For an agent to act as a useful meeting scheduling assistant, it should be able to automatically schedule meetings that are most convenient for its user, without continually requesting additional information. To achieve this, the agent should maintain an up-to-date model of the user's current activities. RCAL alleviates this dependency by automating the acquisition of schedules from the web and other heterogeneous agents.

```
<ical:VCALENDAR="TAC01">
  <dc:title>TradingAgentCompetitionWorkshop</dc:title>
  <ical:VEVENTPROP resource="http://www.tac.org/2001event/df#PainInNEC"/>
  <ical:VEVENTPROP>
    <ical:VEVENTPROP>RetsinaTrading
    <ical:VEVENTPROP>
      <ical:START>
        <ical:TIME><value>20011014T133000</value></ical:TIME>
      </ical:START>
      <!--end not included in this example-->
      <ical:LOCATION resource="HRTampa"/>
      <ical:ATTENDEE resource="http://www.damlrimeu.edu/people/df#s"/>
      <ical:ATTENDEE resource="http://www.damlrimeu.edu/people/df#n"/>
      <ical:DESCRIPTION>Presentation at Retsina</ical:DESCRIPTION>
    </ical:VEVENTPROP>
  </ical:VEVENTPROP>
</ical:VCALENDAR>
```

Figure 1: Schedule on the Semantic Web.

RCAL works synergistically with a commercial Personal Information Manager (PIM). It retrieves appointments and contact details from the PIM, and uses these to reason about available meeting slots. It can also negotiate with other RCAL agents to determine mutually available time slots for meeting requests. This approach, based on the Contract Net Protocol [4], solicits meeting times from each of the agents involved, and evaluates the responding bids to determine a suitable meeting time, which is then propagated back to the other agents. Schedules and contact details found on the Semantic Web can also be imported into the PIM via the semantic web schedule browser (e.g. the user specifies a URL of a conference program, and then selects specific talks to attend at a conference), or by automatically obtaining and importing new schedules shared by a community of agents (e.g. receiving seminar notifications, etc.).

Concepts defined by the iCal ontology are used to mark up the schedule illustrated in Fig. 1. The Dublin Core ontology is also used to provide meta-data about the schedule, such as *title*, *description*, *author* etc. The above schedule illustrates how information can be distributed across different documents, and how concepts can be reused. It contains two events (via the property VEVENT-PROP). The first event is referenced by the resource "<http://www.tac.org>

[/2001event.rdf#PainInNEC](#)", whereas the second, "*RetsinaTrading*", is defined within the document itself. In addition, the location of the event (i.e. the *Hyatt Regency, Tampa, FL*) is also defined elsewhere.




Figure 2: Browsing schedules & invoking context-based services/agents.

As concepts are often referenced by a resource URI, they may also include other information that may be of use to the user. For example, one of the ATTENDEE concepts in Fig. 1 may contain more than a first and last name, such as contact details including an email address, RCAL agent-name, and a WWW Homepage. If these properties are known to RCAL, or are equivalent to known properties (via shared ontologies), then additional services are offered to the user when the user selects a concept (e.g. the user right-clicks the ATTENDEE concept "*Katia Sycara*" in Fig. 2). 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 any service that might be of use to the user. RCAL constructs requests for services based on the properties of the selected concept, and returns a URL of a web page that can then be presented to the user. For example, the LOCATION property in Fig. 1 refers to a GEO concept, *HRTampa*, which includes properties describing the latitude and longitude of the *Hyatt Regency Hotel*. These requests can then be submitted to a middle-agent, which returns the advertisements for the matching services, such as the *RetsinaRestaurantAgent* (Fig. 2). 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. 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.

5. REFERENCES

- [1] A. Ankolekar et. al. DAML-S: Semantic markup for web services. In *International Semantic Web Working Symposium*, 2001.
- [2] T. Berners-Lee, J. Hendler, and O. Lassila. The Semantic Web. *Scientific American*, 284(5):34–43, 2001.
- [3] O. Lassila and R. R. Swick. Resource Description Framework (RDF) Model and Syntax Specification. <http://www.w3.org/TR/REC-rdf-syntax/>, 1999.
- [4] R. G. Smith. The Contract Net Protocol: High-Level Communications and Control in a Distributed Problem Solver. *IEEE Transactions on Computers*, C29(12), 1980.