Available online at www.sciencedirect.com
ScienceDirect
Computer Networks 51 (2007) 4529–4542
Computer Networks
www.elsevier.com/locate/comnet

mRDP: An HTTP-based lightweight semantic discovery protocol

Juan Ignacio Vazquez *, Diego López-de-Ipiña

Faculty of Engineering, University of Deusto, Avda. Universidades, 24, Bilbao 48007, Spain

Available online 24 June 2007

Abstract

Discovery is one of the most important activities in ubiquitous and distributed computing, with a plethora of available protocols. Most of these protocols are designed for one concrete purpose: network nodes discovery, service discovery, search of specific information stored through the network, and so forth.

Designing a single discovery system able to deal with the particularities of many different information structures and purposes seems not feasible. Moreover, these data structures possess some underlaying meanings and relationships that are usually hidden from traditional discovery protocols that use simple text-based matchmaking.

A semantic discovery protocol could solve this problem by taking advantage of semantically annotated data and performing reasoning over the information to obtain additional knowledge that can be crucial in processing the queries.

In this paper, we describe the basics of a novel semantic discovery mechanism called mRDP (Multicast Resource Discovery Protocol) built upon HTTP and Semantic Web technology to provide more powerful discovery capabilities.

Keywords: Web; HTTP; Discovery; Semantic; Protocol

1. Introduction

Discovery has always been a hot topic in networking, generally constituting an initial step to obtain information about available entities in order to perform further interactions. This issue becomes specially important in ubiquitous computing systems or mobile ad-hoc networks, where existing entities must be continuously aware of each other, and adapt the network topology as required.

Although several different discovery protocols have been proposed and used over the last years in concrete architectures, mainly for device and service discovery, there is no common agreement about a unified discovery protocol.

Edwards [9] defines discovery in ubiquitous computing systems as "*a mechanism for dynamically referencing a resource on the network*". Devices and resources come and go on a highly dynamic basis, thus McGrath highlights "spontaneity" and "automatic adaptation" as some of the most important features of the discovery process [18].

* Corresponding author.
E-mail addresses: ivazquez@eside.deusto.es (J.I. Vazquez), dipina@eside.deusto.es (D. López-de-Ipiña).
URLs: http://paginaspersonales.deusto.es/ivazquez (J.I. Vazquez), http://paginaspersonales.deusto.es/dipina (D. López-de-Ipiña).

1389-1286/$ - see front matter
doi:10.1016/j.comnet.2007.06.017

Ubiquitous computing researchers are particularly sensitive to discovery protocols, since an important part of the expected intelligence in the environment is meant to be provided during this initial phase: if the user's PDA is not able to find appropriate devices to communicate with in the environment, it will not be able to provide the required services to the user.

Although this goal may not seem very difficult to achieve – in fact several simple and lightweight discovery protocols have been used in the past – providing a higher level of intelligence during discovery is much more complicated.

For instance, an entity may issue a search request for "devices *inside* a cupboard", "computers *near* the TV" or "digital photo frames with pictures authored *by a friend of* mine".

The above examples are not especially difficult to implement in any ubiquitous computing architecture conveniently prepared to cope with them, especially the meaning of the highlighted terms: if the information is structured in a uniform way (e.g., XML) and stored somewhere, or maybe distributed across the different entities in the network, a query can be created and disseminated to obtain the results.

However, a more difficult approach consists in the possibility of creating a discovery system able to manage any kind of *unexpected* query, dealing with concepts, vocabularies and relationships among the terms that are unknown at the time of system design.

For instance, let's consider the following scenario:

- Objects are tagged in such a way that a containing object (e.g., a wardrobe, or a backpack) is able to know the identity of the objects directly placed inside. An example of such system is formed by objects tagged with barcodes and barcode readers on the containers.
- A PDA is stored in a backpack.
- The backpack is placed in a room.

If a subject is provided with this information and asked to identify available PDAs in the room he will surely point at the PDA inside the backpack as one of the items: location is a transitive property, which means that if the PDA is in the backpack, which in turn is placed in the room, the PDA is located in the room.

However, computing systems are not able to interpret information at as fine a level as humans do; they do not know about the properties of the location concept. Unless explicitly programmed to do so, "location" is a word like any other such as "point to": the PDA is pointing to the backpack, which in turn is pointing to the room, does not imply that the PDA is pointing to the room.

An intelligent discovery mechanism must provide a higher level of context interpretation and knowledge than traditional text pattern matching schemes: it dives through the information relationships, understanding their implications. Fortunately, there is one technology able to provide the required framework for annotating data and relationships and creating knowledge that can be used in the discovery process: the Semantic Web.

The Semantic Web is particularly interesting for discovery because of its *future-proof* characteristics: it can provide a solution framework for "*problems and situations yet to be defined*" [15]. Therefore, it can provide the mechanism to support the creation of undefined search queries on concepts unknown at the moment of design, exploring the connections of information structures to provide a much deeper knowledge, and more refined search results.

This paper presents mRDP, a discovery protocol based on semantic queries. In Section 2 a number of discovery protocols and architectures are mentioned along with their limitations. Section 3 provides an introduction to semantic queries. Section 4 introduces the basics of mRDP – Multicast Resource Discovery Protocol, a lightweight protocol for semantic discovery in local area networks. Sections 5 and 6 describe the format of queries and protocol messages, respectively. In Section 7 measures of mRDP performance are included. Finally, the conclusions and a brief discussion about future lines of semantic discovery are provided.

2. Related work

Different discovery protocols such as SLP [13], Jini [26], UPnP SSDP [12] and other alternatives [9] have been widely used in the past, and their limitations have been clearly identified [6]:

- Lack of rich representation: the existing architectures lack of expressive languages, representation and tools for the broad range of service descriptions.
- Lack of constraint specification and inexact matching: most protocols require exact matching, with a simplistic notion of constraints. Lack of semantic matching.