Peer-to-Peer Remote Service Discovery in Pervasive Computing*

Stefan D. Bruda[1], Farzad Salehi[1], Yasir Malik[2], and Bessam Abdulrazak[2]

[1] Bishop's University, Sherbrooke, QC J1M 1Z7, Canada
[2] Université de Sherbrooke, Sherbrooke, QC J1K 2R1, Canada
stefan@bruda.ca, salehi@cs.ubishops.ca,
{yasir.malik,bessam.abdulrazak}@usherbrooke.ca

Abstract. Service discovery is important in realizing the concept of pervasive computing; consequently, service discovery protocols must be able to work in the heterogeneous environment offered by this computing paradigm. Remote service discovery in particular has not been properly achieved so far. In an attempt to remedy this we propose a new architecture for enabling local service discovery mechanisms to discover services remotely. We first base our architecture on Universal Plug and Play (UPnP) as an example of local service discovery protocols, and Gnutella as an example of peer-to-peer distributed search protocols. We introduce a module called *service mirror builder* to the UPnP protocol, and a remote communication protocol over a Gnutella network. We then note that our architecture is actually independent on any locally deployed discovery protocol and thus supports full interoperability.

1 Introduction

Mark Weiser gave birth to the vision of "ubiquitous computing" or "pervasive computing" (as named throughout this paper). He defined it as follows: *Ubiquitous computing is the method of enhancing computer use by making many computers available throughout the physical environment, but making them effectively invisible to the user* [1]. Computing anytime, anywhere, and in any device means a massive presence of computing devices in the physical world. At the same time, people should be able to access information and computation in a user-centric way i.e., user interaction with such a system must be natural and comfortable.

Pervasive computing offers an environment saturated with sensors, actuators, cameras, and other sorts of computing devices; all these devices should work together and satisfy users' needs with minimal user intervention. Service discovery protocols are one tool that accomplished this. Many service discovery protocols have been designed. The dominant protocols (at least for home appliances) include Microsoft Universal Plug and Play or UPnP [2], Bluetooth Service Discovery Protocol [3], Apple's Bonjour, and Sun's Jini technology [4].

Current service discovery protocols are designed for home or enterprise environments [5]; the pervasive computing environment is however far more heterogeneous

and sophisticated. Furthermore, most service discovery protocols are designed to work only in a local area network (LAN) [6], which is justified for the many services that are physically-oriented (such as video projectors or coffee machines). Still, many other services are not physically-oriented and are accessed by users physically far away from them (such as the digital data in someone's home, or sensors, actuators, and cameras present in a place for security or health care purposes). Computing anywhere is also the very definition of the concept of pervasive computing. While it is not possible to provide all services anywhere, remote access to any services (from anywhere) makes sense. For this purpose service discovery protocols must be able to discover services remotely in order to be able to work in a pervasive computing environment. A combination of existing technologies and services enables some level of remote access, but seamless discovery and control of remote services is currently not possible [6–8].

The objective of this work is to enable local service discovery protocols (such as UPnP) to discover remote services which are not available locally. We lay the basis of such remote service discovery by proposing a suitable architecture. We use UPnP as an example of service discovery protocols. In our architecture each local UPnP network is enhanced by a function called service mirror builder. A service mirror builder presents local services as remote services to other UPnP networks, and also builds mirrors of remote services in its local network. The process of finding a remote service uses the distributed peer-to-peer search protocol Gnutella (though other implementations are also possible).

A service mirror builder is seen as an UPnP enabled device in the local UPnP network. It is worth emphasizing that UPnP is just an example; the service mirror builder can be generally defined as a service discovery enabled device with respect to any service discovery protocol.

1.1 Motivation

Between other things pervasive computing means spatial heterogeneity: some places offer all the needed services and others only have a few services to offer. Therefore a combination of remote and local services is sometimes needed. The following scenarios motivate our quest for remote service discovery.

One example of pervasive computing environment is a *connected (smart) home*, which is a dwelling incorporating a communications network that connects key devices (sensors and actuators, electrical appliances) and allows them to be remotely controlled, monitored or accessed [7]. To realize a smart home we thus need to have a mechanism to access its services remotely. In addition, most of us desire seamless storage, access and consumption of digital content from and to any compatible digital device in a home or smart home; ideally, users should be able to access their residential services from anywhere using any type of terminal [8]. Overall use cases for remote service discovery therefore include lighting, residential climate control, home theater, audio entertainment systems, domestic security, domestic health care systems, etc.

Vendors need to connect to their devices for various purposes such as to update their software or perform routine checks (*remote support*). Security and health care companies in particular need to be in contact with their customers and their products continuously. The information from sensors, actuators and cameras can be monitored

by such companies, which can then take action in case of any threat, but also control devices to be more efficient and usable. The vendors can also advertise features and offer upgrades to their devices (*continuing close presence*).

Massively Multiplayer Games (MMGs) are traditionally supported by a client-server architecture, but such a centralized architecture lacks flexibility and can put communication and computation stress on the servers [9]. To overcome these problems inherent to centralized solutions, peer-to-peer networks are emerging as a promising architecture for MMGs [9]. Running MMGs with the help of remote service discovery and without any centralized coordinator is arguably the best use cases to motivate our research contribution.

2   Preliminaries

2.1   UPnP

The role of Universal Plug and Play (UPnP) [2] is the automatic discovery and configuration of any new devices that connect to a computer network. UPnP supports zero configuration networking or Zeroconf, meaning that UPnP creates an IP network without any need of manual configuration or configuration servers.

UPnP uses the Internet protocol suite. Special features include the following [2]: *media and device independence* (any network media or device which supports IP can be a basis for the establishment of UPnP), *user Interface (UI) control* (devices can have a UI written by XML which is readable by a browser), and *operating system and programming language independence*.

UPnP has three major components: *device* (contains one or more services), *service* (performs actions and shows its state; consists of a state table, control server and event server), and *control point* (a system that discovers and then controls services and devices). The functioning of UPnP then involves six steps:

Addressing. Each device must have a Dynamic Host Configuration Protocol (DHCP) client. When the device connects to the network for the first time it must search for a DHCP server. If a DHCP server exists, then the device receives an IP address this way. Otherwise the device must assign an IP address to itself (Auto-IP [10]).

Discovery. Discovery is the process of discovering the capabilities of the devices on the network. It can take place in two ways.

First, when a new device gets an IP address and so is connected to the network, the device must multicast discovery messages, advertising its embedded devices and services. This process is called *discovery-advertisement*. Any interested control point in the network can listen to these advertisements and then connect and control the originating devices or only some of their services.

Secondly, when a new control point is established in the network. Such a new control point multicasts a Simple Service Discovery Protocol (SSDP) message [11], searching for available devices and services. All devices in the network must listen to this kind of messages and respond to them whenever any of their services or embedded devices matches the criteria from the SSDP messages. This process is called *discovery-search* [2].

Description. Once discovery is complete and the control point knows about the existence of one device or service, it must also find out how to invoke that device or service. The respective control point retrieves the device description from the URL provided by the device in the discovery message. The UPnP description for a device is expressed in XML and includes vendor-specific information, manufacturer information, a list of any embedded devices or services, as well as URLs for control, eventing, and presentation [2].

Control. Now that the control point has a clear overview of the service and knows how to control it, it can send an action request. The control point sends a control message to the device according to the respective service control description. Control messages are expressed in XML. In response, the service will return action specific values or fault codes [2].

Eventing. Services keep control points informed by sending them event messages. Event messages contain the last update of changed state variables in the service. This process is called eventing [2].

Presentation. Some devices have URLs for presentation. Such an URL can be fetched and then presented in a browser by the control point. According to the device capabilities and URL presentation definition, a user can then see the status of the service and even control it [2].

2.2   Gnutella

A distributed network architecture may be called peer-to-peer (P2P) whenever the participants share a part of their own hardware resources (processing power, storage capacity, network link capacity, printers, etc.) with each other in order to provide the service and content offered by the network (e.g., file sharing or shared workspace for collaboration). Furthermore these resources are accessible by other peers directly, without passing through intermediate entities. The participants in such a network are thus resource providers and at the same time resource requesters (the "servent" concept) [12]. Peer-to-peer file sharing is a particular example of peer-to-peer network. Each peer in a peer-to-peer file sharing network is implemented by a client which uses some distributed search protocol to find other peers as well as the files that are being shared by them. Different protocols for distributed search are being used by peer-to-peer file sharing programs, the most prominent being BitTorrent [13] and Gnutella [14].

Because of the distributed nature of Gnutella and its independency from any central servers, a Gnutella network is highly fault-tolerant. Indeed, a network can work continuously despite the fact that different servents go off-line and back on-line [14]. We describe in what follows the Gnutella protocol [14–16]. The first time a servent wants to join a Gnutella network, its client software must *bootstrap* [17] and thus find at least one other servent (node, peer) in the network. The bootstrap can happen automatically or manually, either out of band or using Gnutella Web caches (caches that include a pre-existing list of addresses of possibly working hosts may be shipped with the Gnutella client software or made available over the Web).

A node connected to a Gnutella network keeps in touch periodically with its (directly connected) neighbours through ping messages. These messages are not only replied to (by pong messages) but they are also propagated to the other interconnected servents.

When a servent receives a ping message, it sends it to the nodes to which it is connected. Once the servent finds at least one active peer in the Gnutella network, it can create an updated list of active servents by observing the ping and the corresponding pong messages.

When a client wants to search for a file (or as we will see in Sect. 4 for a service), it sends a query to all its directly connected neighbour servents (except the one which delivered this query message), which forward the query to their neighbours, and so on. This process repeats throughout the network. A *query message* is the primary mechanism for searching the distributed network. If a servent receives a query and finds a match in its directory, it will respond to it with a *query-hit message*. A query-hit contains enough information for the retrieval of the data matching the corresponding query.

To avoid flooding the network the query messages contain a TTL (Time To Live) field. It is possible that one query reaches a servent more than one time. To avoid serving a query more than once, each query is identified by a unique identification (muid). Before processing a query a servent checks the query's muid against a table of previous muids; a hit causes the query to be dropped.

The query-hit can go back along the reverse path of the query to reach the servent which requested it, or it can be sent directly to the requester.

3 Related Work

Along with summarizing previous work on remote service discovery we also anticipate a bit and take the opportunity to compare the previous research with our solution (which will actually be introduced later in Sect. 4).

3.1 Remote Access to UPnP Devices Using the Atom Publishing Protocol

The network topology of one architecture for remote service discovery in UPnP [6] consists of at least two network segments: the home network and the remote network. The architecture assumes that there is an IP tunnelling mechanism such as a Virtual Private Network (VPN) between the two network segments. The architecture introduces a new element called *UPnP Device Aggregator* which is acting as a proxy for the existing standard UPnP devices. *Enhanced UPnP Devices* or *Control Points* are then UPnP devices or control points which are compatible with this remote service discovery architecture. The UPnP Device Aggregator aggregates information about the services and devices in the local network as an Atom feed, which can then be retrieved (using GET commands) by the enhanced UPnP control points in the remote network. Additionally, a UPnP Device Aggregator can receive information from remote Enhanced UPnP Devices and present them to the local control points. This information can be received by the UPnP Device Aggregator via HTTP POST.

The main shortcoming of this architecture is the need for VPN. Indeed, VPN does not scale well, requiring careful administration of IP addresses and subnetworks [8]. VPN also limits the architecture to the domains within the VPN network (limiting heterogeneity). No such limiting factors are present in our architecture, which is substantially more scalable. In addition, all remote service discovery requests are addressed to the

home network, so this architecture can be considered centralized or partially centralized: service coordinators (the UPnP Device Aggregators) register and cache services [7]. By contrast, our architecture is fully distributed: no centralized coordinator is necessary. We note that Gnutella has switched to a hybrid architecture using Ultrapeers [16] for efficiency purposes, but even in this case we obtain a more distributed architecture.

3.2   Presence-Based Remote Service Discovery

An architecture for remote service discovery and control based on presence service (as used in instant messaging and VOIP) was also proposed [8]. A presentity can be anything that can have a presence state (be present or absent); presence information is sent to a presence service, which is a network service that records and distributes presence information. In the remote service discovery architecture based on presence service [8] there are two new functions called *service discovery gateway* and *service virtualizer*. The gateways register local services as presentities in a presence server. They can also retrieve other presentities from the presence server and present them to the service virtualizer. The service virtualizer uses this presence information to virtualize a local service in the local network. That is, a service virtualizer presents a remote service as a local one.

This architecture is partially centralized, as remote service providers and remote service requesters must first find a presence server to register or request a service. Although presence servers (as service coordinators) provide service visibility, the benefit does not come without cost and complexity [7, 18]. By contrast, our architecture is fully distributed.

3.3   Transparent UPnP Interaction between UPnP Gateways

Dynamic Overlay Topology Optimizing Content Search (DOTOCS) [19] enables flexible content searches among UPnP gateways. DOTOCS aims to establish an optimized peer-to-peer overlay network among UPnP gateways. DOTOCS uses a communication protocol between UPnP local networks called transparent interaction solution and described elsewhere [20]: The communication between two connected UPnP local networks across the Internet is accomplished using the Web service technology. A local gateway encapsulates Simple Service Discovery Protocol (SSDP) messages into Simple Object Access Protocol (SOAP) messages and transmit them to another gateway over the global network. A Web service at the destination UPnP gateway extracts the SSDP message and replaces the original IP address (which is not valid in this local network) with the IP address of the gateway itself. The gateway then multicasts this discovery search message in the local UPnP network. If any device responds to that message (meaning that the device has the service demanded by the SSDP message), then the gateway encapsulates that message into another SOAP message and sends it back to the first network. This way one local UPnP network can discover remote services from a different UPnP network.

Scalability between local networks is manageable when this solution is used. However, each gateway multicasts in its local UPnP network any received discovery message (regardless whether the demanded service in that discovery message is locally available

or not). This creates substantial traffic in the local network, most of it useless, which reduces scalability. Our protocol does not multicast remote requests to the local network (for indeed the service mirror builder has already discovered the locally available services), so the local UPnP network will not be loaded with spurious messages. Scalability therefore only depends on the Gnutella network (which is scalable to a high degree).

3.4   Peer-to-Peer Caching of Classes of Services

Chakraborty et. al. [21] propose a distributed service discovery protocol based on the concepts of peer-to-peer caching of service advertisements and group-based intelligent forwarding of service requests. Services are described using the Web Ontology Language (OWL). The semantic class/sub-class hierarchy of OWL is used to describe service groups and then this semantic information is used to selectively forward service requests. OWL-based service description also enables increased flexibility in service matching.

While the hierarchical structure of service classes improves efficiency, the protocol is mainly proactive rather than reactive: services need to advertise actively before becoming accessible to the remote devices. We believe that this can easily cause either increased traffic (when the time delay between successive advertisements is low) or increased response time (when the mentioned delays are large). This is all mitigated by the class/sub-class mechanism, but cannot be eliminated altogether. By contrast our architecture is completely reactive, thus minimizing network traffic.

4   A New Distributed Architecture for Remote Service Discovery

Recall that remote services are not present in the current physical location of the controller but are available to the controller upon request. A control point may also reside in a pervasive computing environment with heterogeneous protocols and networks. Even if some otherwise available services in the local domain could not be accessed because of heterogeneity in protocols (networks, ontologies, etc.), the controller may still be able to remotely access services within its capabilities but far from its physical location. In other words, sometimes service discovery protocols could not see all the available services in their domain, but if they could just bridge to neighbour networks (with the same protocols and ontologies) they could accomplish their tasks.

We propose a new architecture that accomplishes remote service discovery in a fully distributed manner i.e., without the need of any centralized, coordinating entity. Our architecture allows the discovery of services in local and remote domains, and offers a solution for automatic discovery and control of remote services. We use UPnP as an example in our architecture, but in fact we try not to depend on any particular service discovery protocol.

Fig. 1 shows our architecture. There are 5 local networks in the figure, labelled from 1 to 5. Each of these local networks offers local services, devices, and control points. These devices, services, and control points are connected with each other locally through UPnP. In each local network there is one special function called *service*
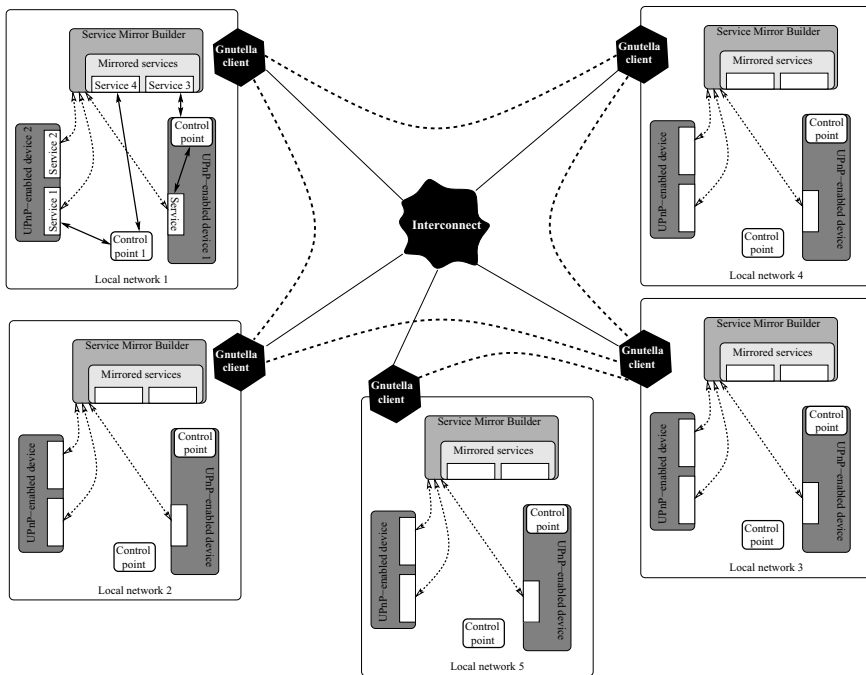
Fig. 1. A distributed architecture for remote service discovery (doted lines connecting local networks show the Gnutella network overlay)

mirror builder. This special function performs the actual remote service discovery and is seen by the local network as just another UPnP-enabled device.

In addition, each local network runs a Gnutella client software. These clients are specialized clients that share local services to the outside world and find services requested by their service mirror builder. The local networks establish a Gnutella network between them. Dotted lines connecting local networks in the figure show the overlay of the Gnutella network.

4.1   Changes to the Local Networks

Our architecture introduces a *service mirror builder* in every local network. The network is an IP based network with all of these devices connected through UPnP (the UPnP protocol with its six steps is described in Sect. 2). *Addressing* is accomplished using the normal UPnP protocol.

Discovery-advertising and *discovery-search* are then performed in the local network as prescribed by the local UPnP protocol. The service mirror builder must be aware of all the available services in the local network, so it never ignores any multicast message. It will also advertise its services (that are all remotely discovered as we will see later) as they become available. During any kind of discovery-search process (that is, whenever a control point becomes interested in a new service) the respective control

point multicasts a discovery message, thus searching for available services and devices in the network. All the devices listen to these messages and respond whenever any of their services match the criteria specified by the request. Additionally, the service mirror builder listens to these messages as well. It checks whether the requested service is in the list of available local services, case in which the message is ignored; otherwise, the service mirror builder performs remote discovery for that service.

Refer to Fig. 1 for a closer look at one of the local networks (namely, local network 1). This network features four components: two UPnP-enabled devices (labelled Device 1 and Device 2), one control point (Control point 1) and one service mirror builder (SMB for short). The service mirror builder typically resides on the smart environment gateway (such as a connected home gateway). Suppose that Device 1 has not introduced its service to other control points except the service mirror builder, and its control point has discovered a mirror of a remote service (Service 3). Device 2 is a UPnP device with 2 embedded services (Service 1 and Service 2) which are also not known to the others. Then Device 2 must inform all the available control points in the network about its services; it does so by multicasting a message and thus advertising Services 1 and 2 (discovery-advertisement). The multicast message will be received by the service mirror builder and by Device 1. The control point of Device 1 is not interested in (or not capable to control) either Service 1 or Service 2 and so it ignores this message. However, the service mirror builder is aware of all the available services in the local network, so it cannot ignore any multicast message. The service mirror builder uses this information for remote service discovery, which will be discussed later. In local network 1 the service mirror builder is interested in Service 1 and Service 2. It then sends a message to device 1 to retrieve the description of the two services as per the description UPnP step.

Suppose now that Device 1, Device 2, and the service mirror builder have all discovered each other. Control point 1 then joins the network and obtains an IP address, but has not discovered any services to control yet. In such a case the newly added control point multicasts a Simple Service Discovery Protocol (SSDP) discovery message, thus searching for available services and devices in the network (discovery-search). The devices in the network listen to these messages and respond whenever any of their services match the criteria specified therein. The service mirror builder listens to all these messages and proceeds to remote discovery for the respective service whenever the requested service is not provided locally. In our example, Service 1 is matched with the request of Control point 1. Therefore Device 2 unicasts a response message to Control point 1. Now that Control point 1 has discovered the service, it asks for a description. Once the description is received, Control point 1 can control Service 1 in Device 2.

Consider now that Control point 1 multicasts a discovery-search message requesting a service which is not locally available (Service 4). The service mirror builder will recognize that this service is not locally available, and so it sends a query for that service to the local Gnutella client. The Gnutella client will then pass that query to the Gnutella network (see Sect. 4.2). Once such a service is found, a mirror of that service is made available in the local network. In local network 1 from Fig. 1 the mirrors of the remote services are shown in the service mirror builder box (Services 3 and 4).
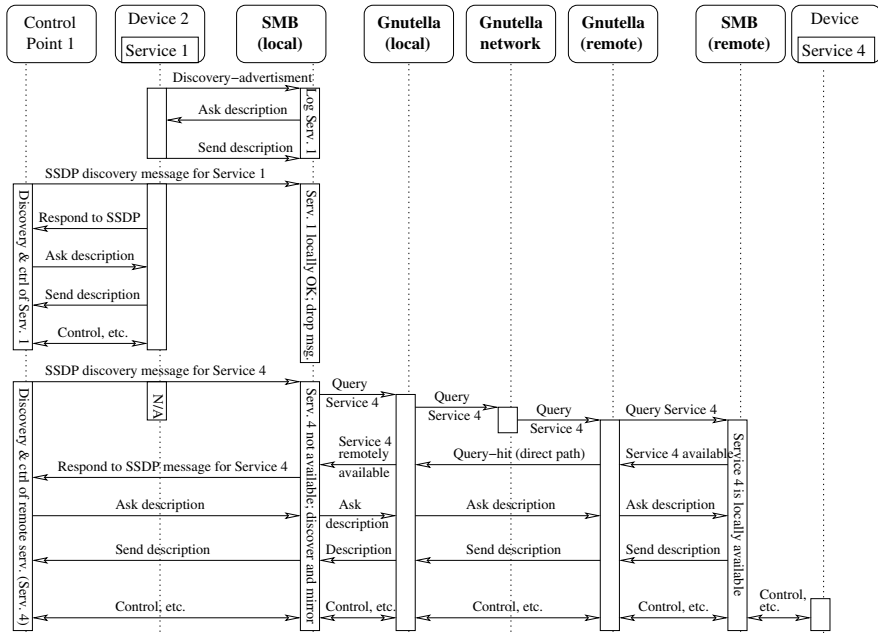
Fig. 2. A sequence diagram detailing the behaviour of our distributed architecture for remote service discovery

After the discovery step (which makes the control points aware of the available services), the control points must find out how to use these available services (*description*). Advertising messages circulated during discovery contain URLs from which the control points can retrieve the description of the respective devices. Once a control point has the device or service description it can invoke actions on that service and get result values in return. Invoking an action in UPnP is a particular instance of Remote Procedure Call [2]. The major focus of this research contribution however is service discovery so we will not discuss service control, eventing and presentation any further.

4.2 Remote Service Discovery

Two major characteristics of pervasive computing are distributedness and mobility. In such an environment we want to connect nodes in a distributed manner and without any dependency to a central server (such as the presence server used in Sect. 3.2). We therefore connect the local networks using Gnutella, a strongly decentralized peer-to-peer system [15].

Gnutella servents can share any type of resources [15]. In our design they are sharing the services available in their local networks. The overlaying Gnutella network (dotted lines in Fig. 1) is established according to the protocol.

Now that both the local networks and the Gnutella network are established, remote service discovery can begin. Such an event happens whenever a control point requests a service which is not locally available. The service mirror builder then activates and tries to remotely discover it.

Each service mirror builder has a cached description of all of the available local services (obtained during the local discovery phase as explained earlier). When a control point requests a service, the service mirror builder checks in its local service directory to see if the service is already available in the local network. If this is not the case, then the service mirror builder proceeds to discover it remotely by sending a request for the respective service to the Gnutella client. The Gnutella client in turn issues a query message asking for the requested service to the Gnutella network according to the Gnutella protocol. When receiving a query, a Gnutella client sends the included service request to the local service mirror builder, which in turn will check the availability of the requested service in its local network. Should the service be locally available, the service mirror builder communicates this to the Gnutella client, which in turn responds with a query-hit message to the original requester. Overall, the query is answered with a query-hit by the nodes that offer the respective service. These nodes also send a service description and other information back to the node that issued the query. This information is then delivered to the service mirror builder of that node, which creates a mirror of the service in the local network. The control points in the local network see the service just likes a local one and can control it in the usual way.

Suppose that some control point (such as Control point 1) requests a service which is not available in any of the participating local networks; in such a case the respective Gnutella client returns no hits. Whenever the service becomes available in the local network, it will be made available through discovery-advertisement; similarly, the Gnutella client will re-issue the corresponding query periodically until either (a) the service becomes available in the local network, (b) the service is discovered remotely, or (c) the control point that requested the service disappears. This mechanism extends the discovery-search mechanism almost transparently (but with some delay).

The functioning of the whole protocol is summarized in Fig. 2.

4.3   Gnutella and UPnP Messages in the New Architecture

We show the possibility of using the Gnutella distributed search protocol to search for services in remote networks (remote services). We do this by discussing the Gnutella and UPnP message structure and the modifications that are needed in our architecture.

In our architecture all local network components communicate and work with each other under UPnP protocol standards. All six steps in UPnP (addressing, discovery, description, control, eventing and presentation) are being done as per the UPnP protocol.

As far as the remote connections are concerned, all servents are working under the Gnutella standards and specification. All Gnutella connect, Gnutella OK, ping and pong messages are exactly according to the available Gnutella protocol. The only differences happen in the Gnutella query and Gnutella query-hit messages (since the original messages are used for requesting for and responding with shared files). We show the structure of these messages in more detail along with recommendations for changing them to work in the new architecture for the purpose of service discovery instead of file sharing.

All of the Gnutella protocol messages, including query and query-hit, include a header with the byte structure described in Table 1(a) [22]. The payload type indicates the type of the message. Other payload types can also be used as long as all the participating servents support them [22]. Payload length shows the size of the payload.

Table 1. Gnutella message header (a) and query message structure (b)

	(a)		(b)
Bytes	**Description**	**Bytes**	**Description**
0-15	Message ID/GUID (Globally Unique ID)	0-1	Minimum speed
16	Payload Type	2	Search criteria
	0x00 = Ping 0x01 = Pong	Rest	Optional extension block
	0x02 = Bye 0x40 = Push		
	0x80 = Query 0x81 = Query-Hit		
17	TTL (Time To Live)		
18	Hops		
19-22	Payload Length		

Table 2. Gnutella query-hit message structure

Bytes	Description	Result set structure:	
0	Number of Hits	**Bytes**	**Description:**
1-2	Port	0-3	File Index
3-6	IP Address	4-7	File Size
7-10	Speed	8-	File Name (null-terminated)
11-	Result Set	x	Extensions Block (null-terminated)

The whole Gnutella message should be no more than 4 kB in size. Immediately following the message header is a payload which can be one of the following messages: ping, pong, query, query-hit and push [22]. This message header structure will remain unchanged in our architecture.

The Query Message. Since queries are broadcast to many nodes, servents normally send query messages that are smaller than 256 bytes; however, query messages can be as large as 4 kB. A query message has the structure shown in Table 1(b) [22]. The rest field of a query message is used for the original query which in our case is a query for a remote service. The allowed extension types in the rest field can be specified using the Gnutella Generic Extension Protocol (GGEP), Hash/URN Gnutella Extensions (HUGE), and XML [22]. The Gnutella Generic Extension Protocol (GGEP) allows arbitrary extensions in a Gnutella message; a GGEP block is a framework for other extensions [22].

In a UPnP network service discovery is accomplished using Simple Service Discovery Protocol (SSDP). All SSDP messages are sent using the HTTP protocol. The HTTP and Gnutella protocols are both application layer protocols. The fundamental data in a SSDP discovery search or discovery-advertisement message (in a UPnP network) contains a few essential specifics about the device or one of its services (its type, universally unique identifier, etc.) [11]. All this information can be readily encoded in a GGEP extension by the service mirror builders and then sent to the Gnutella network agent. Then the Gnutella network agent can put this GGEP-formatted information in the Rest part of a Gnutella query message and send it to the Gnutella network.

The Query-hit Message. The structure of a query-hit message is shown in Table 2 [22]. The result set is used for the response to the query; its structure is also shown in the table.

The first three fields of the result set are defined specifically to hold information about a requested file or file portion, as Gnutella is mainly used for file sharing. In our case it is possible to redefine these fields; to prevent increased complexity and extra work to define a new specification, we recommend that these fields be filled with some default labels. In other words these fields of the result set are simply ignored.

GGEP, HUGE, and plain text metadata are all allowed in the extension block. We recommend that the response messages from service mirror builders be formatted in a GGEP extension and sent back to the network in the extensions field of the query-hit message.

5   Protocol Independence and Performance Considerations

The use of particular protocols (UPnP and Gnutella) are in fact necessary only for illustration purposes. The architecture described here can be in particular adapted in a trivial manner to handle any local service discovery protocol.

Through the use of UPnP we have included in our discussion a comprehensive set of features for the local service discovery protocol, but this is done without loss of generality: those protocols that lack some of the features can still function in our framework, which is agnostic with respect to which of the described services are actually used. Conversely, to our knowledge no protocol provides more features than the ones discussed here.

The peer-to-peer architecture is clearly immaterial to the discussion, as long as the architecture is capable of maintaining an ad-hoc network, propagating requests for distributed search throughout the network, and handling responses to such requests (all peer-to-peer protocols do). The implementer can in effect choose any other peer-to-peer protocol.

The only somehow substantial changes that are necessary to switch protocols are related to the form of the messages passed along between the peer-to-peer servents. The message structure outlined in Sect. 4.3 is specific to Gnutella; should another peer-to-peer network be chosen, this structure will likely change. Further encapsulation might be needed to pass along messages generated by different local service discovery protocol, though this is unlikely.

The performance of our architecture is the same as the performance of the underlying Gnutella (or more generally peer-to-peer) network. All the advantages (and also disadvantages) of Gnutella are thus maintained. This is actually one reason for our choice of peer-to-peer protocol. Indeed, Gnutella is one of the most studied such a protocol. Many changes and refinements have been added to it throughout the years to make it more efficient. Such refinements and techniques include Ultrapeers and Leaves, Distributed Hash Table, Query Routing Protocol (QRP), and so on. This all helps to reduce the traffic in the Gnutella network and generally increase the efficiency of the protocol. These refinements can be trivially added to our architecture. Future refinements of

the protocol can also be readily integrated. It can be argued that Gnutella is not fully distributed, since the bootstrapping process is centralized. However, such a centralized mechanism is used for empirical reasons (works well in most cases), but different bootstrapping techniques exist, including random address probing and network layer probing with anycast or multicast messages [23]. These techniques can be readily used should a fully distributed architecture be necessary.

We already mentioned the advantages of our approach compared to the only other peer-to-peer approach known to us [21]. Indeed, our approach is reactive and so we believe that it minimizes the network traffic in the general case. True, a "sweet spot" can probably be achieved in the ontology-based approach [21] by changing the time delay between advertisements. Such a "sweet spot" has the potential of offering a better performance; however, it is also likely that different circumstances may require different values, making a global solution unlikely. By contrast Gnutella is designed to offer consistent performance in many circumstances; indeed, optimization adjustments are already implemented! We thus believe that overall our solution offers more consistent performance, in addition to being able to benefit from any future improvement of the Gnutella protocol (as mentioned above).

Classifying services based on an ontology [21] does reduce network traffic and has the potential of improving the overall efficiency. We believe that such a technique is orthogonal to the service discovery mechanism. We thus believe that the ontology-based approach can be used equally effectively in out setting. Doing this however is a subject of future research.

6   Conclusions

Service discovery plays an important role in pervasive computing. At the same time pervasive computing creates many challenges for service discovery protocols, one of them being remote service discovery. We introduced a new approach to remote service discovery that is flexible, decentralized, and fully distributed.

The core part of the new architecture is the new function in a UPnP network called service mirror builder and its cooperation with a specialized Gnutella client software to discover remote services and then present these remote services as local ones. Conversely, a service mirror builder can also control local services to serve them as remote services for other, remote service mirror builders. The service mirror builder communicates with the specialized Gnutella client software (from the point of view of the local network however the service mirror builder is just a UPnP-enabled device). We used UPnP for illustration purposes, but the service mirror builder can be defined based on *any* service discovery protocol (Bluetooth, Apple Bonjour, etc.). Our solution is in fact general and not dependent on any particular service discovery protocol.

We propose Gnutella as a distributed search protocol for discovering remote services. The very design of a Gnutella network as a decentralized and distributed protocol moves this remote service discovery architecture one step ahead toward truly distributed computing. Our architecture is also readily amenable to improvements in either the underlying peer-to-peer protocol or the ontology-based classification of services (or indeed both). We therefore believe that our approach offers better compatibility with pervasive computing.

References

1. Weiser, M.: Some computer science issues in ubiquitous computing. Communications of the ACM 36, 75–84 (1993)
2. Understanding Universal Plug and Play (2000), White paper: `http://www.upnp.org/download/UPNP_understandingUPNP.doc`
3. Bluetooth Special Interest Group (SIG): Specification of the Bluetooth System Version 1.1 (2001), `http://www.tscm.com/BluetoothSpec.pdf`
4. Sun Microsystems: Jini Technology Core Platform Specification Version 1.2 (2001), `http://www-csag.ucsd.edu/teaching/cse291s03/Readings/core1_2.pdf`
5. Zhu, F., Mutka, M., Ni, L.: Service discovery in pervasive computing environments. Pervasive Computing 4, 81–90 (2005)
6. Belimpasakis, P., Stirbu, V.: Remote access to universal plug and play (UPnP) devices utilizing the Atom publishing protocol. In: International Conference on Networking and Services, p. 59. IEEE Computer Society (2007)
7. Feng, W.: Remote Service Provision for Connected Homes. PhD thesis, De Montfort University (2010)
8. Häber, A.: Remote Service Discovery and Control for Ubiquitous Service Environments in Next-Generation Networks. PhD thesis, University of Agder (2010)
9. Buyukkaya, E., Abdallah, M., Cavagna, R.: VoroGame: A hybrid P2P architecture for massively multiplayer games. In: 6th IEEE Consumer Communications and Networking Conference (CCNC), pp. 1–5. IEEE (2009)
10. Cheshire, S., Aboba, B., Guttman, E.: Dynamic Configuration of IPv4 Link-Local Addresses. Internet Engineering Task Force, RFC 3927 (2005)
11. UPnP forum: UPnP Device Architecture 1.1 (2008), `http://www.upnp.org/specs/arch/UPnP-arch-DeviceArchitecture-v1.1.pdf`
12. Schollmeier, R.: A definition of peer-to-peer networking for the classification of peer-to-peer architectures and applications. In: 1st International Conference on Peer-to-Peer Computing, pp. 101–102 (2001)
13. Cohen, B.: The BitTorrent Protocol Specification (2008), `http://www.bittorrent.org/beps/bep_0003.html`
14. Clip2 Distributed Search Services: The Gnutella Protocol Specification Version 0.4 (2003), `http://www.stanford.edu/class/cs244b/gnutella_protocol_0.4.pdf`
15. Ilie, D.: Gnutella Network Traffic-Measurements and Characteristics. Master's thesis, Blekinge Tekniska Högskola (2006)
16. Oram, A.: Peer-to-Peer: Harnessing the Benefits of a Disruptive Technology. O'Reilly Media (2001)
17. Gtk-Gnutella: Gnutella Bootstrapping (2011), `http://gtk-gnutella.sourceforge.net/en/?page=bootstrap`
18. Engelstad, P., Zheng, Y., Tore, J.: Service discovery and name resolution architectures for on-demand MANETs. In: 23rd International Conference on Distributed Computing Systems, pp. 736–742. IEEE Computer Society (2003)
19. Kawamoto, E., Kadowaki, K., Koita, T., Sato, K.: Content sharing among UPnP gateways on unstructured P2P network using dynamic overlay topology optimization. In: 6th IEEE Consumer Communications and Networking Conference (CCNC), pp. 1–5. IEEE (2009)

20. Ogawa, M., Hayakawa, H., Koita, T., Sato, K.: Transparent UPnP interactions over global network. In: Proceedings of SPIE, vol. 6794 pp. 67944P (2007)
21. Chakraborty, D., Joshi, A., Yesha, Y., Finin, T.: Toward distributed service discovery in pervasive computing environments. IEEE Transactions on Mobile Computing 5, 97–112 (2006)
22. Klingberg, T., Manfredi, R.: Gnutella 0.6. Network Working Group (2002)
23. Conrad, M., Hof, H.-J.: A generic, self-organizing, and distributed bootstrap service for peer-to-peer networks. In: Hutchison, D., Katz, R.H. (eds.) IWSOS 2007. LNCS, vol. 4725, pp. 59–72. Springer, Heidelberg (2007)