Resource Discovery in Ad-Hoc Grids^{*}

Rafael Moreno-Vozmediano

Dept. de Arquitectura de Computadores y Automática
Universidad Complutense de Madrid. 28040 - Madrid, Spain

rmoreno@dacya.ucm.es

Tel.: (+34) 913947615; Fax: (+34) 913947527

Abstract. The extension of grid computing technology to ad-hoc mobile environments is giving rise to the development of ad-hoc grids, which enable wireless and mobile users to share computing resources, services, and information. However, the adaptation of grid technology to ad-hoc networks is not straightforward, and exhibits numerous difficulties (resource discovery, security, power consumption, QoS, etc.). This paper is focussed on the problem of resource discovery in ad-hoc grids, we study the existing resource and service discovery architectures, analyzing the main limitations of these systems (scalability, discovery delay, adaptation to changing conditions, etc.), and we propose a hybrid mechanism that overcomes these limitations.

1 Introduction

Grid technology enables organizations to share geographically distributed computing and information resources in a secure and efficient manner [1]. Shared resources can be computers, storage devices, data, software applications, or dedicated devices like scientific instruments, sensors, etc. Traditional grid infrastructures are mostly based on wired network resources owned by various individuals and/or institutions, structured in Virtual Organizations, which are subjected to specific sharing policies. Grid middleware provides basic services for resource discovery, resource management, data management, security and communication.

With the proliferation of wireless mobile devices (laptops, PDAs, mobile phones, wireless sensors, etc.), and the development of efficient protocols for communication, routing, and addressing in mobile ad-hoc networks (MANETs), wireless or ad-hoc grids are emerging as a new computing paradigm [2] [3] [4] [5] [6] [7], enabling innovative applications through the efficient sharing of information, computing resources, and services among devices in ad-hoc networks.

However, the development of ad-hoc grids entails new challenges, compared to traditional wired grids. Resource discovery, power consumption, QoS, security, etc. are problems that have still to be solved [3] [4].

In this paper we study in-depth the problem of resource discovery in ad-hoc grids. We classify the existing discovery architectures and we analyze their

^{*} This research was supported by the Ministerio de Educación y Ciencia of Spain through the research grant TIC 2003-01321.

main limitations, such as scalability, discovery delays, bandwidth consumption, adaptation to changing conditions, management complexity, etc. In view of these limitations, we propose a new resource discovery mechanism, based on a hybrid peer-to-peer approach and on the concept of discovery zone, which overcomes the main shortcomings of existing approaches.

2 Classification of Service/Resource Discovery Architectures

Existing service/resource discovery mechanisms can be classified in two main categories:

2.1 Peer-to-Peer Architectures

Peer-to-peer (P2P) architectures use fully distributed mechanisms for resource or service discovery, where networks entities (providers and clients) negotiate on-to-one with each other to discover the available services and their attributes, and to find those services that meet the user requirements. Two basics mechanisms can be used to service or resource discovery in peer-to-peer systems: query mechanisms and advertising mechanisms.

P2P Query-Based Systems (P2P-Query). In P2P-Query, also called active or pull P2P mechanisms, clients send a discovery message to the network, by broadcasting or multicasting, asking for services or resources that match same specific requirements or attributes. Providers respond to the client query by sending a description of the service or resource attributes. Examples of P2P query-based systems are the Service Discovery Protocol (SDP) used in Bluetooth [8], the service discovery mechanism proposal for on-demand ad-hoc networks [9] [10], and the Konark active pull protocol [11].

P2P Advertisement-Based Systems (P2P-Adv). In P2P-Adv, also called passive or push P2P mechanisms, providers advertise periodically, by broadcasting or multicasting the location and attributes of resources and services, so that clients can build a local database with all the resources available on the network. Examples of P2P Advertising mechanisms are the Universal Plug and Play (UPnP) discovery service [12] developed by Microsoft, and the Konark passive push protocol [11].

Peer-to-peer architectures are useful for very dynamic ad-hoc environments, where network infrastructure is unpredictable, and the presence of permanent dedicated directories can not be guaranteed. However, these mechanisms, which are based on broadcasting (flooding) or multicasting, suffer from huge bandwidth usage and very low scalability, so they only suit well for small networks. Advertising mechanisms use much more bandwidth and scale worst than query mechanisms, since unsolicited information is issued periodically to the network. However, they reduce the lookup time, since every client holds updated information about all the resources and services that are available in the network.

2.2 Directory-Based Architectures

Discovery architectures based on directory use a centralized or distributed repository, which aggregates and indexes the information about resources and services offered in the network. Providers register their resources and services with this directory, and clients query the directory to obtain information about resources or services. There are three different general schemes of directory-based systems: centralized directory, distributed flat directory, distributed hierarchical directory.

Central Directory Architecture (CD). CD architecture is based on a central directory that aggregates information from every provider, and respond to queries from every client. Central directory architecture is a simple solution, easy to administrate, but directory can represent a bottleneck and a single point of failure, which causes the whole system's failure. Therefore, this solution does not scale well and is only suitable for small networks. Some examples of discovery mechanisms based on centralized architecture are the Service Location Protocol (SLP) [13] standardized by the Internet Engineering Task Force (IETF, RFC 2608), the Jini system [14], which is a platform-independent service discovery mechanism based on Java and developed by Sun Microsystems, and the Agent-Based Service Discovery proposal [15].

Distributed Flat Directory Architecture (DFD). In DFD architecture several directories cooperate in a peer-to-peer fashion, to maintain a distributed repository of information about resources and services. Flat distributed directories can work in two different ways. Directories can exchange information with all other directories, usually by multicasting, so that each directory maintains a complete database about all resources and services in the network. The Intentional Naming Service (INS) [16] and the Salutation protocol [17] are two examples of discovery mechanisms based on this technique. It is obvious that this solution generates high communication traffic level, and hence it is not scalable. The second alternative is to divide the network in clusters or domains, so that each directory maintains a repository with information about services and resources within the cluster or domain. Information exchange between directories in different clusters can be achieved using a peer-to-peer scheme, but using a lower advertising frequency than within the cluster, like for example the INS/Twine system [18], or can be achieved on-demand, like for example the service locating system based on Virtual Backbone [19]. Although clustered solutions are more scalable and suitable for large networks, they must use complex algorithms to manage clusters (cluster formation, selection of directories, addition and removal of nodes to/from the cluster, etc.), and guarantee cluster stability.

Distributed Hierarchical Directory Architecture (DHD). With DHD architecture, the network is divided in domains with a hierarchical structure (like DNS) and directories have parent and child relationship. This solution is fully scalable, but it enforces a rigid hierarchical network organization, which does not fit well in ad-hoc environments. Some examples of distributed hierarchical

directory architectures are the Monitoring and Discovery Service (MDS) used in Globus [20] [21] and the Secure Service Discovery Service (SDS) developed at UC Berkeley [22].

3 A Hybrid Mechanism for Resource Discovery in Ad-Hoc Grids

In view of the advantages and limitations of the existing discovery architectures, summarized in table 1, we propose a hybrid discovery mechanism, which combines the advantages of peer-to-peer mechanisms (high adaptability for changing conditions, and low management complexity), and the advantage of clustered solutions (high scalability).

This hybrid approach is based on the idea of zone, similar to the concept introduced by the Zone Routing Protocol (ZRP) for ad-hoc networks [23]. A **discovery zone** is defined for each grid node individually, and is composed by all the neighbor nodes whose distance to the node in question does not exceed a certain number of hops, R , where R is the **zone radius**. It is obvious that the discovery zone of neighbor nodes can overlap.

Within the discovery zone of a given node, we can distinguish two kinds of nodes: the interior nodes, whose distance to the central node is lower than R ; and the peripheral nodes, whose distance to the central node is exactly equal to R . Example in Fig. 1 shows the discovery zone of node node A with $R=2$.

The resource discovery mechanism uses a mixed peer-to-peer approach: to discover grid nodes within the zone it uses an advertisement mechanism, and to

Table 1. Main features of discovery mechanisms

	P2P-Query	P2P-Adv	CD	DFD	DHD
Suitability for changing conditions	High	High	Low	Low	Low
Scalability	Low	Low	Low	High	High
Bandwidth consumption	Medium	High	Low	Low	Low
Discovery delay	High	Low	Low	Low	Low
Management complexity	Low	Low	Medium	High	High

Fig. 1. Example of discovery zone for node A, with $R=2$

discover grid nodes out of the zone it uses a query mechanism. Each grid node periodically multicasts advertisement packets with a hop limit of R hops, so that these packets only reach those nodes within the discovery zone. Using this mechanism, every node constructs a database with detailed information about all the neighbors within its zone. If no advertisement messages are received from a given neighbor within a specific period, this node is removed from the database. This restricted multicast technique reduces the bandwidth consumption and provides a low delay mechanism for discovering grid nodes within the zone.

If the number of resources within the discovery zone is not enough to meet the client application requirements, a query mechanism is initiated. In this case, the client's node sends a query message to the peripheral nodes, to obtain information about the grid nodes existing in the adjacent zones. This procedure can be repeated several times by the client's node, to obtain information about grid nodes existing two zones away, three zones away, etc., until the number of discovered resources is enough, or until a maximum discovery delay is exceeded. To implement this behavior, each query message includes a parameter called forwarding distance, which specifies how many times the message must be forwarded by peripheral nodes to the next adjacent peripheral nodes. Figure 2 shows how a query message with Forwarding Distance = 1 is forwarded to peripheral nodes of the client's zone, and the query message with Forwarding Distance = 2 is forwarded to peripheral nodes of adjacent zones.

Fig. 2. forwarding of query messages

The three main messages involved in this discovery mechanism are the **Advertisement message**, which is used by a grid resource to multicast its presence and characteristics to the rest of nodes within its discovery zone. This message can contain static and dynamic information about the resource (CPU type and architecture, CPU count, processor load, OS, total and free memory, total and free disk space, bandwidth network links, software, services, etc.). Advertisement procedure is controlled by two main parameters: the Advertisement Period and

the TTL period. The Advertisement Period specifies how often a grid node multicasts an Advertisement message to the discovery zone. The TTL period specifies how long a node should keep the information advertised by a neighbor, if no Advertisement messages are received from it. The **Query Request message** is sent by the client node to the peripheral nodes to discover resources out of the client's discovery zone. This message must contain the Forwarding Distance parameter, and the client application requirements, i.e., a list of static or dynamic characteristics that the remote grid nodes should meet. During the discovery process, the client node can send different query messages with increasing values of Forwarding Distance, to discover nodes further away. Finally, the **Query Response message** is used by the peripheral nodes to return to the client node a list of resources that meet the user requirements.

The hybrid method proposed is scalable, since multicast advertisement messages are restricted to the client's zone, and query messages do not use flooding, but they are propagated only by peripheral nodes of successive neighboring zones. Furthermore, discovering delays are much lower than pure peer-to-peer query mechanisms, since a peripheral node can provide information about all grid nodes within its zone. This mechanism is very suitable for changing environments, since information in node databases is updated automatically by the advertisement procedure, and it does not require any administration or management effort.

4 Results

Figure 3 shows the results of number of messages (bandwidth consumption) and discovery delay for the ad hoc network in Figure 2, using different discovery mechanisms: P2P-Query mechanism, P2P-Adv mechanism and the proposed hybrid mechanism with zone radius $R=1$ and $R=2$.

The number of messages includes all the messages (advertisement, query request, and query response) that the different mechanisms use to discover all the available resources in the network. For simplicity reasons, the delay for query request/response messages is given in generic time units, and it is computed by assuming that the propagation delay of every link is equivalent to 1 time unit, and the processing time of every query request message is equivalent to 2 time units.

We can observe that the discovery delay of P2P-Adv mechanisms is zero, since every node in the network maintains its own complete database with information about all the resources. However, because this mechanism is based on broadcasting, the number of messages (and hence the bandwidth consumption) is extremely high. On the other hand, the number of messages of the P2P-Query mechanisms is very much lower, but the discovery delay increases significantly. In the middle of these two extremes, the hybrid mechanism exhibits a good trade-off between these two parameters, since it can reduce appreciably the discovery delay, maintaining a low bandwidth consumption.

Fig. 3. Bandwidth consumption and discovery delay results

5 Conclusions and Future Work

Efficient resource discovery is a major challenge in ad-hoc grids. Most of the existing mechanisms for resource and service discovery can be classified in peer-to-peer architectures, and directory-based architectures. While peer-to-peer architectures do not scale well, directory-based systems are too rigid and could not be suitable for mobile ad-hoc environments. In this paper we propose a hybrid resource discovery mechanism that is based in the concept of discovery zone. Although it uses peer-to-peer communication, multicasting is restricted to the discovery zone, and queries are forwarded by peripheral nodes, avoiding flooding. This mechanism is scalable, exhibits low discovery delays, is adaptable to changing conditions, and does not require any management effort.

As future work we plan to introduce query control mechanisms, which try to avoid that a given node could forward the same query request several times, and to prevent query requests from being forwarded to zones already visited.

References

1. I. Foster, C. Kesselman, and S. Tuecke, *The Anatomy of the Grid: Enabling Scalable Virtual Organizations*, Proc. Euro-Par 2001 Parallel Processing, LCNS 2150, Springer-Verlag, 2001, pp. 1-4.
2. M. Gaynor, M. Welsh, S. Moulton, A. Rowan, E. LaCombe, J. Wynne, *Integrating Wireless Sensor Networks with the Grid*, IEEE Internet Computing, special issue on the wireless grid, July/Aug 2004, pp. 32-39.
3. L.W. McKnight, J. Howison, S. Bradner, *Wireless Grids: Distributed Resource Sharing by Mobile, Nomadic, and Fixed Devices*. IEEE Internet Computing, special issue on the wireless grid, Jul./Aug. 2004, pp. 24-31.
4. D.C. Marinescu, G.M. Marinescu, Y. Ji, L. Boloni, H.J. Siegel. *Ad Hoc Grids: Communication and Computing in a Power Constrained Environment*. Proc. 22nd Int. Performance, Computing, and Communications Conf. (IPCCC) 2003, pp. 113-122.

5. K. Amin, G. Laszewski, A.R. Mikler, Toward an Architecture for Ad Hoc Grids, Proc. of the IEEE 12th Int. Conf. on Advanced Computing and Communications, ADCOM 2004.
6. K. Amin, G. Laszewski, M. Sosonkin, A.R. Mikler, M. Hategan, Ad Hoc Grid Security Infrastructure. Proc. of the 6th IEEE/ACM Int. Workshop on Grid Computing, pp. 69- 76, 2005.
7. Z. Li, L. Sun, E.C. Ifeachor, Challenges of Mobile ad-hoc Grids and their Applications in e-Healthcare. Proc. of the 2nd Int. Conf. on Computational Intelligence in Medicine and Healthcare, CIMED 2005.
8. Bluetooth, Service discovery Protocol, Bluetooth Specification Version 1.1, Part E, Feb. 2001.
9. R. Koodli, C. Perkins, Service discovery in on-demand ad hoc networks, IETF Internet Draft (draft-koodli-manet-servicediscovery-00.txt), October 2002.
10. Christian Frank, Holger Karl, Consistency challenges of service discovery in mobile ad hoc networks, Proc. of the 7th ACM Int. symposium on modelling, analysis and simulation of wireless and mobile systems (MSWiM), 2004, pp. 105-114.
11. S. Helal, N. Desai, V. Verma, C. Lee, Konark - A Service Discovery and Delivery Protocol for Ad-hoc Networks, Proc. of the Third IEEE Conference on Wireless Communication Networks (WCNC), New Orleans, March 2003
12. B.A. Miller, T. Nixon, C. Tai, M. D. Wood, Home Networking with Universal Plug and Play, IEEE Communications Magazine, vol. 39, no. 12 (2001), pp. 104-109.
13. E. Guttman, Service Location Protocol: Automatic Discovery of IP Network Services, IEEE Internet Computing. Vol. 3, No. 4 (1999), pp. 71-80
14. Sun Microsystems, Jini. Architecture Specification, Technical Report, version 1.2, 2001.
15. Y. Lu , A. Karmouch, M.Ahmed, R.Impey, Agent-Based Service Discovery in Ad-Hoc Networks, 22nd B. Symp. on Communications, Kingston, Ontario, Canada, 2004.
16. W. Adjie-Winoto, E. Schwartz, H. Balakrishnan, J. Lilley, The design and implementation of an intentional naming system, Proc. of the 17th ACM Symposium on Operating Systems Principles, 34(5), 1999, pp. 186-201.
17. The Salutation Consortium, Salutation Architecture Specification, Technical Report, version 2.0c, 1999.
18. M. Balazinska, H. Balakrishnan, D. Karger, INS/Twine: A Scalable Peer-to-Peer architecture for Intentional Resource Discovery, Int. Conf. on Pervasive Computing, Zurich, Switzerland, 2002.
19. J. Liu, Q. Zhang, W. Zhu, B. Li, Service Locating for Large-Scale Mobile Ad Hoc Network, International Journal of Wireless Information Networks, Vol. 10, No. 1 (2003), pp. 33-40.
20. The Globus Alliance. Globus Toolkit 2.2 MDS Technology Brief, Draft 4, http://www.globus.org/toolkit/docs/2.4/mds/mdstechnologybrief_draft4.pdf, 2003.
21. The Globus Alliance, MDS Functionality in GT3. OGSA Technical Resources, <http://www.globus.org/toolkit/docs/3.0/mds/MDS.html>, 2003.
22. S. Czerwinski, B. Zhao, T. Hodes, A. Joseph, R. Katz, An Architecture for a secure Service Discovery Service, Proc. of the ACM/IEEE MOBICOM, 1999, pp. 24-35.
23. Z.J. Haas, M.R. Pearlman, The performance of query control schemes for the zone routing protocol, IEEE/ACM Transactions on Networking, vol. 9, no. 4 (2001), pp. 427-438.