Organization of Grid Resources in Communities

Asif Akram, CCLRC e-Science Centre, CCLRC Daresbury Laboratory, Warrington, UK, WA4 4AD; E-mail: a.akram@dl.ac.uk
Rob Allan, CCLRC e-Science Centre, CCLRC Daresbury Laboratory, Warrington, UK, WA4 4AD

ABSTRACT

Locating suitable resources within a Grid is a computationally intensive process. An alternative approach is to categorize resources based on the services they provide – leading to the interaction of peers with common goals to form societies/communities. The communities can be adaptive in nature and evolve based on changes in their operating. We have implemented JXTA prototype to illustrate the concepts of community formation in which Peers offering different services can be grouped together based on different criteria.

1. INTRODUCTION

Emerging distributed computing paradigms, such as Grid Computing, comprise of dynamic and distributed resources/peers. Resource discovery in such systems is a time-consuming process with network overhead. The numbers of interactions are likely to increase exponentially as the numbers of peers grow. Restricting interactions between set of peers is a key factor to scale the resource discovery problem. Davis and Smith referred it as the "connection problem" [1], where peers need to find other suitable peers to co-operate with, assist, or interact with. Peers are categorized based on criteria i.e. type and quality of service, etc. Any initial cost in categorizing peers result discovering "preferable" resource with minimum discovery cost subsequently – thereby leading to the development of "communities". "Focused Addressing" [2] is one solution to the connection problem where requests are sent to particular subset of peers, believed to assist the requesting peer.

Individual peers, although selfish, are expected to interact with each other in some way. Each peer prefers to be in environment where it may be easily discovered by a suitable user, and can locate other peers with minimum efforts. This analogy helps us to define two terms, Expertise and Interest [13], [14]. Expertise of a peer is the basic service provided by that peer and Interest of a peer is the service/services provided by other peers which are supportive to its main service. Each community has one Service Peer with dual responsibility of managing the member peers and providing communication source with external environment. Interaction between different communities is only through the Service Peers.

2. TYPE OF COMMUNITIES

Individual autonomous peers have expertise and interests in specific resource/s. Based on these expertise and interests, peers are grouped together, but expertise and interests are not the only criteria for categorizing peers. Communities/societies can be of different types as mentioned below:

Competing Community: In a Competing Community each peer has the same expertise – although some service attributes may vary. Similarity in services may develop competition amongst member peers, as members compete each other to get selected.

Co-Operative Community: In Co-Operative communities peers provide different services, which must be used alongside services of other members. Hence, when any peer is selected, then the possibility of selection of another member peer providing utility service/s increases. This mutual co-operation is suitable for peers which provide simple services.

Goal Oriented Community: This is collection of peers work together to achieve a particular goal. Goal oriented communities are important in self-organizing systems, where interactions between member peers are not pre-defined, but the services required are. In such instances, member peers may interact with each other in arbitrary ways to achieve a given end result.

Ad Hoc Community: In ad hoc communities peers interact directly with each other without interference and involvement of a Service Peer. Peers belonging to different communities providing supporting services form the basis of an ad hoc community.

Domain-Oriented Community: Such a community is formed by linking together similar-minded organizations and institutions, instead of the services they provide, such as academic communities, research communities, and open-source communities. Hence these communities are domain-oriented rather than service-oriented.

Virtual Community: The Virtual community is a community of communities. This effect is achieved by leasing out the member Peer to other community for certain time period, before that lease period either Service Peer requests to renew the lease of corresponding Peer or it can't use the service of the Peer directly.

Sharing Community: In this type of community different Communities share their resources with each other; this sharing of resources is not restricted to member Peers but includes core and optional services. Community A may have QoS monitoring module, which it shares with Community B assuming either Community B doesn't have such module or Community A may has more advanced monitoring module.

3. COMMUNITY FORMATION

Community formation involves three basic processes and in all of them the middleware plays a major role; utilizing different core services provided by the community. These three processes are "Initiation Process", "Maintenance Process", and "Formation Process" for collaboration on an ongoing basis between the members and with other communities.

3.1 Initiation Process

The community formation process can be initiated manually or by any peer. The first task of the newly created peer is to search for existing communities which may have interest in its services. The peer queries the middleware for appropriate communities; on failure it may request to create a new community. The middleware searches its own knowledge base to match the services of the peer with requirements of existing communities. The peer can refuse to join any existing community and can insist for the creating of new community. The newly formed created community (i.e. manually or on the request of any peer) search for other peers and services with complementary and competitive resources for long term partnership. The community search for new members by querying the peers registered within the middleware. The main achievement of this process is to promote mutual trust between the members, negotiate consistent rules and regulations i.e. pre-qualification criteria for new entrants, Quality of Service (QoS) [11], sharing costs and profits. The community formation process either creates new Service Peer or requesting Peer acts as coordinator.

3.2 Maintenance Process

After the initial formation of community, the coordinator is responsible for maintaining and improving the collaboration, which involves many different sub-management tasks, conducted concurrently. The coordinator observes the changes in the internal and external environment and adapts according to the changes (updating its knowledge base) to maintain its effectiveness in the distributed environment. The coordinator monitors the performance within the community and provides basic services to members to achieve the required performance. It

is important to keep track of the resources and core competencies of partners and their performance. To improve overall performance of the community, its coordinator is always looking for new members with missing or complementary resources. Although the maintenance process is mainly concerned with the local functioning of the community but in this stage the coordinator also communicates with the middleware to search for new members.

3.3 Formation Process

The community is opportunist and always exploits the potential for collaboration. When a new collaboration opportunity arises it is the coordinator identifies the required individual activities to match the opportunity. The first step for the coordinator is to create a "workflow" based on the available resources, competencies, strengths and weaknesses of the members. The workflow may require collaboration with external communities to "buy in" services missing within the community. No matter how complicated final workflow is; whether it utilizes only internal resources or involves an external collaboration, coordinator develops teamwork for achieving set goals. Once the workflow is formed and tasks are allocated and scheduled among participating members, the coordinator monitors the performance of individual members and quality of tasks. Post-Management formation process involves allocation of additional resources, replacement of under performing member/s, re-scheduling of tasks within the workflow, etc.

4. ARCHITECTURE OF TOOLKIT

The architecture for the community formation should be simple and supportive to the main purpose of their formation. The proposed architecture consists of three main components. These components are discussed following the Tool Selection:

4.1 Tool Selection

The efforts to design and implementation the system can be drastically reduced by selecting appropriate technology. The initial prototype is developed in JXTA [11]. JXTA (jxta.org) is an open source P2P framework initiated by Sun Microsystems. The JXTA protocols are independent of any programming language, and multiple implementations exist for different environments which make it best choice for prototype. The JXTA network consists of a series of interconnected nodes, or Peers. A JXTA Peers is "any entity capable of performing some useful work and communicating the results of that work to another entity over a network". Peers can self-organize into Peers Groups, which provide a common set of services.

JXTA has the concept of Peer and Peer Group which match to our vision of '*Peer*' and '*Communities*', which makes JXTA as best choice for implementing our prototype.

4.2 Middleware

JXTA had default peer group and every peer joins this default peer group after booting within the JXTA network. This default peer group is called the *NetPeerGroup* or *WorldPeerGroup*. The middleware is the extension of default peer group.

The middleware required to support the community formation is more than simple distributed registry. The customized middleware provides more specific search capabilities and match making. It provides the interface to create new community. The toolkit does not impose any restriction on the communities i.e. the nature or role of communities, the services they offer, why and when these communities are created. The framework supports the creation of communities and the definition of membership policy. It is up to cooperating peers to define communities, join communities, and leave communities.

The middleware provides the much required environment to peer and community for advertising their capabilities. All queries and match making is done against these advertisements.

4.3 Peer

A Peer can be simple service, resource on the computer or any hand held device. Each peer operates independently and asynchronously from all other peers, and is uniquely identified by a Peer ID. All peers are automatically members of the framework which extends the default NetPeerGroup. Peers may opt to join and leave customized or user groups/communities at will. In order to join any community, a peer must discover the community through the search capabilities provided by the middleware. Once the suitable communities are discovered then peer apply for the membership.

Peers publish one or more interfaces for its different services and resources. Each published interface is advertised as a peer endpoint, which uniquely identifies the network interface. Peer endpoints are used by other peers to establish direct point-to-point connections between two peers.

4.4 Community

Community is a temporary or permanent coalition of geographically dispersed individuals, group's organizational units or entire organizations that pool resources, capabilities and information to achieve common objectives. Each community has a similar architecture, with one Service Peer, which manages the whole community. Normally Peers from different communities are not allowed to communicate directly and the communication should be done through the Service Peer except for the Ad-Hoc community.

A community offer different services and resources; utilized only by the member peers. The peer looking for specific service should locate and join the appropriate community offering required service.

Communities may strongly enforce a membership requirement. This defines the boundaries for a secure environment where content can be accessed only by member peers. Communities can provide services and thus participate as a single entity in the formation of further communities.

4.5 Service Peer

The Service Peer is a community coordinator, which manages the proper working of the community, provides essential resources to its members and is source of communication with other communities. The Service Peer is the extension of the Rendezvous Peer provided by JXTA environment and has an optimized routing mechanism for an efficient propagation of messages. If the Service Peer cannot locate the destination of message or query, the request will be forwarded to other known Service Peers.

Service Peer maintains a local view of the environment, a list of known Service Peers and communities. Service Peers maintain information of member peers and a restricted set of other communities; this interest is governed by the expertise and interest of other communities. This restricted list of other communities is used to develop referral mechanism.

The Service Peer coordinates the resources and services within the group to achieve set objectives, maintain membership policies, monitors member peers. The Service Peer may not itself offer different services but may only coordinate services offered by different specialized member peers.

5. PROTOTYPE AND RESULTS

For simulation purposes the JXTA prototype has been developed with option for creating Groups and Peers along with their description. This description is used as the one of the membership criteria, when any Peer applies for the membership. The community is created along with randomly generated External Rating and community assigns randomly generated Internal Rating to all of its members at the time of membership. Peers apply for membership based on high external rating of the community. The community grants membership based on overall rating of the peer and description of the peer.

Prototype was evaluated with different set of parameters i.e. maximum number of member Peers in a community, maximum number of communities joined by single peer etc. Evaluation results were quite encouraging, and similar pattern was observed by changing the set of parameters. In the beginning of the evaluation following four steps were quite frequent:

- Selection of communities by peers
- Compatibility checks i.e. rating and description by peers.
- Request for membership by peers
- Membership confirmation from the community.

The preliminary result with different set of constraints is shown below in the tabular and graphical form:

In the table above the left column 'Groups Discovered' compares with the ratio of total membership requests and the accepted requests. It is obvious from the

Table 1. Membership acceptance-request ratio

Groups Discovered	Membership Acceptance/Request Ratio		
	Group Size (50)	Group Size (80)	Group Size (100)
0 - 100	80/100 = 0.8	100/100 = 1.0	100/100 = 1.0
100 - 200	60/80 = 0.75	76/90 = 0.85	90/100 = 0.9
200 – 300	30/50 = 0.6	50/75 = 0.66	63/85 = 0.74
300 - 400	10/25 = 0.4	28/50 = 0.45	28/55 = 0.5
400 - 500	0/5 = 0	3/18 = 0.16	5/21 = 0.23

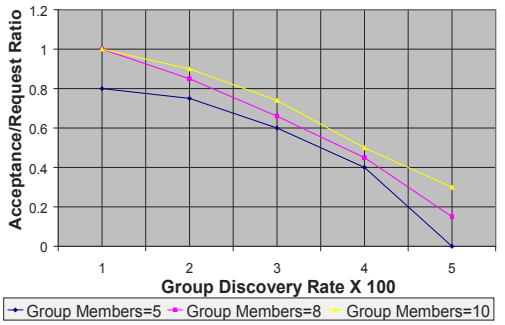

above table initially peer applies for most of discovered communities but this trend decreases when more communities are discovered. With the passage of time frequency of request for membership decreases and the rate of membership granted by communities decreases because either all member peers have better rating or their expertise matches the interest of community. Once System becomes stable then even peers don't apply for membership or changes group.

The time required by the environment to achieve stable state depends on the constraints set by system i.e. rate of community discovery, number of member peers in a single community and number of communities joined by single peer. We are confidant that organizing resources into different communities will give new dimension to Grid Computing.

6. CONCLUSION AND SUMMARY

In this paper, we have presented the concept of categorizing peers in communities on the basis of their expertise and interests. Social networks are a natural way for people to go about seeking information. Organizing peers in one form or another makes the discovery of resources efficient, whilst minimizing computational overheads. Categorizing the peers in communities is simple, open and easy to implement, and the initial overhead of developing communities pays-off latter at the time of resource discovery. Communities are more stable, and stability increases with the passage of time, communities have a simple learning time and are more adaptive to operate in a dynamic environment. We have proposed the external and internal rating for communities and peers respectively which may be used to support a given Quality of Service, effective participation of autonomous peers and better interaction among communities and member peers. Finally, we discuss the different services required to manage the group and requirements of the member peers. A JXTA implementation of a prototype system is discussed to describe the salient features of our approach. A key theme of this work is to determine how communities should be structured to support resource discovery, and how particular roles within a community can be used to determine interactions between participants within a community, and those between participants across community. This work extends techniques and results discussed in [12].

7. REFERENCES

[1] Davis, R. and R. G. Smith, *Negotiation as a Metaphor fro Distributed Problem Solving*. Artificial Intelligence 20, 63-109, 1983.

[2] Parunak, H. V. D., *Distributed Artificial Intelligence*, Chapt. Manufacturing Experionce With the Contract Net, pp. 285-310, Research Notes in Artificial Intelligence. Los Altos, CA: Morgan Kaufmann Publishers, 1987.

[3] P2People http://www.iti.gr/db.php/en/projects/P2PEOPLE.html

[4] Kautz, Henry, Milewski, Al, and Selman Bart, *Agent Amplified Communication*. AAAI '95 Spring Symposium Workshop Notes on Information Gathering in Distributed, Heterogeneous Environments, Stanford, CA.

[5] Kuokka, Daniel, and Harada, Larry, *Matchmaking for Information Agents"*. Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI) '95, 1995.

[6] Leonard Foner: *Clustering and Information Sharing in an Ecology of Cooperating Agents, or How to Gossip without Spilling the Beans*. Conference on Computers, Freedom and Privacy, 1995.

[7] UP2P http://chat.carleton.ca/~narthorn/project/UP2POverview.html

[8] Catalin Dumitrescu (UChicago), Ian Foster (UChicago/Argonne), Mike Wilde (UChicago/Argonne), *Policy-based Resource Allocation for Virtual Organizations.* iVDGL/GriPhyN Presentation, Fall 2002

[9] Omer Rana, Asif Akram, Rashid Al-Ali, David W.Walker, Gregor von Laszewski, Kaizar Amin, *Quality of Services Based Grid Communities. Web Services and Agent Systems*. Kluwer Academic (2004)

[10] GNUTELLA. Gnutella Home. http://www.gnutella.com/

[11] JABBER SOFTWARE FOUNDATION. Jabber IM. http://www.jabber.org/about/overview.html

[12]Steven Lynden and Omer Rana, *Coordinated Learning to support Resource Management in Computational Grids* Second International Conference on Peer-to-Peer Computing, (2002) 81-89

[13]Leonard Foner. Yenta: *A multi-agent, referral-based matchmaking system*. In Proceedings of the 1st International Conference on Autonomous Agents, (1997) 301–307.

[14]Bin Yu and Munindar P. Singh: *Searching Social Networks.* Proceedings of Second International Joint Conference on Autonomous Agents and Multi-Agent Systems, 2003.