

P2P Architecture for Scientific Collaboration

José Mitre, Leandro Navarro-Moldes
Polytechnic University of Catalonia, Spain
{jmitre, leandro}@ac.upc.es

Abstract

P2P networks are often associated with file exchange applications among private users. However, their features make them suitable for other uses. In this paper we present a P2P architecture for Scientific Collaboration Networks, which takes advantage of the properties inherent in these social networks –small-world, clustering, community structure, assortative mixing, preferential attachment and small and stable groups– in order to obtain better performance, efficient use of resources and system resilience.

1. Introduction

Many P2P applications are created for sharing any kind of files, and therefore they are not optimized for any particular community. Therefore, they don't take gain from the user behavior patterns in specific communities.

In this work we direct our attention to Scientific Collaboration Networks (SCN), which have natural topologic properties that make them quite well structured and dissimilar in several aspects from traditional file sharing communities using P2P applications.

Along with other authors [1,2], we claim that the properties found in social networks have a great influence on the way these networks operate and that these properties offer valuable information to be exploited by distributed and P2P applications.

Here we present a P2P architecture for scientific collaboration which takes advantage of the properties of the social network formed by scientists working together, trying at the same time to preserve the advantages and simplicity of general purpose unstructured and structured (e.g. DHT based) P2P networks.

The paper is organized as follows: section 2 shows SCN features and patterns. Section 3 describes how our architecture takes advantage of these features. In section 4 we compare our proposal with some works focused on collaboration. Section 5 presents our conclusions.

2. SCN features

In the last fifty years, many studies have been aimed at finding the properties of social networks and SCN. These efforts arose not only from the interest inherent in

patterns of human interaction, but also from the structure of the resulting network, which has important implications, for example, in the dissemination of information.

A social network is a group of people, known as the actors, linked together by some pattern of interaction. These networks are normally represented by graphs called sociograms. A sociogram is a set of points (or vertices) denoting people, joined by lines (or edges) denoting links.

Then, a SCN is a social network with scientists as actors, where links denote collaboration; e.g. working together on a paper during a certain period of time [3,4,5]. In contrast with other works, we consider that links are not only formed by co-authorship, but also by many other types of relations, since a large part of scientific communication occurs in private conversations; not just the actual write-up process, but extends to the whole research process.

Next we describe some useful properties and patterns of social collaboration networks, which will be used throughout this paper.

Assortative Mixing: A social network is said to show assortative mixing if in that network the people wishing to associate with others have something in common. Newman [6] demonstrates that assortative mixing is present in many networks and shows how their effect on the structure and behavior of the network can be measured and examined.

Preferential Attachment: The majority of real networks describe open systems in which growth exists due to the constant addition of new nodes. Beginning with a small nucleus made up of a group of nodes, the number of nodes increases throughout the life-span of the network due to the subsequent addition of new nodes. In the majority of social networks, the addition of new nodes occurs by preferential connectivity [4,7], in such a way that the new nodes are connected to other nodes by preference; to nodes with a greater degree, for example, or to those that are most popular.

Degree Distribution: many real social networks, including SCN, the degree distribution follows a power law [4], which indicates a heterogeneous topology in which the majority of nodes have a small degree and a small fraction of highly connected nodes, unlike classical P2P networks, being random networks, are statistically homogeneous, with a Poisson degree distribution.

Community Structure: [6,8,9] Is the property of many social networks for forming communities through the union of people in groups. Grouping occurs for many reasons: shared interest, working for the same company, geographical proximity, etc. In many social networks it is possible for people of a similar type to be drawn together and then to divide up naturally into groups of a particular type, so that the density of links within the group is greater than the density of the links among them [6]. In computer networks where collaboration exists, the association between people may be given not only by geographic proximity but also by proximity of subject matter, for example.

Clustering: the probability of two people meeting if they share one or more mutual acquaintances. For example, in SCN scientists tend to introduce their collaborators to scientists belonging to other groups, thus fomenting new collaboration and thereby increasing the clustering coefficient [7].

Affiliation Networks: is [10] a network in which the actors are joined together by common membership of groups or clubs of some kind. Some studies show collaboration groups of academics, actors and business people as affiliation networks. We could say that SCN are also affiliation networks because in the majority of cases, in order to belong to a research group some kind of invitation or official application is necessary for the collaboration to take place. In real life, for example, it is very common to see scientists joining interest groups such as professional associations like IEEE or ACM. In our architecture we exploit this property in the design of a connection mechanism (section 3.1) to support grouping of scientists by affiliation into interest groups.

The above properties clearly show the important role of interest groups in SCN. This enables us to assume that in most cases the relevant information for a group of scientists making up a group will be found within the limits of that group; that is, if a scientist is looking for information it is highly probable that such information can be located within the group to which the scientist belongs.

Small World: SCN form “Small Worlds”. Typically, participants are separated by short paths [4] of known intermediates. Clearly, news of important findings and scientific information can circulate more quickly in a network where scientists are more closely connected.

SCN have useful patterns of behavior in the number and selection of collaborators, and in network robustness.

The number of collaborators of a scientist is small compared with other social networks [3]. This is perhaps due to the fact that scientists often collaborate with the same group of colleagues rather than writing each paper with different or new scientists (high clustering coefficient, long term links), although there may be additional undiscovered and beneficial links between any

two members of the same community (i.e. weak links [21]) that a collaborative application could exploit.

Given the probability distribution of links between any two given scientists, the probability of having a direct link is much greater if the two scientists in question share other mutual colleagues, or if they both have common interests. This property contributes to make SCN highly “clustered”, and therefore it helps to define clear boundaries between groups of interest.

Therefore we propose a search mechanism based on scoped flooding within groups (small worlds, high clustering coefficient, and discovery of weak links). This mechanism is scalable enough and imposes a low network overhead given that search messages will not reach the whole network: they will only appear within the group boundary, as explained in section 3.3.

Because of their high degree of robustness [11], SCN are fault-tolerant; that is, network operation is not interrupted by failure or disconnection of subsets of network nodes. Nevertheless, they are extremely vulnerable to directed attacks, since there is a danger that highly connected nodes can fail. Object replication among neighbor nodes on a small-network topology reduces this vulnerability (section 3.2).

As one may see, many of these features can be used to either incorporate improvements in existing generic P2P networks, or create new ad-hoc P2P networks specific for collaboration groups. Our intention was to design a SCN-aware (specific) P2P architecture, with the aim of facilitating collaboration among scientists who might be geographically far apart, so they could exchange information between each other without having to pass through any dedicated server.

The architecture proposed provides the basis for a modular collaborative application. It is currently being implemented in Java with JXTA [20]. The resulting application will support scientific collaboration activities such as subscription to scientific communities, paper production, diffusion of awareness information on related publications, call for papers, events, news, etc.

3. P2P architecture for SCN

Based on the previous properties and patterns, we describe a P2P architecture for SCN.

We will call *SCN* a network of computers using our P2P model in which the scientists collaborate.

A *servent* is a computer connected to an *SCN*. The *servents* will provide interfaces by means of which scientists can exchange messages, share information, carry out searches, compare data, and other collaborative tasks.

Two *servents* are *neighbors* if they have a mutual link generated by a direct collaboration.

A *group* is a sub-network of the *SCN* formed by the *servents* associated to scientists who share interest in

common topics. We claim that new *servents* joining a group must follow the same rules of behavior as in real life; that is, by affiliation or invitation of *servents* to a given group. Thus the *SCN* topology would be similar to the topological structure of the social network (affiliation networks, preferential attachment) of scientists working together in real life.

A *member* is a *servent* belonging to a given group. All the *servents* should, by default, be *members* of at least one *group*. Disconnected *servents* will continue to be *members* unless they explicitly withdraw.

Every *servent* holds a list of known groups (GroupId List) and a list of group members identified by their ServentId. These lists may be incomplete, and they are kept consistent using an epidemic consistency algorithm such as the TSAE weak consistency algorithm [12].

Our architecture has three fundamental mechanisms for carrying out cooperation functions: 1) connection and join, 2) replication, and 3) search.

3.1. Connection and join mechanism

In many cooperation networks, users connect and disconnect from the network several times a day, in some cases only once a day, depending on the activity of the group. It is therefore easy to see that a cooperation network must have mechanisms that manage the connection and disconnection of nodes from the network. From the foregoing and for simplicity, we distinguish two types of connection: joining and connection.

3.1.1. Joining. Is the first connection to a group and it is a special case, since then the new node must obtain membership information from the group. In order for the topology generated by our architecture to maintain the same properties as those of a real *SCN*, e.g. Clustering and small-world, the *servents* must have means of connecting to *groups* that are similar to those used in real life (affiliation network). We therefore assume that when a new *servent* joins a *group* it should be by invitation or by application from the new *servent* to the *group*. For a *servent* to be connected to the *SCN* for the first time, the scientist must either establish contact with an existing group or create a new group. The *servent* must provide a suitable interface in order to carry out both operations.

When a *servent* creates a new *group*, he generates a unique GroupId that identifies the *group* throughout the *SCN*. The *servent* has a ServentId that identifies him, adding to the ServentId list of the group and sending a message to a number of *servents* of the other groups, using an epidemic dissemination algorithm (TSAE), to notify them of his existence. They will feed the initiating *servent* in turn with information about the existing groups in the whole network.

If a *servent* wishes to join to one or more existing groups, he first receives authorization from any *member*

of that group and receives the potentially incomplete group's ServentId list. Once the scientist has chosen the group to which he wishes to belong, the *servent* sends a message to any *member* of that group to apply to such group. If a *member* accepts the application to join, the ServentId of the new *servent* is added to the ServentId list of each member of the group, using the epidemic algorithm to spread the new ServentId.

If a scientist no longer wishes to belong to a group, he sends a message with his ServentId, via the *servent*, to other (a few neighbors + epidemic propagation) *members* of the group or groups to which he belongs in order to cancel membership. The other *servents* must delete the ServentId from the group's ServentId list. The *member* will then cease to belong to that group.

3.1.2. Connection. This operation is used for any further connection after joining a group and after having been disconnected for some time. When a *servent* is connecting he sends a message to all his *neighbors* (eventually by epidemic propagation, it will be known by all the members of the group) informing them that a connection has taken place. Once the *neighbors* have received the message, they all update their local ServentId list. The connecting *servent* will update his own ServentId list by sending a request to any *neighbor*.

In order to know about potential object changes that may have occurred while he was disconnected, the connecting *servent* will launch a search operation (see Section 3.3) for events that might have taken place during his absence.

When a *servent* is instructed to disconnect from the network, he must immediately inform all (a few *neighbors* + epidemic propagation) connected *members* that he is about to leave the network, in order to keep the ServentId list up-to-date. In the case of some connection failure, if a *servent* sending a message receives no reply from another *servent*, the sender will assume that some fault has occurred in the connection with the recipient, and will then proceed to update the list of connected *servents*, indicating that a *servent* is not connected, or informing other *members* of the change in the ServentId list by epidemic propagation. In this way the list will eventually be up-to-date.

These mechanisms have similarities with JXTA, however all our *servents* have the same function (in JXTA there are simple, rendezvous and router peers). In addition we include a GroupId and ServentId lists to provide network and *servent* status information.

3.2. Replication mechanism

As previously mentioned (Section 2), given that *SCN* are affiliation networks, that possess a high degree of clustering, and have a community structure, we may assume the existence of groups, and given that these

groups are made up of scientists with common interests, and since the information relevant to a group of scientists can usually be found within their group, we claim that when searching for an object occurs, it is highly probable that the object can be found within a group interested in this object, in few hops (small-world). Object replication will improve object availability, increase system resilience even during directed attacks to high degree *servents*, and it will improve the performance of search operations without overloading the network. Replication could be carried out solely for the members of the group where the object originates, not necessarily for *servents* outside the interest group or even the entire network. This is because these group members will have the greatest interest in the object, since it was generated within their group (assortative mixing).

Actually little investigation [13,16] has been carried out concerning replication mechanisms on P2P networks. These mechanisms extend from a simple path replication like Freenet (replicate along the path from the requester to the provider); and random replication (same number of replicas as in path replication, but replicas are placed randomly among the sites probed) to strategies like uniform replication (all data objects have the same degree of replication), proportional replication (more popular data objects have more replicas) and Square-Root Allocation [13], which is a strategy somewhere in between uniform and proportional replication.

However, each replication strategy introduces an additional cost, not only as to replica storage, but also in the network traffic overload generated by sending data to the peers selected to hold a replica. In [14] we did a cost analysis introduced by replication and we found that this cost depends on the network diameter. Given that the diameter of a group of interest is much smaller than of the entire network, the cost of the replication inside this group is also lower.

We also have seen that the frequency with which objects of interest for a particular group are created and modified is in fact low, and the greater part of communication consists of the exchange of ideas via e-mail or chat which do not need to be replicated. This has been confirmed by the analysis of one year event log for the activity performed by a collaborative social network of people using BSCW, an application for collaborative work support. It shows that the number of reading events is several magnitude orders higher than the number of writing or modification events [15].

Unlike [13], where in order to render searches more efficient, but without taking into account the topological properties of the network, partial replication mechanisms are proposed, we propose making the most of existing community structures to minimize the need to search for objects within a group, since all the objects of interest for that group will be replicated to a number of group

members. In this way, the cost generated by searching (Section 3.3) and replication will be minimized.

We now present a way of managing replication in our *SCN*: Let $g=\{A,B,C,D,E,F\}$ where $A,B...F$ are *servents* belonging to group g , and B,C,D are the *neighbours* of *servent* A . When *servent* A creates a new object or when he modifies an object already existing in the *SCN*, the given object is replicated only to their *neighbours* B , C and D , since they have higher need of that object than any other *servent*, given that B , C and D directly collaborate with A . This proximity replication criteria guarantees that the immediate collaborators will have a replica of the object of interest (assortative mixing). Given that the number of replicas is directly related to the *servent* degree, high degree *servents* will have more replicas, making the *SCN* resistant to failures or directed attacks.

Simulation results confirm that the number of replicas of an object grows quickly with the number of related search operations, and with the degree (number of neighbours) of the originating node which is correlated with popularity.

3.3. Searching mechanism

We have already mentioned that, because of the proposed replication mechanism and the *SCN* capacity for forming small-world communities, the cost of searching based on flooding (number of search messages circulating in the network) is drastically reduced.

In this architecture two types of flooding search can be distinguished – local and external.

3.3.1. Local search. Search undertaken by a *servent* within his group, for three reasons: 1) When a new *servent* joins a group and needs to know about all the objects shared by the group. So when a new *servent* receives a message on concluding the initial connection process (Section 3.1), he asks other *servents* for the objects shared by the group, 2) When a *servent* has reconnected after having been disconnected for a certain time; the *servent* seeks to update information (new or modified objects) generated in his group during his absence, and 3) When a *servent* needs a particular *object*, he will carry out a local search for that object by sending a query to his *neighbors*.

The *servents* receiving the query message will send information about the object to the requesting *servent* if they have the target object. Otherwise they will flood the query to the neighbors, restricted to the group where the query originated.

3.3.2. External search. Search carried out outside the group to which the *servent* initiating the search belongs. This situation may arise when a *servent* needs an object that is not available from any of the group members, and must therefore look for it in other groups throughout the *SCN*. The user must explicitly undertake this kind of

search, restricted to a list of groups of his choice where the user considers the object may be located.

To make searching more efficient than P2P algorithms based on flooding, *servents* have general information (GroupId) about each existing group within the *SCN*. The *servent* initiating the search locates and sends the query to at least one *servent* from each selected group where a local search will be performed, without further propagation to external *servents* (other groups). This differs from flooding algorithms such as Gnutella with super-peers, where queries propagate up to a maximum number of hops (TTL).

Table 1. Relation mechanisms - SCN properties

| Mechanism | Properties | Effect |
|-----------------|--|--|
| Connection | Clustering, assortative mixing, community structure, affiliation networks, preferential attachment | Interest groups joining by affiliation |
| Replication | Clustering, Assortative Mixing, Community Structure, Preferential attachment | Interest-based replication into groups |
| Flooding Search | Clustering, Assortative Mixing, Community Structure, Small world | Topological and thematically Proximity. Small search path length |

Both in local and external search, given that *servents* have object replicas, the probability of locating a desired object through any member of the group where the object originates, will be quite high and search messages will go directly to those *servents* who are most likely in possession of the object.

A summary of relation between the proposed mechanism and SCN properties is shown in Table 1.

It is also worth pointing out that, while our proposal for solving this problem is an initial approximation, and therefore relatively simple, it will be refined by employing additional techniques based on validation results. Initial simulation results on a J-Sim based simulator have been satisfactory in terms of scalability and availability. Using the Newman's algorithm [6] we have simulated networks generated randomly with properties such as: clustering, community structure and small-world. For each topology we run 1,000 differently seeded simulations, consisting of 100 requests (one for each *servent*) for a single object created on a random *servent*.

In each simulation cycle, we randomly designate a *servent* to be the object initiating a search, among those without a replica: at the end of the simulation, every *servent* will have done just one search and will hold one replica.

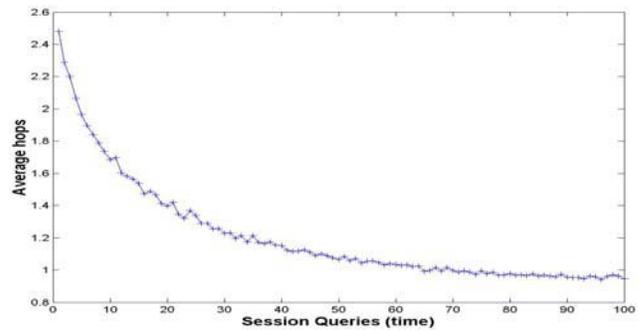


Figure 1. Average number of hops

Our simulation results, see Figure 1, reveal that the greatest long-path length to reach a replica is very small (average values with small dispersion) for interest clusters of 100 *servents*. Roughly after 50% of *servents* have executed a query and have got a replica the long-path to reach a replica is almost 1. Based on other studies [24], the characteristic diameter in Gnutella is smaller than 12 hops and over 95% of the nodes are at most 7 hops away from one another. In our case, the diameter is initially smaller than 3 for each cluster of 100 *servents* and almost 90% of the *servents* are at most 2 hops away (we assume a network of the scale of Gnutella may be composed of many clusters and the cluster size does not change the results significantly). Nevertheless, with less than 3 hops a query can always be resolved and on average in 1.2 hops (4 in Gnutella). In addition, popular objects (high number of searches) are easier to find (more replicas) than non-popular objects (low number of searches). Our search strategy is thus less costly $\sim O(1)$ (objects located in a single jump in most cases) than classical flooding $\sim O(N)$ and can be more flexible than DHT, typically $\sim O(\log N)$.

4. Related work

Iamnitchi et al. [5] put forward ideas about making use of the small-world property and SCN clustering. Some mechanisms are proposed to facilitate searching, but without suggesting any particular protocol. In our work we propose an architecture using replication and considering additional SCN properties.

Other [17,18,19] related work concentrates on identifying clusters of interest to improve the performance of search process so that queries can be steered to peers that are more likely to have an answer.

Unlike these works, we do not identify clusters, given that clusters are formed by users through explicit affiliation with groups (an SCN feature). Cluster identification algorithms could assign a user node to a cluster with only a subset of files of interest. Letting the user select which groups wants to join guarantees he will

be a member of communities of his interest, and have the relevant documents close at hand.

On the other hand, these works leave aside the relevant fact that computer networks give support to social networks with distinctive statistical properties. Our proposal is based on the properties belonging to social collaboration networks, and in particular SCN.

5. Conclusions

A great deal of research work seeks to develop better methods of locating data in P2P networks. These efforts are aimed at improving scalability, greater reliability under dynamic conditions, more efficient searching, and improved performance. The main problem with these systems, however, is that they ignore the fact that computer networks, such as P2P, are made up of people who in turn form social networks with statistical properties which affect the way these networks function.

In this work, we present a proposal for a new community aware P2P architecture for collaboration among scientists using the social network topology, and exploiting inherent characteristics of such networks: small-world, clustering, community structure, assortative mixing, preferential attachment and small and stable groups.

We also show how the combination of replication mechanisms with SCN properties can increase system resilience, and assist in appreciably reducing message overload in the network, compared with overload generated by searching in traditional P2P protocols, thus improving performance.

Although this work is focused on scientific collaboration networks, we believe that many of the ideas set out in this paper can also be applied to the design of other community-specific collaborative networks.

At present, we are working on the implementation of a simulator and a prototype that will enable us to assess improvements in possible optimizations of the architecture.

6. Acknowledgements

This work has been partially supported by the Mexican Institute of Petroleum (Instituto Mexicano del Petróleo) and the Spanish MCYT project TIC2002-04258-C03-01.

7. References

[1] R. Albert, A.L. Barabási, "Statistical mechanics of complex networks". *Rev. Mod. Phys.*, vol. 74, pp. 47-97, 2002.
[2] V. Krebs, "The social Life of Routers, Applying Knowledge of Human Networks to the Design of Computer Networks". *The Internet Protocol Journal*, Vol 3, Num. 4, 14-25, Dec. 2000.

[3] M. E. J. Newman "The structure of scientific collaboration networks", *Proc. Natl. Acad. Sci.*, 2001, pp. 409-415.
[4] A. L. Barabási, H. Jeong, Z. Nédá, E. Ravasz, A. Schubert, T. Vicsek, "Evolution of the social network of scientific collaborations". *arXiv:cond-mat/0104162*, 2001.
[5] A. Iamnitchi, M. Ripeanu, I. Foster "Locating Data in (Small-World?) Peer-to-Peer Scientific Collaborations". *1st International Workshop on Peer-to-Peer Systems*, Springer-Verlag, 2002.
[6] M. E. J. Newman and M. Girvan, "Mixing patterns and community structure in networks". *Proceedings of the XVIII Sitges Conference on Statistical Mechanics*, Springer-Verlag, Berlin, 2003.
[7] M. E. J. Newman, "Clustering and preferential attachment in growing networks". *cond-mat/0104209*.
[8] E. M. Jin, M. Girvan, and M. E. J. Newman, "The Structure of Growing Social Networks". *Santa Fe Institute working papers*, 2001.
[9] H. Zhou, "Distance, dissimilarity index, and network community structure". *arxiv.org/physics/0302032*.
[10] M. E. J. Newman, D. Watts, S. Strogatz, "Random graph models of social networks". *Proc. Natl. Acad. Sci.*, USA 99, 2566-2572 (2002).
[11] R. Albert, H. Jeong, and A.-L. Barabasi, "Error and attack tolerance of complex networks". *Nature* 406, 378-382 July 2000.
[12] R. A. Golding, "Weak-Consistency Group Communication and Membership". *PhD thesis*, University of California, Santa Cruz, Computer and Information Sciences Technical Report UCSC-CRL-92-52, December 1992.
[13] Q. Lv, P. Cao, E. Cohen, K. Li, and Scott Shenker. "Search and Replication in Unstructured Peer-to-Peer Networks". *Proc. ACMICS* 2002.
[14] J. Mitre, L. Navarro. "P2P scientific collaboration: structured topology and replication". *Technical Report No. UPC-DAC-2003-56*. Department of Computer Architecture, Polytechnic University of Catalonia, Spain, 2003.
[15] J. M. Marquès, L. Navarro. "WWG: a Distributed Infrastructure to support groups". *Proceedings of the ACM Conference: Group 2001 (Group'01)*.
[16] K. Ranganathan, A. Iamnitchi and I. Foster, "Improving data availability through dynamic model-driven replication in large peer-to-peer communities," in *GP2PC*, May 2002.
[17] A. Iamnitchi. "Resource Discovery in Large Resource-Sharing Environments". *PhD Thesis*. Department of Computer Science The University of Chicago Dec 2003.
[18] K. Sripanidkulchai, B. Maggs, and H. Zhang, "Enabling efficient content location and retrieval in peer-to-peer systems by exploiting locality in interests," *ACM Computer Communication Review*, vol. 32, Jan. 2002.
[19] E. Cohen, A. Fiat, and H. Kaplan, "A Case for Associative Peer to Peer Overlays". *Proc. of HotNets-I*. 2002.
[20] JXTA Homepage: www.jxta.org.
[21] Granovetter, M "The strength of weak ties: A network theory revisited". *American Journal of Sociology*, 78 (6): 1360-1380, 1973.
[22] J-sim homepage: <http://www.j-sim.org>
[23] K. Sripanidkulchai, B. Maggs, and H. Zhang "Efficient Content Location Using Interest-Based Locality in Peer-to-Peer Systems", *Infocom* 2003.