Semantic Gossiping

Karl Aberer, Manfred Hauswirth
Swiss Federal Institute of Technology (EPFL)
CH-1015 Lausanne, Switzerland
{Karl.Aberer, Manfred.Hauswirth}@epfl.ch

Abstract

We state that the problem of semantic interoperability in information search on the Internet is solved today mostly by means of centralization, both at a system and at a logical level. A new brand of system architectures, peer-to-peer systems, indicates that the principle of decentralization might lead to solutions to many problems that scale well to very large numbers of users. In this position statement we develop a scenario of how decentralization could be working for addressing the problem of information search at a global scale. In this scenario we extend the architecture of a well-known peer-to-peer file sharing system, namely Gnutella, to a peer-to-peer information sharing system enabling semantic interoperability that is driven in a bottom-up manner by the participating peers. Such a system could readily be used in order to study semantic interoperability as a global scale phenomenon taking place in a social network of information sharing peers.

1 Introduction

Today Internet-scale information search is implemented by systems that follow a centralized system architecture. We can distinguish two major categories of systems for Internet information search:

1. General-purpose full-text search engines
2. Domain-specific portals or data warehouses

Search engines such as Google, AltaVista, Excite and alike enable full-text search on unstructured data. Structural information, such as HTML markup or link information, is frequently used to improve the ranking schemes, thus it is used for an interpretation of data that is based on quantitative measures rather than qualitative statements. Because of a lack of application-specific schemas in the data sources (i.e. semantic markup in XML speak) more "intelligent" ways of searching the Web using search engines will remain a challenge for the future (and probably forever).

Domain-specific portals such as CiteSeer (www.researchindex.com, publication data), SRS (srs.ebi.ac.uk, biology) or streetprices.com (ecommerce) integrate data sources on the Internet and store them in a central data warehouse. The data is converted to a common data schema, which is usually of very simple to medium complexity. This approach adopts a simple form of wrapper-mediator architecture and requires typically substantial development effort for the automatic or semi-automatic generation of mappings from the source databases (or Web documents) into the global schema.

Both approaches solve the problem of semantic interoperability in widely distributed information systems by strong centralization, both at the logical level and at the implementation level. For other types of systems that require semantic interoperability, but are inherently not centralizable at the system level, such as electronic business systems, standardization is today the solution of choice to solve the interoperability problem. We can consider this as a form of logical centralization.

Recently a new class of systems has been appearing on the Internet, which we can consider as completely different approach to solving the problem of information search on the Internet: Peer-to-Peer systems [5]. They adhere to the prinicple of decentralization, whereas today the emphasis is on decentralization at the system level. To date peer-to-peer systems are mainly used for file sharing. Examples of popular (true) peer-to-peer systems are Gnutella [9] and FreeNet [6]. Their system implementation is completely decentralized, as every peer plays the same role in the system (symmetry) and there exist no global coordination or database. The peers have to adhere to system-specific communication protocols and implement internally system-specific request routing and data caching strategies. At the logical level peer-to-peer systems typically do not offer a lot. Content is represented by assigning a meaningful filename to shared files, for example, "King-Crimson-21st-Century-Schizoid-Man.mp3", and search is possible by using substring or equality predicates on this text string. Thus semantics is hidden in textual representations, comparable to the situation with Internet search engines. The specific approach to evaluate search requests by distributing them over the peer network is one of the major differences among the different systems.

- Gnutella uses a broadcasting strategy (gossiping). In that way many (or even most of the peers receive the search request, but a very large number of messages is generated
- FreeNet uses self-organizing routing tables at the peers to perform a more directed routing of the requests. The construction of the routing tables is driven by the requests passing by and tends to cluster related data items such that search speeds up as the network evolves and the number of messages tends to become lower.

Another type of peer-2-peer systems, of which Napster (www.napster.com) and FastTrack (www.fasttrack.nu) are the most popular examples, employs centralized directories in order to locate files. Napster relies (better: relied) on a single server site, whereas FastTrack dynamically assigns the server role to so-called super-peers. Together with this centralization at the system level we also see the use of a standardized schema to represent metadata on the shared files. We do not consider these systems as "true" peer-to-peer systems, in particular since they centralize especially that part of the system that is responsible for information search and semantic interoperability.

In parallel, Sun has developed its JXTA architecture [7] for peer-to-peer systems which aims at standardizing their architecture. Among others, they also propose standardized interfaces for forwarding search requests among peers [10], irrespective of what the structure and content of search messages is.

The emergence of this new type of systems raises the interesting question whether the paradigm of decentralization would not also be applicable and potentially helpful for tackling the problem of interoperability at a logical level, i.e. the problem of semantic interoperability. We will sketch in the following a general outline of how this approach could be taken up.

2 Peer-to-peer semantic gossiping)

In the following we intend to depart from the pure content-based textual search paradigm that underlies todays P2P system, and develop a scenario where peers use schemas in order to make information available in a semantically more meaningful manner. One immediate application that one can see even in the context of todays P2P systems, would be the provisioning of structured metadata that is used in order to annotate files (media files, documents etc.). For example, using XML as syntax and XML Schema or XML DTD as schema language, we could annotate music files as illustrated in the following example.

```
<Song>
  <Title>
    21st century schizoid man
  </Title>
  <Interpreter>
    King Crimson
  </Interpreter>
  <Size>
    23456789
  </Size>
</Song>
```

Using such metadata annotations would be a simple extension of the current practice of using textual strings for representing media content and could be adopted in existing implementations of P2P systems with minimal effort. Once such a representation is given, search requests can be not only posed as textual strings (which continues to be a possiblity for backward compatibility), but also by using a structured query language, like XPath. So a search requests could be

```
/Song[contains(Title, "schizoid") AND Size < 1000000]
```

The request can easily be processed against the local data collection of a peer. Of course, immediately we encounter the problem of semantic interoperability. The evaluation of the request will only be successful and/or meaningful when the requestor and the requestee share the same schema for structuring the annotations. Since we start from the assumption that there exists no central authority in a P2P system, which could standardize these schemas, we have to assume that different schemas exist.

Looking at the life cycle of a peer-to-peer system we can expect that the following basic situations occur:

1. Peers that own substantial amounts of data and have probably annotated it already will develop a proprietary metadata schema and publish it.

2. Peers that join the network without having put in place their own metadata schema will adopt some existing, published schema, when they encounter a peer already participating in the network.

3. Peers that are meeting and where each is using its specific metadata schema are willing to develop mapping functions between their schemas, both for data objects and queries.

As a result we arrive at a situation where there exist subsets (or clusters) $C_1, \ldots, C_n$ in the peer network. Peers within the same cluster share the same schema and the clusters are connected among each other wherever a "translation link" exists. Taking a quantitative perspective on the evolution of such a network (and drawing from experience with existing peer-to-peer networks) we can expect the following to happen:

- Peers that are sharing the same schema will be much stronger connected among each other than peers with different schemas, as they can interact more simply and less expensively. By "connected" we mean, that they are aware of each others existence and regularly forward to each other messages.
- The cluster sizes will follow a Power-law distribution [9], as there will exist on the one side of the spectrum a few schemas which become popular and on the other side of the spectrum many proprietary schemas which are supported only by few peers.
- The diameters of the network graphs will be short (i.e. logarithmic in the number of nodes). This will be true both for the whole network, for each cluster and for the inter-cluster network. Experience shows that these graphs typically exhibit small-world properties [8].

Now that we have a rough understanding of the possible network structure we want to develop a scenario on the behavior of peers in the network. The basic problem we address is the strategy to be used for answering search requests. We will base this discussion on the existing approach that is used in Gnutella, namely using a broadcast strategy. We are aware that this is extremely inefficient in terms of messages generated, but

- the strategy seems to a certain degree to be considered as acceptable since Gnutella is used in practice.
- broadcast is inefficient for search, but on the other hand very efficient in terms of updates as they can be performed independently at the peers.
- optimization of such an approach is an exciting research issue for the future (for which we already have first results [3]) and not an issue for this position paper.

A broadcast strategy for searching a data item in the network is based on three elements

1. a set of neighbors each peer knows and to which a peer forwards a search request. In Gnutella this set is for example typically of size 4.
2. a time-to-live (TTL) for messages that is the maximal number of hops a request message is forwarded in the network. For Gnutella this number is typically set to 7.
3. a message identifier that is used to detect loops, such that messages that come back to a peer that has already forwarded the message are not forwarded once more.

In principle exactly the same strategy could be applied also for our more general scenario of a network with semantic clusters and translation links among clusters. Peers would forward messages to their neighbors, and in case the neighbor is in another semantic cluster the necessary translations would be performed. However, we think that some care needs to be taken in that new situation, since forwarding a message within the same cluster and between clusters has a different quality. Multiple translations during inter-cluster forwarding might lead to a "semantic degradation" of the requests and schemas and alternative translation paths might lead to inconsistent results. "Back-translations" of requests, where a request is forwarded from a cluster C_i to another cluster C_j and at a later stage returns back to C_i should be avoided.

Therefore a better strategy is to use the intra-cluster network in order to achieve full coverage of the network, and use inter-cluster links as little as possible, just enough to "infect" new clusters with a search request. In order to achieve such a behavior we propose the following (practical) broadcast strategy.

- The number of neighbors a peer maintains within the same cluster is higher than the number of neighbors in other clusters, e.g. 4 neighbors within the cluster and 0 or 1 neighbor to another cluster. Since intra-links are more costly to establish this also corresponds to the natural evolution.
- each message is equipped with a local-time-to-live (LTTL) and a global-time-to live (GTTL). The LTTL is reset always to the maximal value when the message changes the cluster and reduced by one when the messages is forwarded within the same cluster, whereas the GTTL is reduced by one whenever the message changes the cluster. The GTTL can also be used for controlling the semantic distance that is considered as desirable or acceptable for answering the request.
- Messages keep a record of the schemas of the clusters they have visited and are not forwarded to clusters in which they have already occurred. For practical purposes we identify schemas by a hash value (which is useful also in other contexts, e.g. to locate schemas that have been published). Message identifiers would not be sufficient for that purpose since the messages can leave and enter the cluster through different peers.

By putting into place such an infrastructure we can hope to establish a laboratory for studying of how peers (which are of course instantiations of human users) interact, when they have the possibility to interact in a semantically more meaningful manner. It would be for example interesting to see whether specific schemas start to dominate the network, or multiple schemas connected by gateways could co-exist, or whether, for example, the network would partition into completely disconnected sub-networks. Essentially these processes will be driven by individual decisions of peers. They will be taking into account the basic trade-off of the cost of adapting their own schema to some other (and so adhering to some established schema), or producing the necessary translation to some other schema in order to stay connected to rest of the network. Some applications of such an approach could be easily put in place for domains like scientific data, media data or scientific publications.

Performance is a concern for the efficiency of the network. We have already mentioned the main disadvantage of using a broadcasting approach. For the simple cases of searching data items by using their data key or for simple textual search we have developed methods that exhibit good performance characteristics [1, 2] (i.e. logarithmic in search time, message number and storage cost at peers) while maintaining the functional characteristics of a Gnutella-like network and not limiting the peer's autonomy. The extensions of theses methods to the scenario described above will be of course an interesting research challenge.

3 Conclusion

If we would have to bring the message of this paper to the point we could summarize it as follows: it might be beneficial to move the focus from a mostly constructive approach, where the emphasis is put on developing ever more complex solutions for bridging semantical gaps through potentially complex and difficult to find mappings, to a rather more experimental approach by providing simple tools and infrastructures where users can start to work and interact, and let the complexity emerge as a result of the composite behavior of all participants. This is of interest from a research perspective as it opens some quite novel types questions, including social and economic dimensions [4], as well as from an application perspective as there surely exist application classes where such an approach could be successfully working.

References

[1] Karl Aberer. *P-Grid: A self-organizing access structure for P2P information systems.* Proceedings of the Sixth International Conference on Cooperative Information Systems (CoopIS 2001), Trento, Italy, 2001.

[2] Karl Aberer. *Scalable data access in P2P systems using unbalanced search trees.* Workshop on Distributed Access Structures (WDAS), Paris, 2002.

[3] Karl Aberer, Manfred Hauswirth, Magdalena Punceva, and Roman Schmidt. *Improving data access in P2P systems.* IEEE Internet Computing, 6(1), Jan./Feb. 2002.

[4] E. Adar and B. A. Huberman: *Free riding on Gnutella* Technical report, Xerox PARC, 10 Aug. 2000.

[5] D. Clark. *Face-to-Face with Peer-to-Peer Networking.* IEEE Computer, January 2001.

[6] Ian Clarke, Scott G. Miller, Theodore W. Hong, Oskar Sandberg, Brandon Wiley. *Protecting Free Expression Online with Freenet* IEEE Internet Computing, 6(1), Jan./Feb. 2002.

[7] Li Gong. *JXTA: A Network Programming Environment.* IEEE Internet Computing, 5(3):88-95, May/June 2001.

[8] Jon Kleinberg. *The Small-World Phenomenon: An Algorithmic Perspective.* Technical report 99-1776. Cornell Computer Science, October 1999. http://www.cs.cornell.edu/home/kleinber/swn.pdf

[9] Matei Ripeanu, Adriana Iamnitchi, Ian Foster *Mapping the Gnutella Network* IEEE Internet Computing, 6(1), Jan./Feb. 2002.

[10] Steve Waterhouse, David M. Doolin, Gene Kan, Yaroslav Faybishenko *Distributed Search in P2P Networks* IEEE Internet Computing, 6(1), Jan./Feb. 2002.