Power-Law Chord Architecture in P2P Overlays

Salma Ktari, Artur Hecker and Houda Labiod
Telecom ParisTech, ENST France
{ktari, hecker, labiod}@enst.fr

ABSTRACT
In this paper, we propose to extend the topology of Chord to a bidirectional graph to provide a super-peer based lookup algorithm and an efficient resource localization service at a very little additional cost.

Keywords
DHT, flooding, super peer, scale free network, DHT

1. INTRODUCTION
Peer-to-peer networking systems consist of a large number of nodes or computers that operate in a decentralized manner to provide reliable global services, such as query resolutions or distributed computing. Existing P2P algorithms assume that all peers are uniform in resources. Messages are routed on the overlay without considering the differences of capabilities among participating peers. However, the heterogeneity nature of participating nodes is quite important. Therefore, it is possible to improve the performance of these algorithms by taking into account the heterogeneous nature of P2P systems. Several approaches propose to organize the topology through a two-level hierarchy by introducing the concept of super nodes (SN) [1]. The idea is to assign a large number of low-capacity nodes to one or more stronger node, which then becomes SN. The super-peer solution helps low capacity users. However, the design introduces non-uniformity, as an explicit hierarchy is imposed on nodes. Furthermore, the SN selection problem is highly challenging because in the P2P environment, a number of super nodes must be selected from a potentially big distributed system, in which neither the node characteristics nor the network topology are known a priori.

In this paper we propose S-Chord, a new super-peer based lookup algorithm, without imposing a hierarchical architecture. The proposed algorithm uses a biased key-based routing and a power law like broadcast mechanism on top of a structured P2P network (Chord [3]). S-Chord search achieves a high query success rate and short routes, causing practically no additional cost.

CoNEXT 2008, December 9-12, 2008, Madrid, Spain.

2. SUPER PEER ROUTING IN S-CHORD
2.1 Motivation
A scale free network [2] is a network whose degree distribution follows a power law (PL). That is, the probability that a randomly chosen node has degree k is given by $P(k) = k^{-\alpha}$. In such networks, few nodes exhibit extremely high connectivity, while the majority is poorly connected. Scale-free networks exhibit many advantageous properties, such as small diameter (for $2 < \alpha < 3$, $d \sim \log \log N$), which allows fast searches, tolerance to random node deletions, and a natural hierarchy for an optimal usage of heterogeneous computing [4]. Such topologies can be exploited to provide scalable global services in highly dynamic and ad hoc environments.

To identify the topologically important nodes that deal with large numbers of connections, we consider two metrics: the outdegree and the indegree for each node. We observe the fact that, although in conventional DHTs nodes maintain a well known number of outdegree neighbors (typically, outgoing degree of each node equals $O(\log N)$); the indegree distribution is neither uniform nor normal. We measured the number of indegree links of each Chord node for a network size of 20000 and calculated the number of nodes corresponding to each indegree value. Figure 1 (a) shows the resulting indegree distribution (linear and log scale). From the chart, we can see that it is similar to a power law (PL) distribution especially for indegree values exceeding 30 (but a correct parameter estimation is work in progress). Indeed, few nodes exhibit extremely high connectivity, while the majority is poorly connected. As a result, Chord topology seems to be is a PL network, where most nodes have only log N neighbors, but some nodes have substantially more “contacts”. Nevertheless, Chord does not use its presumed PL knowledge for routing, but limits itself to $\log N$ outdegree neighbors (and ignores indegree neighbors). Therefore, it could be possible to design structured P2P networks that allegedly admit PL-like degree distributions from the conventional structured DHT. Our idea is to extend Chord to use its full neighborhood knowledge for routing and, if our assumption is correct, to establish a functioning PL network. Highly connected nodes may be considered as super nodes as they already have larger neighborhood knowledge.

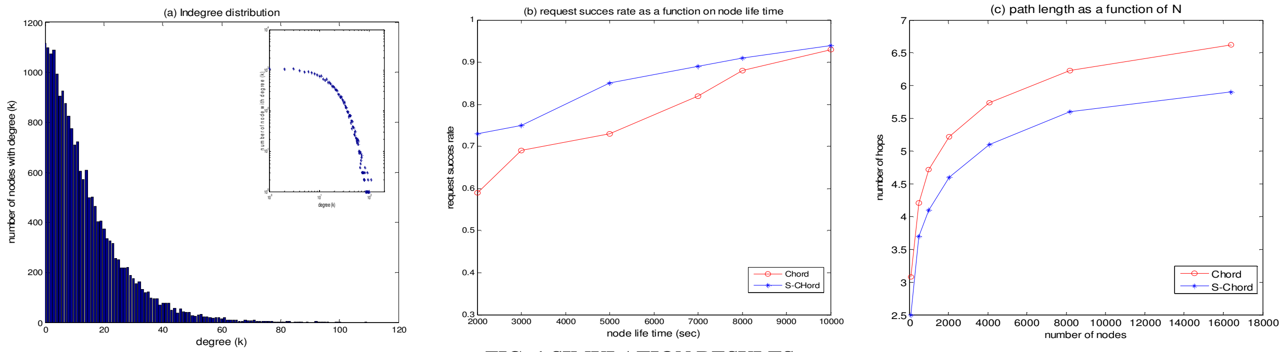

FIG. 1 SIMULATION RESULTS

2.2 Our proposal

Additional data structure: In S-Chord, a node maintains a new table, that we call a *finger selector table* (FST) in addition to the classic *finger table*. If there is a finger link from one node to its finger, we add a reverse link that links the same two nodes, but in the reverse direction. These reverse links connect the node to its *finger selectors* (FS), i.e. the nodes that selected it as a finger. In addition to the finger selector ID information, the FST stores the predecessor ID of the corresponding node.

Knowing the FS predecessor ID, every node may know the data interval associated to each FS entry, which corresponds to the index of the content of the corresponding node, and that without exchanging content information with neighbors, and consequently without generating extra overhead. The interesting property of this addition is that it comes at almost no additional cost.

Routing: Structured DHT networks perform key based routing. In fact, the query follows a calculated path according to the encountered routing tables and the searched key, regardless of the node's connectivity.

In order to shorten the query search process in Chord, we propose to use the finger selector table in addition to the finger table for the query routing. Therefore, when receiving a query, the intermediate node examines first its FST: the node compares the searched key to the IDs of each finger selector entry and the corresponding predecessor. For a given FST entry, if the searched key lies between these two IDs, the query is transmitted directly to this finger selector, as it is the owner of the searched key. Thus, the lookup search is resolved. If not, the node forwards the query to its closest preceding finger or finger selector. The search is performed in the entire neighbourhood.

S-Chord adds bidirectional links to enlarge the neighborhood knowledge especially for topologically important nodes. Knowing the predecessor ID information, the search is performed in the entire neighbourhood as a local broadcast, but that without flooding the neighbors or exchanging content information. Therefore, de facto highly connected Chord nodes (i.e. nodes with a high indegree) are automatically considered super nodes because they use a larger knowledge of the network. These nodes, when used for routing, would accelerate the query search process.

2.3 Simulation results

We compare the performance of our proposed search model S-Chord to the classic Chord DHT. We simulated a network with $N = 2^k$ nodes. We varied k from 9 to 14. Figure 4 (c) shows that S-Chord exhibits an improved behavior compared to standard Chord in term of average search path length. The facts that the query traverses highly connected nodes, that every node keeps a sort of index of its neighbors' shared content, and that routing is performed in both directions, accelerate considerably the query process. Figure 4 (b) plots the request success rate (RSR) as a function of node's mean life time. We note that S-Chord outperforms Chord and especially for shorter life time values (over 73% for all values). For such short lifetimes, in Chord, the success ratio decreases to less than 60%.

3. CONCLUSION

In this paper, we propose a new super-peer based lookup algorithm without imposing a hierarchical architecture. S-Chord employs KBR routing and provides a power law like broadcast mechanism, to perform query search. Primary simulation results show improvements over classic Chord. The proposed model might be considered as an alternative to artificially imposing a hierarchical structure. As a future work, it would be interesting to apply our model to different DHTs and evaluate the performance results. For this purpose, we have to adapt our scheme to flexible DHT geometries such as tree for Pastry or butterfly for Viceroy.

4. REFERENCES

[1] Yuhong Liu, Chris Gauthierdickey, "Scalable supernode selection in peer-to-peer overlay networks". *2nd International Workshop on Hot Topics in P2P Systems.*

[2] Diestel, Reinhard (2005), *Graph Theory* (3rd ed.), Berlin, New York: Springer-Verlag, ISBN 978-3-540-26183-4.

[3] I. Stoica, R. Morris, "Chord: a scalable peer-to-peer lookup service for internet applications", *Proc.ACM SIGCOMM 2001 Technical Conf., San Diego, USA, August 2001.*

[4] Barabási, Albert-László "Scale-Free Networks". *Scientific American*, 288:60-69, May 2003.