Volume 2, No. 5, Sept-Oct 2011
ISSN No. 0976-5697
International Journal of Advanced Research in Computer Science
REVIEW ARTICLE
Available Online at www.ijarcs.info

An account of Peer to Peer and Overlay Networks

Adarsh Nair
Computer Engineering Department,
University of Mumbai, India
adarsh.nair@ieee.org

Abstract: Computing needs and environments today are much more complicated than when compared to the primitive networking methodologies. Hence an increased interest in Peer to Peer and Overlay networks has come about due to their self governance, scalability, robustness and localized control. In this paper I present a review of both peer to peer and overlay networks with an example each and give a viewpoint as to why overlay networks are better than peer to peer networks when it comes to scalability.

Keywords: Peer to Peer Networks, Overlay Networks, Distributed Hash Table, Time to Live, Decentralization

I. INTRODUCTION

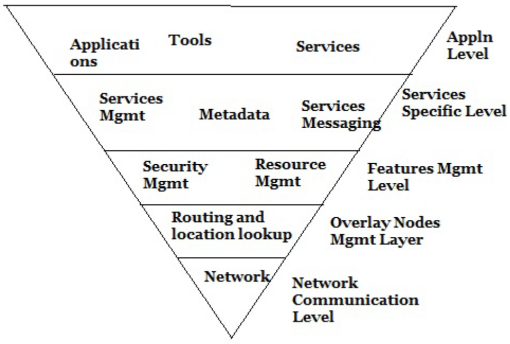

Figure1.P2P and Overlay Network Hierarchy Model

Peer-to-peer overlay systems go beyond services offered by client-server systems by allowing a client to also be a server. It allows access to its resources by other systems and supports resource-sharing, The dynamic nature of peers poses challenges in communication paradigm. The Network Communication Layer is at the bottom and controls the inter network communication. The Overlay Nodes Management layer covers the management of peers, which include discovery of peers and routing algorithms for optimization.

The Features Management layer deals with the security, reliability, fault resiliency and aggregated resource availability aspects of maintaining the robustness of P2P systems. The Services Specific layer supports the underlying P2P infrastructure and the application-specific components through scheduling of parallel and computation-intensive tasks, content and file management. Metadata describes the content stored across the P2P peers and the location information. The Application-level layer is concerned with tools, applications and services that are implemented with specific functionalities on top of the underlying P2P overlay infrastructure.

II. PEER TO PEER SYSTEMS

Peer-to-peer systems [1] are distributed systems that operate without centralized organization or control. Peers are equally privileged, equipotent participants in the application. Peers make a portion of their resources, such as processing power, disk storage or network bandwidth, directly available to other network participants, without the need for central coordination by servers or stable hosts. There are two classes of P2P networks: Structured and Unstructured.

The technical meaning of Structured is that the P2P network topology is tightly controlled and content are placed not at random peers but at specified locations that will make subsequent queries more efficient. Such Structured P2P systems use the Distributed Hash Table (DHT)[2] as a platform, in which data object location information is placed at the peers with identifiers corresponding to the data object's unique key. DHT-based systems have a property that consistently assigns uniform random IDs to the set of peers into a large space of identifiers.

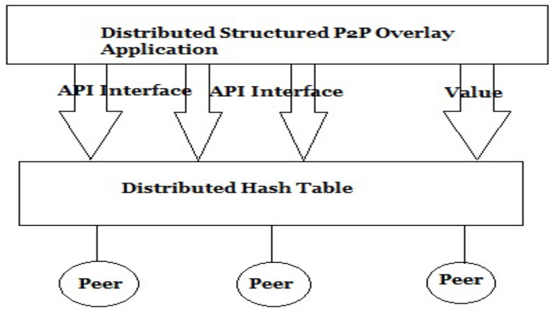

Figure 2. DHT Model

DHT research was originally motivated, in part, by peer-to-peer systems such as Freenet, gnutella, and Napster, which took advantage of resources distributed across the Internet to provide a single useful application. DHTs characteristically emphasize Decentralization[3] i.e.the nodes collectively form the system without any central coordination, fault tolerance i.e.the system should be reliable even with nodes continuously joining, leaving, and failing and scalability i.e.the system should function efficiently even with millions of nodes.A key technique used to achieve these goals is that any one node needs to coordinate with only a few other nodes in the system.

Unstructured peer-to-peer networks do not provide any algorithm for organization or optimization of network connections. An unstructured P2P network is formed when the overlay links are established arbitrarily. Such networks can be easily constructed as a new peer that wants to join the network can copy existing links of another node and then form its own links over time.

An Unstructured P2P system[4] is composed of peers joining the network with some loose rules, without any prior knowledge of the topology. The network uses flooding as the mechanism to send queries. While flooding-based techniques are effective for locating highly replicated items and are resilient to peers joining and leaving the system, they are poorly suited for locating rare items.

Current Internet P2P applications typically provide locator functions using time-to-live (TTL) controlled-flooding mechanisms. With this approach, the querying node wraps the query in a single message and sends it to all known neighbours. The neighbours then check to see whether they can reply to the query by matching it to keys in their internal database. If they find a match, they reply; otherwise, they forward the query to their own neighbours and increase the message's hop count. If the hop count passes the TTL limit,forwarding stops.

Unstructured P2P networks face one basic problem: peers readily become overloaded, therefore, the system does not scale when handling a high rate of aggregate queries and sudden increase in system size. Although Structured P2P networks can efficiently locate rare items since the key-based routing is scalable, they incur significantly higher overheads than Unstructured P2P networks .Hence over the Internet today, the decentralized Unstructured P2P overlay networks are more commonly used.

Overlay networks can address these issues. Overlay networks have a network semantics layer above the basic transport protocol level that organizes the network topology according to the nodes' content, implementing a distributed hash table abstraction that provides load balancing, query forwarding, and bounded lookup times.

III. WORKING OF P2P STYLE SEARCH

Let node Na is requesting the associated value of a key located in Nb . Nodes that can't answer the query forward it to their neighbours, eventually reaching Nb, which returns the result directly to the requesting node; the concentric circles indicate the number of message hops. Even though only Nb can answer the query, all the nodes within TTL-range must process it. Also, if the value had been stored in node Nx, the query result would not be found unless the message's TTL was set to a higher value.

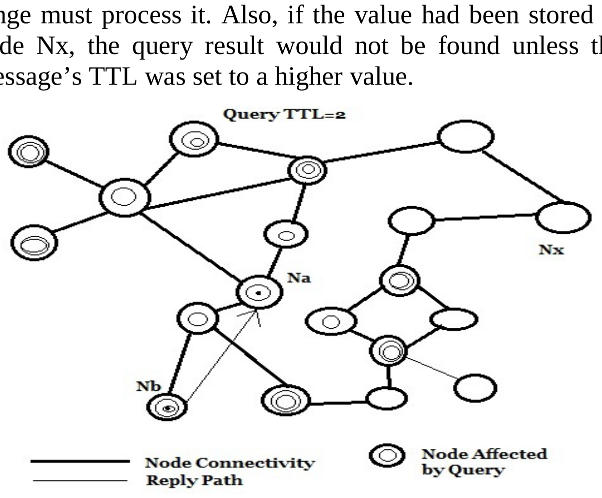

Figure 3. Sample TTL-based P2P network and query. The Na node transmits a query requesting the value of a key located in Nb.Concentric circles indicate the number of message hops.

Although some networks adapt to the underlying physical topology, such optimization is not required for the algorithm to operate properly. These networks are unstructured: nodes attach to the network according to measures unrelated to content, such as join-order, connection speed,and even physical proximity, creating a random connection topology. Although this approach makes maintaining connections simpler, it has two problems:

a. Content location and network topology are uncorrelated. Network searches are essentially open ended, forcing protocols to use TTL measures to control message propagation and avoid flooding the whole network. Thus, available content might not be accessible to all network nodes, and a query hit cannot be guaranteed even if the target node is connected to the network.
b. The network is random. As a result, searching for a particular element within the horizon has a theoretical limit of N hops, where N is the number of nodes within the query's reach. In practice, however, the networks typically traverse different sections of the graph in parallel, reducing lookup times. Still, strictly speaking, queries on an unstructured P2P network tend to have lookup complexity of the order of N, or O(N), hops.

IV. OVERLAY NETWORKS

Overlay networks create a virtual topology over a physical topology. It is a computer network which is built on top of another network. . Nodes in the overlay can be thought of as being connected by virtual or logical links, each of which corresponds to a path, perhaps through many physical links, in the underlying network. For example, distributed systems such as cloud computing, peer-to-peer networks, and client-server applications are overlay networks because their nodes run on top of the Internet. The Internet was built as an overlay upon the telephone network. Overlay networks share four main qualities:

a. Guaranteed data retrieval
b. Provable lookup-time horizons (typically O(log N) with N being the number of network nodes)
c. Automatic load balancing
d. Self-organization

Overlay networks define neighbour nodes by content stored, and hence they can change search from a standard graph-traversal problem into a localized iterative process. In this process, each hop brings the query closer to its target set of hops, which can be calculated according to a mathematical function. an overlay network operates like a distributed hash table by allowing key insertion, querying, and removal. This reduces the overall network load and makes the query process deterministic.

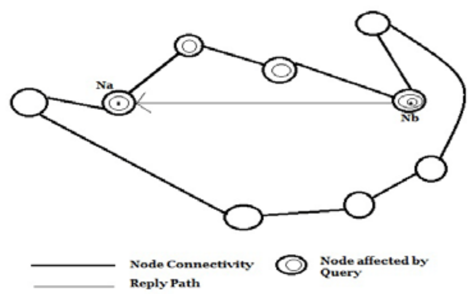

Figure 4. The Na node transmits a query requesting the value of a key

located in Nb.Concentric circles indicate the number of message hops.

An overlay network's connectivity pattern is different from that obtained using a TTL-based algorithm in that it is structured and typically symmetrical. The structure is based on one or more mathematical functions that determine how the nodes are connected. The network's structure contributes to the overlays' bound lookup times. When nodes fail, overlay network algorithms provide mechanisms that let the network recover and recreate or maintain an appropriate network structure. An important difference between overlay networks and unstructured P2P networks is that overlays lookup data on the basis of identifiers derived from the content, and thus don't directly support keyword-based searching. 3 simple rules to create an overlay topology are :
Each overlay node has two neighbours: the node whose value is the next available (higher) integer, and the node whose value is the previous available (lower) integer.

If the current node is the network's lowest or highest identifier, one of the neighbours will be the opposite value in the available node range(that is, the highest or the lowest, respectively).

To join the network, a node must perform a broadcast to find another network node. The incoming node can then use the search function to find the network "slot" where it should insert itself.

The node that initiates the query determines the relation between its own value and the target value. If the target value is higher than the node's value, the node passes the request to its higher-value neighbor, if it is lower it passes it to its lower value network. This local decision process continues until the request reaches the destination node, which replies directly to the requester, sending its physical network address for additional operations.

Figure 4 shows a hypothetical overlay network built using our simple example algorithm propagating a query. Because overlay nodes are connected according to the content stored in them,queries can be routed efficiently to the target.This example is unrealistic because the search time is bound but linear (the maximum number of hops is N/2), which creates unacceptable lookup times. It also fails to deal with recovery and possible loops created by missing nodes in the topology. However, the example does show how a set of simple rules lets nodes use their content to self-organize and provide bound lookup times. Current overlay networks are useful for applications that require reliable, highly scalable, and selforganizing storage and lookup for unique key–value pairs. This includes distributed databases, processing clusters, and deterministic search applications.

V. CONCLUSION

Researchers are using overlay networks in diverse applications, ranging from Internet routing to distributed network storage. The overlay-based Internet Indirection Infrastructure (i3) routing system, for example, aims to simplify network services' deployment and management by decoupling the acts of sending and receiving.6 This additional level of indirection allows for more flexibility in node mobility, and in service location and deployment. Researchers are also successfully deploying overlay networks as part of distributed storage systems. Overlay network algorithms are the subject of ongoing research and development. In particular, researchers are working to reduce network operation costs, such as multiple concurrent node join and leave, fault tolerance, security, and physical proximity (by modifying the overlay to adapt better to the underlying physical topology).

VI. REFERENCES

[1] Rüdiger Schollmeier, A Definition of Peer-to- Peer Networking for the Classification of Peer-to-Peer Architectures and Applications, Proceedings of the First International Conference on Peer-to-Peer Computing, IEEE (2002).

[2] Kelaskar, M.; Matossian, V.; Mehra, P.; Paul, D.; Parashar, M. (2002), A Study of Discovery Mechanisms for Peer-to-Peer Application

[3] Hari Balakrishnan, M. Frans Kaashoek, David Karger, Robert Morris, and Ion Stoica. Looking up data in P2P systems. In Communications of the ACM, February 2003.

[4] Lua, Eng Keong; Crowcroft, Jon; Pias, Marcelo; Sharma, Ravi; Lim, Steven (2005). "A survey and comparison of peer-to-peer overlay network schemes.

[5] Gareth Tyson, Andreas Mauthe, Sebastian Kaune, Mu Mu and Thomas Plagemann. Corelli: A Dynamic Replication Service for Supporting Latency-Dependent Content in Community Networks. In Proc. 16th ACM/SPIE Multimedia Computing and Networking Conference (MMCN), San Jose, CA (2009)

[6] Foundation of Peer-to-Peer Computing, Special Issue, Elsevier Journal of Computer Communication, (Ed) Javed I. Khan and Adam Wierzbicki, Volume 31, Issue 2, February 2008

[7] Kelaskar, M.; Matossian, V.; Mehra, P.; Paul, D.; Parashar, M. (2002), A Study of Discovery Mechanisms for Peer-to-Peer Application

[8] A. Barabasi and R. Albert, "Emergence of scaling in random networks,"Science, vol. 286, no. 509, 1999.

[9] L. Adamic, R. Lukose, A. Puniyani, and B. Huberman, "Search inpower-law networks," Physical Review E, vol. 64, 2001.

[10] B. Yang and H. Garcia-Molina, "Efficient search in peer-to-peer networks," in Proceedings of the 22nd IEEE International Conferenceon Distributed Computing Systems (ICDCS), July 2002.

[11] H. J. R. Albert and A. Barabasi, "Attack and tolerance in complex ´networks," Nature, vol. 406, no. 378, 2000. N. J. A. Harvey, M. B. Jones, S. S. amd M. Theimer, and A. Wolman,

[12] "Skipnet: A scalable overlay network with practical locality properties,"in Proceedings of the 4th USENIX Symposium on Internet Technologiesand Systems (USITS), Seattle, WA, USA, March 2003.

[13] B. T. Loo, R. Huebsch, I. Stoica, and J. M. Hellerstein, "The case for a hybrid p2p search infrastructure," in Proceedings of the 3rd International Workshop on Peer-to-Peer Systems (IPTPS), San Diego, California, USA, February 26-27 2004.

[14] S. Bellovin, "Security aspects of napster and gnutella," in Proceedings of the 2001 Usenix Annual Technical Conference, Boston, Massachusetts, USA, June 2001.

[15] G. Hardin, "The tragedy of the commons," Science, vol. 162, pp. 1243–1248, 1968.

[16] E. K. Lua, J. Crowcroft, and M. Pias, "Highways: Proximity clustering for scalable peer-to-peer network," in Proceedings

of the IEEE Fourth International Conference on Peer-to-Peer Computing (P2P'04), August 25-28 2004, pp. 266–267.

[17] M. Costa, M. Castro, A. Rowstron, and P. Key, "PIC: Practical Internet Coordinates for Distance Estimation," in 24th IEEE International Conference on Distributed Computing Systems (ICDCS' 04), Tokyo,Japan, March 2004.

[18] F. Dabek, R. Cox, F. Kaashoek, and R. Morris, "Vivaldi: A decentralized network coordinate system," in Proceedings of the ACM SIGCOMM '04 Conference, Portland, Oregon, August 2004.

[19] Y. Shavitt and T. Tankel, "Big-bang simulation for embedding network distances in euclidean space," in Proceedings of the IEEE INFOCOM '03 Conference, San Francisco, California, USA, March 30 - April 3 2003.

[20] "On the curvature of the internet and its usage for overlay construction and distance estimation," in Proceedings of the IEEE INFOCOM '04 Conference, Hong Kong, March 7-11 2004.

[21] M. Castro, M. Costa, and A. Rowstron, "Performance and dependability of structured peer-to-peer overlays," in Proceedings of the 2004 International Conference on Dependable Systems and Networks, Palazzo dei Congressi, Florence, Italy, June 28 - July 1 2004.

[22] D. Liben-Nowell, H. Balakrishnan, and D. Karger, "Analysis of the evolution of peer-to-peer systems," in Proceedings of the Annual ACM Symposium on Principles of Distributed Computing, Monterey, California, USA, 2002.

[23] L. Alima, S. El-Ansary, P. Brand, and S. Haridi, "Dks(n,k,f): a family of low communication, scalable and fault-tolerant infrastructures for p2p applications," in Proceedings of the 3rd IEEE/ACM International Symposium on Cluster Computing and the Grid, Monterey, California, USA, 2003, pp. 344–350.

[24] X. Li and C. Plaxton, "On name resolution in peer to peer networks," in Proceedings of the 2nd ACM International workshop on principles of mobile computing, Monterey, California, USA, 2002, pp. 82–89.

[25] F. Kaashoek and D. Karger, "Koorde: A simple degree-optimal hash table," in Proceedings of the 2nd International Workshop on Peer-toPeer Systems (IPTPS 03), Berkeley, CA, USA, February 20-21 2003.

[26] I. Abraham, D. Malkhi, and O. Dubzinski, "Land: Stretch (1+epsilon) locality aware networks for dhts," in Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA'04), New Orleans, LA., USA, 2004.

[27] N. D. de Bruijn, "A combinatorial problem," Koninklijke Netherlands:

[28] Academe Van Wetenschappen, vol. 49, pp. 758–764, 1946.

[29] M. Naor and U. Wieder, "Novel architectures for p2p applications: the continuous-discrete approach," in Proceedings of the 15th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA 2003), San Diego, California, USA, June 7-9 2003, pp. 50–59.

[30] D. Loguinov, A. Kumar, and S. Ganesh, "Graph-theoretic analysis of structured peer-to-peer systems: routing distances and fault resilience," in Proceedings of the ACM SIGCOMM, Karlsruhe, Germany, August 25-29 2003, pp. 395–406.

