

Trivial Support For Continous Queries in Formless Peer-To-Peer Networks

¹M. Chandrasekhar Varma, ²M. M. Balakrishna

^{1,2}Dept. of CSE, Akula Gopayya College of Engg. and Tech., Prathipadu, Tadepalligudem, AP, India

Abstract

This paper presents CoQUOS a scalable and lightweight middleware to support continuous queries in unstructured P2P networks. A key strength of the CoQUOS system is that it can be implemented on any unstructured overlay network. Moreover, CoQUOS preserves the simplicity and flexibility of the overlay network. The CoQUOS middleware is a completely decentralized scheme to register a query at different regions of the P2P network. This mechanism includes two novel components, namely cluster resilient random walk and dynamic probability-based query registration technique. The loosely-coupled and highly dynamic nature of underlying P2P network poses several additional challenges. This paper focuses on the issues that are of particular importance to the performance of the CoQUOS system, namely Churn of the P2P overlay and Load distribution among peers.

Keywords

Continuous queries, Publish-Subscribe Systems, Cluster-Resilient Random Walk, Loosely-Coupled Networks, Peer-To-Peer Networks


I. Introduction

Peer-to-peer (P2P) networking is one of the applications that are based on overlays. Unstructured architectures are attractive because of their simplicity and their high robustness. In unstructured systems, there is neither a centralized directory nor any control over the network topology or resource placement. When a new peer joins the P2P network, it forms connections with other peers freely, e.g., it selects arbitrary peers as neighbours. In order to publish its resources, a peer usually just stores them locally or places them on randomly chosen peers. Generally, unstructured overlays have loose guarantees for resource discovery, and it is possible that a file is not found although it exists in the network. It is often believed that due to the absence of topological constraints such unstructured systems have a better performance and require less maintenance overhead in highly dynamic environments where peers join and leave frequently and concurrently. Basically, there are two fundamental routing operations: flooding and random walks. In flooding, a search packet with a limited time-to-live (TTL) maximally 10 hops in Gnutella for example, is repeatedly forwarded to all neighbors, while in random walks, a packet is only forwarded to one randomly chosen neighboring peer.

Most unstructured P2P content distribution systems only support a very simple model for data sharing and discovery called the ad hoc query model. A peer that is interested in discovering data items initiates a query with a set of search parameters, which is then circulated among the peers according to the specific query forwarding mechanism employed by the network. A peer receiving a query responds to the query initiator, if it has any content satisfying the search criterion. Once a query has been processed at a node, it is removed from the local buffers. Therefore, a query exists within the P2P network only until it is propagated to various nodes and processed by them. Once a query completes its circulation, the system essentially forgets it. The ad hoc query model provides no support for peers to advertise or announce the

data-items they own to other interested peers.

Pub-sub system enables its users to register subscriptions expressing their interests and to announce the occurrence of certain events by publishing them. The pub-sub system matches incoming announcements to the existing subscriptions and notifies the users that have registered the matching subscriptions. But it requires specialized overlay topologies, intricate indexing and routing mechanisms and also complex algorithms.



II. Related Works

The work presented in this paper is primarily related to two fields, namely P2P networks and publish subscribe systems (event-delivery systems), both of which have been very active areas of research in the past few years. Pub-sub systems can be classified into two broad categories:

1. topic-based and
2. content-based. With the aim of enhancing scalability, efficiency and scalability several distributed pub-sub systems have been proposed. Recently P2P computing models have been utilized for this purpose. Researchers have explored two strategies for constructing P2P-based pub-sub systems, namely
 - adopting a structured P2P network as the underlying substrate, and utilizing its indexing schemes for mapping subscriptions and events to nodes of the P2P systems,
 - organizing the nodes of the P2P system into specialized topologies and/or embedding application specific distributed index structures within nodes of the P2P network.




Fig. 2:

The CoQUOS system differs from the above works in terms of motivation, goals and approach. The goal of the above systems is to improve the various performance parameters of pub-sub systems and they use P2P-based techniques as a means towards this end. In contrast, goal is to enhance the P2P data sharing systems, and continuous queries is a means towards that end. Second, the above pub-sub systems cannot be implemented on top of generic P2P networks, they need specialized overlays. CoQUOS middleware does not need any distributed indexing structures, nor does it impose any topological constraints on the overlay network. Recently, researchers have studied random walk as a scalable and efficient alternative to flooding-based searching in unstructured P2P networks. Several variants of random walk have also been designed for this purpose. The cluster resilient random walk scheme that we have proposed for continuous query propagation attempts to alleviate this problem by favoring out-of-cluster nodes.

III. Proposed Methodology

Considering the complexities involved in developing full fledged pub-sub systems on top of generic unstructured overlays, this paper presents a scalable and lightweight middleware that supports continuous queries and advertisements in unstructured P2P networks. Two important features distinguish the CoQUOS system. First, it does not impose any topological constraints on the underlying P2P network and can be implemented as an independent module in any unstructured overlay network. Second, the CoQUOS system is very lightweight as it does not require complex index structures or routing techniques. Thus, it preserves the design simplicity of the unstructured overlay networks as well as their flexibility towards transient node populations. This approach is a completely decentralized technique for registering and storing copies of a continuous query at various regions of the P2P network. The CoQUOS system utilizes the stored query replicas for notifying the source peers of matching advertisements that were issued by other nodes in those regions.

A. This Paper Makes the Following Technical Contributions

1. First, I present the architectural design of the CoQUOS system for supporting continuous queries in unstructured P2P networks. In my architecture each peer maintains a set of continuous queries and notifies the respective query issuers of any matching data items that it discovers. This architecture includes a completely decentralized mechanism for registering a query at various regions of the P2P network.
2. Second, I propose a novel Cluster Resilient Random Walk (CRW) technique for propagating a query to various regions of the network. While preserving the overall framework of random walks, the CRW scheme favors neighbours that are more likely to send the message deeper into the network.
3. Third, I design a dynamic probability scheme for ensuring that query registrations are well distributed along the path of the query. In this scheme, a query that has not been registered in the past several hops has a higher chance of getting registered in its next hop.
4. The main focus in this paper is on the design of decentralized strategy for selecting good beacon node sets for the continuous query. The characteristic of beacon node is listed below.

B. Characteristics of a Good Beacon Set

The beacon nodes of a query should be distributed in all major regions of the overlay network. This property is essential for

achieving high notification success rates as the advertisements are circulated through very limited broadcast around the advertising peer.

To prevent duplicate notifications, the beacon nodes of a query should not be located too close to one another.

Factors such as churn in the overlay network and the load imbalance among the network nodes can have considerable impact on the performance of the CoQUOS system. The loosely-coupled poses several additional challenges. This paper also focuses on Churn of the overlay and Load distribution among peers. These two issues are very important to the performance of the CoQUOS systems

IV. Load Distribution

Achieving good load distribution among peers is important requirement for the performance of the CoQUOS system. The number of queries and the numbers of notifications sent out per unit time by various nodes represent two key load metrics for CoQUOS system. These load parameters can vary among the nodes of the CoQUOS system. Ensuring good load balancing in decentralized, loosely coupled systems such as unstructured P2P overlays is challenging. The goal for the CoQUOS system is to ensure good load balancing in an efficient and scalable manner. The CoQUOS system includes two schemes for balancing notification loads among its peers.

- Active load balancing: In this scheme, an over loaded peer ask one or more of its less-loaded neighbours to take over a few of the queries that are current registered at the peer.
- Passive load balancing In this scheme, an over-loaded peer avoids registering any new queries until the load becomes more balanced.




Fig. 3:

message rate is the predominant component of the cumulative message rate. The messages due to system maintenance operations form a very small fraction of the cumulative message load demonstrating the lightweight nature of the CoQUOS

V. Existing System

Despite their popularity, most of the current unstructured P2P content distribution systems suffer from certain serious limitations. One such limitation is their simple, on demand mechanism for content discovery. Peers in these systems discover data items by circulating queries within the overlay network. A peer receiving a query responds back to the initiating node if it has any matching content. Upon processing a query, the recipient node removes it from its local buffers¹. Thus, a query expires after it completes

its circulation within the network. In other words, the network forgets the queries once they have completed their circulation. For clarity purposes, we call this the ad hoc query model, and we refer to the queries as ad hoc queries.

A. Disadvantages

1. One such limitation is their simple, on demand mechanism for content discovery.
2. However, this approach is unviable. Besides heavy messaging overheads, this scheme could overwhelm the peers with unwanted advertisements.
3. The ad hoc query model suffers from two main shortcomings. First, an ad hoc query is only capable of searching and retrieving content that exists in the P2P network at the time the query was issued.

VI. Proposed System

We Propose, focus on an alternate notification paradigm called the continuous query model. Similar to content-based pub-sub systems this model provides a mechanism through which peers can register their queries, which are maintained in the network for extended durations of time. However, in contrast to traditional pub-sub model, a system implementing the continuous query model provides a best-effort notification service for the registered queries informing their initiating nodes of new content that may have been added in the recent past.

Peers in these systems discover data items by circulating queries within the overlay network. A peer receiving a query responds back to the initiating node if it has any matching content. Upon processing a query, the recipient node removes it from its local buffers¹. Thus, a query expires after it completes its circulation within the network. In other words, the network forgets the queries once they have completed their circulation. For clarity purposes, we call this the ad hoc query model, and we refer to the queries as ad hoc queries.

A. Advantages

1. First, we present a novel query propagation technique called Cluster Resilient Random Walk (CRW). This technique retains the overall framework of the random walk paradigm. However, at each step of propagation, CRW favors neighbors that are more likely to send messages deeper into the network thereby enabling the continuous queries to reach different topological regions of the overlay network.
2. Second, a dynamic probability scheme is proposed for enabling the recipients of a continuous query to make independent decisions on whether to register the query. In this scheme, a query that has not been registered in the past several hops has a higher chance of getting registered in its next hop, which ensures that registrations are well distributed along the path of a query message.
3. Third, we discuss a passive replication-based scheme for preserving high notification effectiveness of the system even when the underlying P2P network experiences significant churn.

B. Module Description

1. Cluster Resilient Random Walk
2. Dynamic Probability Scheme
3. Passive Replication
4. Overlay Churn

1. Cluster Resilient Random Walk

Random walk corresponds to a depth first traversal of the network, and a message propagated through random walks has a higher probability of reaching remote regions of the network than its flooding-based counterpart. In this paper we use the terms random walk and pure random walk (PRW) interchangeably.

The above property of the random walk makes it an attractive paradigm for propagating continuous queries. Unfortunately, the random walk protocol suffers from one significant drawback that undermines its utility for propagating queries in the CoQUOS system.

we have designed a novel query dissemination scheme called Cluster Resilient Random Walk (CRW). This scheme is motivated by a crucial observation: Two peers belonging to the same cluster generally have large numbers of common neighbors.

The discussion in the previous section highlights the crucial role played by the beacon nodes in notifying the source peer of matching data items. Hence, the choice of beacon nodes would have a significant impact on the notification success rates of a continuous query. So, an important research question is: how do we select the set of peers that will serve as the beacon nodes of a query?. In other words, which set of peers

Let $NbrList(P_j)$ denote the list of neighbors of P_j . Let the peers $P_k, P_{k+1}, P_{k+2}, \dots, P_{k+l}$ denote the neighbors of the node P_j . Let $UniqueNbrs(P_{k+1}, P_j)$ denote the set of neighboring peers of P_{k+1} that are not the neighbors of P_j (i.e., $UniqueNbrs(P_{k+1}, P_j) = NbrList(P_{k+1}) - (NbrList(P_j) \cap NbrList(P_{k+1}))$). Suppose the node P_j receives a query message from a neighboring peer P_k . In the CRW algorithm, the probability of a neighbor P_{k+1} receiving the message in the next hop (represented as $FwdProbability(P_{K+1})$) is proportional to

$$\left(\frac{UniqueNbrs(P_{k+1}, P_j)}{NbrList(P_{k+1})} \right)^\lambda \cdot \left| \begin{array}{l} \text{Normalizing over all the neighbouring nodes of } p_j \text{ except } p_k \\ FwdProbability(P_{K+1}) \end{array} \right|$$

$$= \frac{\left(\frac{UniqueNbrs(P_{k+1}, P_j)}{NbrList(P_{k+1})} \right)^\lambda}{\sum_{P_i \in \{NbrList(P_j) - P_k\}} \left(\frac{UniqueNbrs(P_i, P_j)}{NbrList(P_i)} \right)^\lambda} \quad (1)$$




Fig. 4:

Flooding-based broadcast is an option for circulating continuous queries. However, this would be analogous to a breadth first traversal of the network. As previous studies have reported, in this

scheme, messages remain in close vicinity of the source node and do not go deep into the network [18-19]. Random walk is another message propagation paradigm that has received considerable attention from the P2P research community [18, 23]. In the context of P2P networks, random walk works as follows: When a peer node P_i receives a message whose TTL has not expired, it selects one of its neighbors completely at random and forwards the message to that peer. Since, at each step the message is forwarded to only one neighbor, the message load imposed by random walk is very low. Random walk corresponds to a depth-first traversal of the network, and a message propagated through random walks has a higher probability of reaching remote regions of the network than its flooding-based counterpart. In this paper, we use the terms random walk and Pure Random Walk (PRW) interchangeably. The above property of the random walk makes it an attractive paradigm for propagating continuous queries. Unfortunately, the random walk protocol suffers from one significant drawback that undermines its utility for propagating queries in the CoQUOS system. Prior studies have shown that the ability of random walk in propagating messages to remote topological regions diminishes significantly on overlay networks that exhibit significant degrees of node clustering [18]. In these networks there are distinct clusters of nodes with large numbers of connections among




Fig. 5:

them, whereas the connections flowing across clusters are comparatively small in number. For these networks, the random walk protocol suffers from the following drawback: When a message enters a cluster, it is likely to keep circulating within the cluster for large number of hops before exiting the cluster. Thus, the message spends a significant fraction of its TTL before reaching a different topological region of the overlay. In Fig. 2, a random walk message that reaches peer P_4 has a high probability of visiting majority of the other peers in P_4 's cluster (i.e., peers P_0, P_1, P_2, P_3), possibly multiple times, before going to other regions of the network. In other words, the message gets trapped in the cluster for considerable number of hops thereby adversely affecting its ability to reach remote regions of the overlay. Toward overcoming the above drawback, we have designed a novel query dissemination scheme called CRW. This scheme is motivated by a crucial observation: Two peers belonging to the same cluster generally have large numbers of common neighbors. Thus, a peer P_j has a lesser likelihood of being in the cluster of another peer P_i , if a

A. Dynamic Probability Scheme

The CRW scheme provides a mechanism for propagating a continuous query. But, how does a node receiving this message decide whether to register the query? A straightforward solution would be to register a query at every node it visits. However, this would result in large numbers of unnecessary subscriptions, which affects the efficiency of the network.

The reason is that for some continuous queries a long series of peers in the path of the query message may all decide not to register the query, whereas another sequence of consecutive nodes may all decide to host the query. The announcements originated near the dry patches of a query's path might fail to reach any of its beacon nodes, thus leading to low success rates. Considering these requirements, we have designed a novel dynamic probability-based technique (DP scheme, for short) for peers to decide whether to register a continuous query. However, the registration probability of a query varies as the query traverses along its route. The central idea of the dynamic probability scheme can be summarized as follows:

The probability of registering a query at a peer node would be high if the query has not been registered at the nodes it visited in the recent past. In contrast, if the query has been registered at a node that visited in the past few hops, the probability of it getting registered at the Current peer would be low.

The CRW scheme provides a mechanism for propagating a continuous query. But, how does a node receiving this message decide whether to register the query? A straightforward solution would be to register a query at every node it visits. However, this would result in large numbers of unnecessary subscriptions, which affects the efficiency of the network. Alternatively, each peer receiving a query message can decide register it with a certain fixed probability, say R_p . We call this scheme the fixed probability-based query registration scheme (FP scheme). Although this strategy seems intuitive, it cannot guarantee high notification success rates for every query. The experiments in Section 6 confirms our contention in this regard. The reason is that for some continuous queries a long series of peers in the path of the query message may all decide not to register the query, whereas another sequence of consecutive nodes may all decide to host the query. The announcements originated near the dry patches of a query's path might fail to reach any of its beacon nodes, thus leading to low success rates.




Fig. 6:

Considering these requirements, we have designed a novel DP-based technique for peers to decide whether to register a continuous query. As in the fixed probability scheme, a peer receiving a query-message registers it with certain probability. However, the

registration probability of a query varies as the query traverses along its route. When a peer issues a query, the registration probability is set to an initial value (called initial probability). When a peer P_i receives a query message Q_m , it registers the query with probability $R_{p\delta QmP}$. If P_i registers the query, it also resets the value of $R_{p\delta QmP}$ to the default initial value, before forwarding the message to one of its neighbors. On the other hand, if P_i decides not to register the query, it increments $R_{p\delta QmP}$ by a predetermined amount (called probability increment) before forwarding the query to one of its neighbors. Thus, the registration probability value associated with a query message keeps increasing until it gets registered at a peer, at which point it falls suddenly to the default initial value. The number of beacon nodes of a query can be controlled through the initial probability and probability increment parameters. Higher values of these parameters result in larger number of subscriptions and vice-versa. Experiments show that our decentralized beacon node selection scheme comprising of CRW and the DP query registration scheme, not only yields significant improvement in the overall success rates, but also ensures reasonably high individual success rates for all queries.

B. Passive Replication

We discuss a passive replication-based scheme for preserving high notification effectiveness of the system even when the underlying P2P network experiences significant churn. This churn of the overlay network can adversely impact the success of continuous queries and announcements. When a node P_i gracefully leaves the system, it asks one of its neighbors to handle all registered queries at P_i and also notifies all the beacon nodes with queries issued by P_i to remove the queries. However, when P_i exits the system unexpectedly, all the registrations are lost and the notification success rates of the respective queries and the matching announcements drop. Thus, effective mechanisms are needed to alleviate the negative effects of churn in the overlay network.




Fig. 7:

C. Overlay Churn

In order to counter the adverse effects of network churn, we have designed a low-cost technique wherein the query registrations present on a peer are replicated on one or more of its neighbors. Failures are detected through periodic exchange of heartbeat messages between the beacon node and the peers maintaining its replicas. The beacon node that does not respond to two consecutive messages is assumed to have failed. In the interest of better load

distribution, two or more neighbors may takeover subsets of the queries registered at the failed node. The communication costs of maintaining query replicas are optimized through lazy replication and piggybacking the information they utilize. They only require interactions between neighboring peers, thus making them suitable for generic unstructured P2P networks. In both schemes, neighboring peers periodically (at the end of pre-specified cycles) exchange information about their loads. Based on the load information obtained from its neighbors, a peer decides whether it is overloaded.

The objective of this paper is to introduce a model to guide the analysis of the impact of churn in P2P networks. Using this model, a variety of node membership scenarios is created. These scenarios are used to capture and analyze the performance trends of chord, a Distributed Hash Table (DHT) based resource lookup protocol for peer-to-peer overlay networks. The performance study focuses both on the performance of routing and content retrieval. This study also identifies the limitations of various churn-alleviating mechanisms, frequently proposed in the literature. The study highlights the importance of the content nature and access pattern on the performance of P2P, DHT-based overlay networks. The results show that the type of content being accessed and the way the content is accessed has a significant impact on the performance of P2P networks on a peer P_i are replicated at r_f (replication factor) of its neighbors. If P_i fails, the failure will be noticed by a neighbor, say P_k , that maintains a replica of subscriptions registered at P_i . P_k claims ownership of the queries registered at the failed node by sending messages to other neighbors of P_i , and receiving consent from them. Simultaneous ownership claims by multiple neighbors will be resolved in favor of the peer with smallest ID. Once P_k receives consent from other neighbors, it assumes the query notification functionalities of P_i (i.e., P_k becomes the new beacon node of the queries that were registered at P_i) and notifies the source nodes of the queries about the takeover. Failures are detected through periodic exchange of heartbeat messages between the beacon node and the peers maintaining its replicas. The beacon node that does not respond to two consecutive messages is assumed to have failed.




Fig. 8:

D. Load Balancing

Achieving good load distribution among peers is another important requirement for the performance of the CoQUOS system. The number of queries and the numbers of notifications sent out per unit time by various nodes represent two key load metrics for CoQUOS system. These load parameters can vary widely among the nodes of the CoQUOS system due to variety of reasons, including

topological characteristics of the network, skewed announcement and query popularities, variation in the resource availabilities at peers or a combination of these factors. Irrespective of the cause, load imbalances not only degrade the performance of the system, but may also cause overloaded peers to exit the network. Ensuring good load balancing in decentralized, loosely coupled systems such as unstructured P2P overlays is challenging. In fact, achieving optimal load balancing on a global scale may turn out to be prohibitively expensive as it would require collection and maintenance of load information on a global scale. Skip graphs have been used to alleviate some of the problems associated with global load balancing in the context of range data [16]. However, skip graphs require maintenance of circular linked lists consisting of all nodes in the system. Hence, they are not suitable for a large-scale, dynamic environment like unstructured P2P content distribution platforms.




Fig. 9:

Considering the above discussion, our goal for the CoQUOS system is not to achieve perfect load balancing. Rather, our aim is to ensure good load balancing in an efficient and scalable manner. The CoQUOS system includes two schemes for balancing notification loads among its peers. Each of these schemes can be used exclusively, or they may be used simultaneously. Keeping in mind the decentralized and loosely coupled nature of the P2P network, these techniques have been designed to be localized both in terms of their operations and the information they utilize.

They only require interactions among neighbouring peers, thus making them suitable for generic unstructured P2P networks. In both schemes, neighboring peers periodically (at the end of prespecified cycles) exchange information about their loads. Based on the load information obtained from its neighbors, a peer decides whether it is overloaded. A peer P_i is overloaded if the ratio of P_i 's load in the previous cycle to P_i 's cumulative load on the peer P_i in the previous cycle due to all the queries registered at P_i . Let denote the set of queries that would be offloaded at the end of the cycle. Q_i are formed such that

VII. Conclusion

when new matching items appear can significantly improve their utility and effectiveness. While the pub-sub paradigm can provide this capability, implementing pub-sub systems on unstructured overlays is often a very complex endeavor. The continuous query paradigm studied in this paper is similar to pub-sub, but it provides best-effort notification service. We presented the design and evaluation of a lightweight system, called CoQUOS, which

supports continuous queries in unstructured P2P networks. The CoQUOS system incorporates several novel features such as cluster resilient random walk for query propagation, dynamic probability scheme for query registration, and a lazy replication technique for countering network churn. The proposed techniques have been evaluated through theoretical analysis and simulation based experiments.

References

- [1] J. Katz, M. Yung, "Scalable Protocols for Authenticated Group Key Exchange", *J. Cryptology*, Vol. 20, pp. 85-113, 2007.
- [2] W.S. Ng, B.C. Ooi, K.-L. Tan, A. Zhou, "PeerDB: A P2P-Based System for Distributed Data Sharing", *Proc. 19th Int'l Conf. Data Eng. (ICDE)*, 2003.
- [3] G. Banavar, T. Chandra, B. Mukherjee, J. Nagarajao, R.E. Strom, D.C. Sturman, "An Efficient Multicast Protocol for ContentBased Publish-Subscribe Systems", *Proc. 19th IEEE Int'l Conf. Distributed Computing Systems (ICDCS '99)*, 1999.
- [4] T. Bu, D.F. Towsley, "On Distinguishing between Internet Power Law Topology Generators", *Proc. IEEE INFOCOM*, 2002.
- [5] A. Carzaniga, D.S. Rosenblum, A.L. Wolf, "Design and Evaluation of a Wide-Area Event Notification Service", *ACM Trans. Computer Systems*, Vol. 19, No. 3, pp. 332-383, 2001.
- [6] Y. Chawathe, S. Ratnasamy, L. Breslau, N. Lanham, S. Shenker, "Making Gnutella-Like P2P Systems Scalable", *Proc. ACM SIGCOMM '03*, 2003.
- [7] J. Chen, L. Ramaswamy, A. Meka, "Message Diffusion in Unstructured Overlay Networks", *Proc. Sixth IEEE Int'l Symp. Network Computing and Applications (NCA)*, 2007.
- [8] P.A. Chirita, S. Idreos, M. Koubarakis, W. Nejdl, "Publish/Subscribe for RDF-Based P2P Networks", *Proc. First European Semantic Web Symp.*, May 2004.
- [9] M. Zhong, K. Shen, "Popularity-Biased Random Walks for Peer-to-Peer Search under the Square-Root Principle", *Proc. Fifth Int'l Workshop Peer-to-Peer Systems (IPTPS)*, 2006.
- [10] P.T. Eugster, R. Guerraoui, C.H. Damm, "On Objects and Events", *Proc. 16th ACM SIGPLAN Conf. Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, 2001.
- [11] P. Ganesan, M. Bawa, H. Garcia-Molina, "Online Balancing of Range-Partitioned Data with Applications to Peer-to-Peer Systems", *Proc. 30th Int'l Conf. Very Large Data Bases (VLDB)*, 2004.
- [12] B. Gedik, L. Liu, "A Scalable Peer-to-Peer Architecture for Distributed Information Monitoring Applications", *IEEE Trans. Computers*, Vol. 54, No. 6, pp. 767-782, June 2005.
- [13] C. Gkantsidis, M. Mihail, A. Saberi, "Random Walks in Peer-to-Peer Networks", *Proc. IEEE INFOCOM '04*, 2004.
- [14] C. Gkantsidis, M. Mihail, A. Saberi, "Hybrid Search Schemes for Unstructured Peer-to-Peer Networks", *Proc. IEEE INFOCOM*, 2005.
- [15] A. Gupta, O.D. Sahin, D. Agrawal, A.E. Abbadi, "Meghdoot: Content-Based Publish/Subscribe over P2P Networks", *Proc. ACM/IFIP/USENIX Int'l Middleware Conf. '04*, 2004.

- [16] V. Kalogeraki, D. Gunopoulos, D. Zeinalipour-Yazti, "A Local Search Mechanism for Peer-to-Peer Networks", Proc. 11th Int'l Conf. Information and Knowledge Management (CIKM), 2002.
- [17] J. Kannan, B. Yang, S. Shenker, P. Sharma, S. Banerjee, S. Basu, S.-J. Lee, "SmartSeer: Using a DHT to Process Continuous Queries over Peer-to-Peer Networks", Proc. IEEE INFOCOM, 2006.
- [18] Q. Lv, P. Cao, E. Cohen, K. Li, S. Shenker, "Search and Replication in Unstructured Peer-to-Peer Networks", Proc. 16th Ann. ACM Int'l Conf. Supercomputing, 2002.
- [19] B. Arai, G. Das, D. Gunopulos, V. Kalogeraki, "Approximating Aggregation Queries in Peer-to-Peer Networks", Proc. 22nd Int'l Conf. Data Eng. (ICDE), 2006.
- [20] L. Ramaswamy, J. Chen, P. Parate, "CoQUOS: Lightweight Support for Continuous Queries in Unstructured Overlays", Proc. IEEE Int'l Parallel and Distributed Processing Symp., 2007.
- [21] S. Ratnasamy, P. Francis, M. Handley, R. Karp, S. Schenker, "A Scalable Content-Addressable Network", Proc. ACM SIGCOMM '01, Aug. 2001.
- [22] P. Reynolds, A. Vahdat, "Efficient Peer-to-Peer Keyword Searching", Proc. ACM/IFIP/USENIX Int'l Middleware Conf. '03, 2003.
- [23] B. Segall, D. Arnold, J. Boot, M. Henderson, T. Phelps, "Content Based Routing with Elvin4", Proc. Australian Unix System User Group Conf., 2000.
- [24] K. Sripanidkulchai, "The Popularity of Gnutella Queries and Its Implications on Scalability", O'Reilly's www.openp2p.com website. Feb. 2001.
- [25] I. Stoica, R. Morris, D. Karger, M.F. Kaashoek, H. Balakrishnan, "Chord: A Scalable Peer-to-Peer Lookup Service for Internet Applications", Proc. ACM SIGCOMM '01, Aug. 2001.
- [26] P. Triantafillou, I. Aekaterinidis, "Content-Based Publish/Subscribe over Structured P2P Networks", Proc. Third Int'l Workshop Distributed Event-Based Systems (DEBS), 2004.
- [27] P.A. Chirita, S. Idreos, M. Koubarakis, W. Nejdl, "Designing Semantic Publish/Subscribe Networks Using Super-Peers", Semantic Web and Peer-to-Peer, Springer Verlag, 2005.
- [28] R. Baldoni, C. Marchetti, A. Virgillito, R. Vitenberg, "Content Based Publish-Subscribe over Structured Overlay Networks", Proc. 25th IEEE Int'l Conf. Distributed Computing Systems (ICDCS), 2005.



Mr. M. ChandraSekhar Varma is a student of Akula Gopayya College of Engineering & Technology, Tadepalligudem. Presently he is pursuing his M.Tech (C.S.E) from this college and he received his B.Tech (CS&IT) from Sri Vasavi Engineering College, Tadepalligudem. His area of interest includes Data Warehouse and Data Mining, information security, Data Communications & Networks and other

advances in Computer Applications.



Mr. M. M. Balakrishna well known Author and excellent teacher Received M.Tech (CSE). He is working as a H.O.D for the Department of M.Tech., Computer Science & Engineering, Akula Gopayya College of Engineering & Technology, Tadepalligudem. He has vast teaching experience in engineering college. To his credit couple of publications both National & International conferences /

journals. His area of Interest includes Data Warehouse and Data Mining, information security, Data Communications & Networks and other advances in Computer Applications. He has guided many projects for Engineering Students.