

Content-Based Peer-to-Peer Network Overlay for Full-Text Federated Search

Jie Lu Jamie Callan

Language Technologies Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA
{jielu, callan}@cs.cmu.edu

Abstract

Peer-to-peer network overlays have mostly been designed to support search over document names, identifiers, or keywords from a small or controlled vocabulary. In this paper we propose a content-based P2P network overlay for full-text federated search over heterogeneous, open-domain contents. Local algorithms are developed to dynamically construct a network overlay with content-based locality and content-based small-world properties. Experimental results using P2P testbeds of real documents demonstrate the effectiveness of our approach.

1. INTRODUCTION

Peer-to-peer (P2P) networks integrate autonomous computing resources without requiring a central coordinating authority, which makes them a potentially robust and scalable model for providing federated search to large-scale networks of text digital libraries. Search in current P2P networks has mostly relied on matching over document names, identifiers, or keywords from a small or controlled vocabulary. Recently, a set of solutions to *full-text federated search* in P2P networks has been proposed that searches the full body of text documents and returns results with relevance-based rankings (Lu & Callan, 2006a). Multiple regional directory services (“*hubs*”) in a hierarchical P2P architecture collectively provide resource location and result merging services based on full-text resource representations, and act as gateways between peers providing contents (“*information providers*”) and those making information requests (“*information consumers*”).

The main communication channels for federated search are logical connections between peers established at the protocol layer, which build a *network overlay* on top of the underlying physical network. Peer connections in most operational P2P systems either are established in an arbitrary manner, or are determined based on locations of peer identifiers in the key space of the Distributed Hash Table architectures. Search in these network overlays can be inaccurate and inefficient because relevant contents are likely to be scattered, and a decentralized algorithm with only local information may lack sufficient navigational “*cues*” to find the short paths from source to target. Maximizing the power of full-text federated search requires a P2P network overlay with organized content distribution and good network navigability. Previous work on full-text federated search used network overlays that organize contents based on statically generated clusters (Lu & Callan, 2006a), which is not appropriate for dynamic P2P networks. Therefore, we were motivated to develop solutions to dynamic overlay construction for full-text federated search.

An organized content distribution is one that “*clusters*” peers with similar contents together (“*content-based locality*”) so that relevant contents are likely to be concentrated in a small part of the network. Several methods have been developed to construct content-based clusters in P2P network overlays (Crespo & García-Molina, 2002; Schlosser et al., 2002; Khambatti et al., 2002; Löser et al., 2003; Asvanund, 2004; Lu & Callan, 2006a). However, because they either rely on global knowledge for clustering, or use inefficient distributed discovery to group peers with similar

contents, they are most appropriate for small, homogeneous, or closed-domain contents; none of them are suitable for large networks of heterogeneous and open-domain contents.

Good network navigability is often associated with *small-world* properties (Watts & Strogatz, 1998; Kleinberg, 2000), which are characterized by each peer having *local* connections to close peers and a few *long-range* connections to remote peers. Peer distance can be measured based on content similarity (“*content-based small-world*”) or other attributes such as latency. Several approaches exist to construct small-world network overlays. The approach that requires global knowledge about peer distance (Watts & Strogatz, 1998; Kleinberg, 2000; Stoica et al., 2001; Manku et al., 2003; Li et al., 2004) is not appropriate for large-scale and dynamic distributed environments. Methods that select long-range connections uniformly (Sakaryan & Unger, 2003; Merugu et al., 2004) may not enable good navigability because, as (Kleinberg, 2000) points out, long-range connections should be established based on a power-law distribution of peer distance to guarantee good navigability. Repeatedly rewiring existing connections of randomly selected peers to reduce their path lengths in order to converge to small-world properties (Manna & Kabakcioglu, 2003; Clauset & Christopher, 2004) may be quite slow for large-scale P2P networks. None of these approaches are appropriate for constructing small-world overlays with good navigability in large P2P networks.

We refer to a network overlay with both content-based locality and content-based small-world properties as a *content-based P2P network overlay*. In this paper, we propose local algorithms to construct content-based P2P network overlays to support full-text federated search in large networks. In contrast to existing work on network overlays designed to support search over document names, identifiers, or keywords from a small or controlled vocabulary, our work provides effective and efficient solutions for peers to dynamically self-organize into a content-based P2P network overlay without global knowledge or centralized control. This approach is well-matched to networks with heterogeneous, open-domain and often unstructured contents.

For content-based locality, we present a distributed mechanism that connects information providers containing similar contents to the same hub to form content-based clusters, which implicitly defines *content areas*. An adaptive clustering strategy is used to cultivate multiple sub-clusters of information providers within each hub’s content-based cluster, so that the granularity of each content area can be dynamically adjusted according to actual contents and network conditions. Our approach also provides efficiency and scalability by propagating providers’ content information along selective paths instead of broadcasting in the network, and avoiding skewed cluster size distribution to enable better load balance.

For content-based small-world properties, we develop an algorithm that avoids the weaknesses of previous methods by dynamically adjusting local and long-range connections based on limited local views of the network, and establishing long-range connections using a power-law distribution of content distance. It also balances hub degrees to avoid potential bottlenecks to information flow, and reduce the network’s susceptibility to malicious attacks on highly connected hubs.

Experimental results using two P2P testbeds of different contents and sizes are included to demonstrate that better full-text federated search can be achieved in the network overlays constructed using our approach.

In the following section we describe background and related work on P2P network overlays. Our approach to creating a content-based P2P network overlay for full-text federated search is presented in Sections 3 and 4; Section 3 focuses on content-based locality and Section 4 is devoted to content-based small-world properties. Section 5 describes evaluation methodology. We present experimental settings and results in Sections 6 and 7. Section 8 concludes.

2. BACKGROUND AND RELATED WORK

P2P network overlays generally belong to four basic types of P2P architectures: brokered, completely decentralized, hierarchical, and structured P2P architectures. A brokered P2P architecture relies on a single, centralized directory service (“*broker*”) for query routing. In a completely decentralized P2P architecture, every peer functions as directory service and participates in query routing (Gnutella v0.4). A hierarchical P2P architecture (also called “*super-peer network*”) uses multiple peers with more processing power and connection bandwidth for directory services (“*hubs*”, “*super peers*”) and shields the other peers (“*leaves*”) from irrelevant query traffic; leaf peers (providers and consumers) typically only connect with hubs, and hubs connect with leaves and other hubs (Löser et al., 2003; Asvanund, 2004; Lu & Callan, 2006a). A structured P2P architecture partitions the content space and uses a Distributed Hash Table (DHT) to distribute directory information among peers so that each peer provides directory service for a particular region of the content space (Stoica et al., 2001).

In terms of network overlays, a brokered P2P architecture doesn’t provide any flexibility in allowing different overlays since it requires every peer to connect to the centralized directory service. In contrast, in the other three types of P2P architectures, peers have some freedom in choosing their neighbors, and different relations between connected peers regarding their contents (or other attributes) lead to different network overlays. Below we focus on network overlays with content-based locality or small-world properties.

A network overlay is said to exhibit *content-based locality* if peers with similar contents have short paths between them. Because documents relevant to a given query tend to be similar to one another, content-based locality makes locating most relevant contents efficient. Content-based locality in a P2P network is typically established by forming content-based clusters among peers and keeping peers that belong to the same cluster close to one another in the network overlay. There are two main approaches to constructing content-based clusters in P2P network overlays. The first approach uses a global classification hierarchy/ontology to define clusters and assigns peers to clusters based on semantic annotations of their contents. Examples include Semantic Overlay Networks (SONs) in a completely decentralized P2P network (Crespo & García-Molina, 2002), concept clusters in the hypercube network overlay (Schlosser et al., 2002), Semantic Overlay Clusters (SOCs) in a hierarchical P2P network (Löser et al., 2003), and semantic clustering in a structured P2P network (Li et al., 2004). However, clustering based on predefined semantic annotations is not always appropriate for networks containing heterogeneous, open-domain contents. The second approach relies on distributed discovery to find peers with similar contents without explicitly defining clusters. For instance, peers in a completely decentralized P2P network form implicit content-based clusters by periodically exchanging keyword-based descriptions of their contents and looking for overlap in descriptions (Khambatti et al., 2002). To self-organize peers into content-based clusters in a hierarchical P2P network, a method is proposed in (Asvanund, 2004) for each peer to actively seek out new peers and replace its existing neighbors

with those more similar in content. Distributed discovery without guidance is effective, but is slow and incurs high communication cost, which may become a burden for peers with limited resources.

In a network overlay with *small-world properties*, any two peers are likely to be connected through a short sequence of intermediaries. A small-world network overlay is desirable because it not only guarantees the existence of short paths between peers without requiring a large number of connections, but also provides the potential of using a decentralized algorithm with only local information to find these short paths for efficient query routing. Previous work on small-world network overlays can be grouped into three approaches. The first approach assumes that each peer knows its distance (based on content or other attributes) to all other peers, and constructs a small-world network overlay by each peer choosing several closest peers to establish local connections and a few remote peers for long-range connections (Watts & Strogatz, 1998; Kleinberg, 2000; Stoica et al., 2001; Manku et al., 2003; Li et al., 2004). Requiring global knowledge about peer distance makes this approach inappropriate for real P2P networks that are large-scale and dynamic. The second approach adaptively updates each peer's local and long-range connections based on periodically collected information in local network regions without requiring global knowledge (Sakaryan & Unger, 2003; Merugu et al., 2004). However, peers uniformly select among remote peers to establish long-range connections, which provides no theoretic guarantee on network navigability, as discussed in (Kleinberg, 2000). The third approach repeatedly rewires existing connections of randomly selected peers to reduce their path lengths so that the network overlay can converge to one with small-world properties (Manna & Kabakcioglu, 2003; Clauset & Christopher, 2004). The convergence may be quite slow and the conditions for rewiring connections may be too restrictive for real P2P environments.

3. NETWORK OVERLAY WITH CONTENT-BASED LOCALITY

We focus on constructing content-based network overlays in hierarchical P2P architectures designed to support full-text federated search (Lu & Callan, 2006a). Content-based locality is formed at the *hub-provider topology* level by requiring each hub's neighboring providers to form a cohesive content-based cluster. We do not assume the availability of a global classification hierarchy, ontology, or semantic annotations; each content-cluster is described using the full-text resource description of the hub associated with it, obtained by aggregating the resource descriptions of its neighboring providers. To minimize system overhead on constructing the network overlay, the same resource descriptions required by full-text federated search are used here; they describe contents with aggregate bag-of-words representations (unigram language models), and measure similarity based on their K-L divergence values (Lu & Callan, 2006a).

In Sections 3.1-3.4, we present the design of our method for constructing a network overlay with content-based locality, in consideration of the unique characteristics of P2P networks and heterogeneous, open-domain contents. Section 3.5 describes the algorithm in detail.

3.1 Decentralization and Role Differentiation

When a provider requests to join the network, since there is no centralized server to decide which cluster(s) it should join, decisions must be made in a decentralized manner about which hubs the provider can connect to so as to maintain content-based locality in the hub-provider topology. Because hubs have more processing power and connection bandwidth, they play a more active role in constructing network overlay.

To decide the connection(s) for a provider, a joint effort of the provider and the hubs that receive this provider's resource description is required. Each provider is responsible for providing its resource description to the network and making the final decision about which hub(s) to connect to. Each hub that receives the provider's resource description is responsible for calculating the degree of match between the provider's content and the hub's content-based cluster, providing this information to the provider to facilitate its decision making, and propagating the provider's resource description to other hubs.

3.2 Dynamic Adaptive Clustering

Since it is difficult to obtain a partition of the content space beforehand for digital libraries of unstructured text documents in open domains, the *content area* covered by each hub cannot be predetermined. Instead, it can only be determined implicitly by the contents of the providers already connecting to the hub. As the hub accepts into its content-based cluster more providers whose contents are similar to what it already covers, its content area may be updated dynamically to integrate the contents of these new members. Therefore, in contrast to an explicit fixed clustering policy, each hub uses an implicit adaptive clustering criterion, which is more autonomous and self-adjusting.

In an ideal situation, each hub in the network is responsible for covering a specific content area and uses a similarity threshold to decide whether to accept a provider into its cluster and integrate the provider's content into its content area. When a provider's content is sufficiently dissimilar to all existing content areas, an unoccupied hub is contacted to create a new content-based cluster to accommodate this provider. This approach may work well when appropriate threshold values are chosen to control the granularity of the content area covered by each hub so that the number of content areas matches the number of hubs in the network. However, in real, operational environments, because the number of content areas cannot be predetermined for open-domain contents, and the number of hubs is often limited, it is difficult to choose similarity threshold values in advance to satisfy the above condition, particularly when there is no central coordination and control. On the one hand, tighter similarity thresholds result in more cohesive and homogenous clusters representing narrower content areas so that more hubs than those available may be needed to cover the contents in the network. On the other hand, looser threshold values decrease the homogeneity of content-based clusters and thus reduce the degree of content-based locality. To solve this problem, instead of solely relying on a single similarity threshold to determine the granularity of a hub's content area, we use a more flexible clustering strategy which cultivates multiple sub-clusters within each hub's content-based cluster, and spins off a sub-cluster to create a new content area when and only when the sub-cluster has grown to a certain size. With this strategy, while the granularity of each sub-cluster may be relatively fixed by using a similarity threshold, the granularity of each hub's content area can be dynamically adjusted by increasing its number of sub-clusters or transferring its sub-clusters to other hubs. Although such dynamic adaptive clustering may decrease the homogeneity of some content-based clusters, the potentially small reduction in content-based locality is offset by its ability to adjust autonomously based on actual network conditions.

3.3 Load Balancing

Because the amount of contents available in the network is often biased for different topics, hubs that cover popular content areas may be overwhelmed by the hub-provider connections they need

to maintain and the query load they have to handle. To avoid overloading some hubs with popular contents and information requests while wasting other hubs' resources on marginal or unpopular contents, hubs need to average their responsibilities in order to achieve load balance in the network. Hubs' responsibility of serving popular contents can be balanced by transferring sub-clusters of heavily connected hubs to lightly connected ones. The degree of content-based locality is expected to be affected only slightly if hubs take advantage of content-based small-world properties at the hub level to locate appropriate recipient hubs based on the content areas they cover. By using distributed cultivation of sub-clusters, hubs can also share the responsibility of serving unpopular contents instead of dedicating a single hub to handle them.

3.4 Selective Propagation of Providers' Contents

Because the construction of the network overlay requires additional computation and communication costs, it is more cost-efficient to propagate the resource description of each joining provider only to those hubs whose content areas are likely to match the provider's content. Hubs need to know about the contents covered in their neighborhoods¹ so as to decide where to propagate each provider's resource description. Since hubs need to collect neighborhood content information anyway for query routing among hubs (Lu & Callan, 2006a), selective propagation of providers' resource descriptions adds few additional costs. Selective propagation of providers' resource descriptions is also desired for scalability because it shields hubs with dissimilar content areas from unnecessary overhead.

3.5 Algorithm

An information provider's content is described using its full-text resource description. A sub-cluster is represented by the aggregation of the resource descriptions of its provider members. A hub's content area is represented by its resource description, generated by aggregating the representations of its sub-clusters. We use two similarity thresholds to distinguish among three levels of similarity between a provider's content and the contents covered by a sub-cluster: *high*, *marginal*, and *low*.² A provider's first priority is to join the most similar sub-cluster among all the sub-clusters to which it has a *high* similarity level. If it fails to find sub-clusters with *high* similarity, but has at least one sub-cluster with *marginal* similarity, it will join the content-based cluster of the hub that has the most similar sub-cluster with *marginal* similarity by initiating a new sub-cluster at this hub. The distinction between *high* and *marginal* similarity values is for controlling the granularity of each sub-cluster while allowing the granularity of a hub's content-based cluster to dynamically change by including more sub-clusters. If the provider has a *low* similarity level to all existing sub-clusters in the network, it requests an empty³ hub to initiate a new sub-cluster within a new content-based cluster. When all the hubs are non-empty, the provider initiates a new sub-cluster within the cluster of the hub that has the most similar sub-cluster.

The join process of an information provider P proceeds as follows. If it was active in the network before, it first tries to connect to the hubs it previously connected to. If it fails or if it is completely new to the network, it finds an initial set of hubs by querying host-cache servers or by pinging the

¹ A neighborhood of a hub is the set of hubs that a message can reach by following the path from this hub to one of its hub neighbors and further traveling a number of hops.

² Although it is common and often desirable to use global similarity thresholds in cooperative P2P environments, in theory each hub can have its own similarity thresholds adjusted locally (e.g., based on its workload).

³ A hub is "empty" if it doesn't connect to any information provider. Otherwise, it is non-empty.

network. Then it arbitrarily chooses a hub from the list to connect to temporarily, which is responsible for acquiring P 's resource description and starting its propagation among the hubs within a certain radius specified by the TTL (Time-To-Live) of the message. Each non-empty hub that receives P 's resource description executes two operations. First, the hub computes the similarity between P 's resource description and the description of each of its sub-clusters, and sends back to P a message to include the value of the highest similarity along with the similarity level based on its thresholds. Second, the hub selects some neighboring hubs to propagate P 's resource description based on the similarity between the resource descriptions of P and each of its neighborhoods. Each empty hub that receives P 's resource description directly forwards it to its hub neighbors. P collects the messages from the hubs within a predetermined amount of time, and chooses a hub to connect based on the strategies described in the previous paragraph. It is straightforward to extend the strategies if P is allowed to join multiple content-based clusters.

If the size of a sub-cluster within a hub's content-based cluster exceeds a certain limit to indicate the emergence of a new content area, the hub propagates a message among the hubs requesting an empty hub to take over. The transfer of the sub-cluster is conducted by connecting the provider members of the sub-cluster to the chosen empty hub to generate a new content-based cluster, and disconnecting them from the original hub. The attempt of initiating a new content area fails if no empty hub is available and the hub may try again sometime later.

If a hub becomes heavily connected due to the large number or sizes of its sub-clusters, it may propagate the descriptions of its sub-clusters towards selective directions at the hub level based on neighborhood descriptions to request other hubs covering similar content areas to share the load, and transfer one or more of its sub-clusters to the willing hubs. Alternatively, a hub may pose a limit on the size of its content-based cluster and stop accepting new providers once the limit has been reached. The former approach may result in higher degree of content-based locality at the expense of higher communication costs than the latter approach.

4. NETWORK OVERLAY WITH CONTENT-BASED SMALL-WORLD PROPERTIES

Content-based small-world is constructed at the *hub-hub topology* level by each hub establishing *local* connections to hubs with similar content areas and a few *long-range* connections to hubs with dissimilar content areas. Our method of constructing network overlay with content-based small-world properties has a number of specific features that distinguish it from other methods developed for small-world network overlays. Sections 4.1-4.3 describe these features; Section 4.4 describes the algorithm that implements the method.

4.1 Dynamic Adaptation

A P2P network is dynamic in nature; hubs may join and leave the network, the content area of each hub may change over time as providers connect to and disconnect from it, and new content areas may emerge in the network. Hence each hub needs to adjust its connections to other hubs periodically to accommodate these changes and maintain content-based small-world properties.

4.2 Utilization of Limited Local Information

Due to the decentralized nature of a P2P network, no global information about the content distance between peers is readily available, and acquiring such global information is often inadmissible

because of high cost. Therefore, each hub can only utilize the content information of other hubs in its local neighborhood to find similar (close) and dissimilar (remote) hubs.

4.3 Degree Balancing

Because connections at the hub level are major channels for query routing and propagation of resource descriptions, hubs that are highly connected may become potential bottlenecks which will restrict the information flow in the network. In addition, an attack resulting in the removal of these highly connected hubs could cripple the network. In our approach, special effort is made to balance connection degrees without undermining content-based small-world properties.

4.4 Algorithm

The dynamic construction of a content-based small-world network overlay at the hub-hub topology level proceeds as follows. When a hub H joins the network, if it was active in the network before, it first tries to connect to the hubs it previously connected to. If it fails or if it is completely new to the network, it obtains a list of existing hubs in the network by querying host-cache servers or by pinging the network. Because it doesn't have any providers connecting to it yet, it has an empty resource description. Since the similarity between a hub with an empty resource description and any other hub is undefined, H can only randomly choose its hub neighbors at the moment. H 's resource description is initialized when it responds to a provider's join request or a hub's transfer request to start a new content-based cluster and connect to the joining provider or the providers in the transferred sub-cluster.

Each non-empty hub operates independently in a decentralized manner to select its own hub neighbors based on its local view of the content areas available in the network. Given the dynamic conditions of a P2P network, a hub H periodically evaluates its content similarity to its direct hub neighbors and their direct hub neighbors (i.e., hubs within two hops from it) using the hubs' resource descriptions exchanged among them, and adjusts its outgoing connections to link to several most similar hubs and a few dissimilar hubs using the following procedure:

1. H uses a threshold ρ^* to distinguish between similar and dissimilar hubs among the hubs within two hops from it;
2. H connects to M_{ol} most similar hubs whose incoming hub connection capacities have not reached their maximally allowed values M_i , where M_{ol} is the maximum number of outgoing local hub connections H can have;
3. If all of H 's similar hubs have reached their maximum incoming hub connection capacities M_i , H requests them to recommend their similar hub neighbors, which may be repeated recursively until H establishes at least one local hub connection or the number of requests reaches a limit before it succeeds in finding any similar hub available for connection; and
4. H selects M_{og} dissimilar hubs that have not reached maximum incoming hub connection capacities M_i with probability:

$$P(H_{i \in \{j: NH(H, H_j) \leq 2\}} | CD(H, H_i) \geq \rho^*) = cCD(H, H_i)^{-\beta}$$

where M_{og} is the maximum number of outgoing long-range ("global") hub connections H can have, NH is the number of hops between hubs, CD is the content distance (the inverse of content similarity) between hubs, calculated based on the K-L divergence between hubs' resource descriptions, β is an exponent that essentially controls the "distance scale" of long-

range connections (i.e., smaller β biases towards greater content distance and thus more dissimilar hubs, and larger β biases towards smaller content distance and thus less dissimilar hubs), and c is a normalizing constant.

By periodically adapting each hub's outgoing⁴ connections at the hub level, the hub-hub topology effectively maintains content-based small-world properties in a dynamic environment.⁵ Since each hub adjusts its connections only based on its local knowledge of the hubs that are located within two hops from it and possibly their recommended local contacts, no global information or control are necessary for overlay construction. The step of limiting each hub's connection capacity and recommending other similar hubs when its own connection capacity becomes full helps in distributing connections at the hub level in a less skewed manner to avoid concentrating a large number of connections at a few hubs. Establishing long-range connections based on a power-law distribution of content similarity enables each hub to have nearly uniformly distributed long-range hub connections over all "distance scales" ("similarity scales"), which allows hubs to route any query efficiently towards its targeted content area (Kleinberg, 2000).

5. EVALUATION METHODOLOGY

We used an old P2P testbed created from the TREC WT10g dataset (Lu & Callan, 2003) and a new P2P testbed created from the .GOV2 test collection (Lu & Callan, 2006b) to evaluate the performance of our approach to constructing a content-based P2P network overlay. 2,500 collections, each consisting of documents crawled from a real Web site, were extracted from the TREC WT10g dataset to define 2,500 information providers in a medium-sized hierarchical P2P network with a total number of 1,421,088 documents. The same URL-based partitioning approach was used to divide the .GOV2 dataset, and 25,000 collections with a total number of 11,218,349 documents were selected to define the information providers in a large P2P network. The number of hubs was chosen to be 32 for the medium-sized network. The number of hubs in the large-sized network was 256, eight times as many as the number of hubs in the medium-sized network.

The queries provided by the U. S. National Institute for Standards and Technology (NIST) for the WT10g and .GOV2 test collections are TREC topics 451-550 and 701-800 respectively with the standard TREC relevance assessments. We refer to the title fields of these queries used in our experiments as *TREC queries*. Two larger sets of queries were also used to evaluate federated search performance on more queries. For the medium-sized network, because there are no available user queries targeted specifically at the collections selected from WT10g, key terms were extracted from the documents in WT10g and used as queries (Lu & Callan, 2003). 1,000 queries were selected from these automatically generated queries, which we refer to as *WT10g queries*. For the large-sized network, we selected 1,000 queries from a query set provided by AOL, which were queries with user-clicked results from the *.gov domain, collected during a one-month period on AOL Search in April 2006. We refer to these queries as *GOV queries*.

⁴ The directions of hub-hub connections are only used for overlay construction. They are ignored when hub-hub connections are used as data channels to exchange messages between hubs.

⁵ Although each hub only chooses among the hubs in its local network region to establish local and long-range connections, experimental results show that hubs are able to connect to "globally" close and remote hubs because periodic connection adjustment keeps bringing new hubs to local network regions.

The methods described in Sections 3 and 4 were applied to construct a content-based network overlay with regulated connections at both the hub-provider topology level and the hub-hub topology level. The quality of the constructed network overlay was first evaluated by verifying whether it exhibited the targeted properties, specifically content-based locality, content-based small-world, and balanced connection degrees.

One measure of content-based locality was the cohesion of content-based clusters. The K-L divergence between the descriptions of each information provider and its connecting hub was calculated, and the distribution of the calculated divergence values for all connecting provider-hub pairs was used to evaluate the degree of cohesion of content-based clusters. The second measure we used for content-based locality was the degree of concentration of providers with relevant contents (“*relevant providers*”) among different hubs. Given a set of queries with relevance judgments, we ranked the hubs by the number of relevant providers contained in their clusters, and defined \hat{R}_n to be the percentage of the relevant providers that had been accumulated via the n top-ranked hubs. The values of \hat{R}_n for different queries were averaged. A similar metric has been used to evaluate the effectiveness of resource selection in traditional distributed information retrieval (French et al, 1998). If relevant providers were highly concentrated, then only a small number of hubs were needed to cover them, implying that \hat{R}_n would have large values for small n .

Content-based small-world was measured by the association between content similarities and path lengths of hub pairs. Both the content distance (the K-L divergence between hub descriptions) and the path length (number of hops) between each pair of hubs were calculated, and the relation between content distances of hub pairs and their path lengths was studied to see if the average path length between any pair of hubs was small irrespective of their content distance, a characteristic of content-based small-world.

To evaluate the effectiveness of our attempts on load balancing, we also included in evaluation the distribution of cluster sizes and the degree distribution at the hub level.

The ultimate evaluation of the constructed content-based network overlay was its support for federated search. The performance of federated search was measured by search accuracy as well as efficiency. For each query’s result, its precisions at ranks 1-30 were averaged to get the *average precision over ranks 1-30*. The average precisions for the results of a set of queries were further averaged. A similar measure is used in (Stenmark, 2005) to evaluate ranked retrieval performance. *Set-based recall* (the percentage of the relevant documents returned) was another measure used to evaluate accuracy.

Measuring search accuracy required relevance judgments. The standard relevance assessments supplied by NIST were used for TREC queries. For WT10g and GOV queries without human relevance judgments, we chose to use the retrieval results from a single large collection as the baseline (“*single collection*” baseline) to measure how well federated search in P2P networks could locate those documents considered very relevant by centralized search. The single large collection was constructed by aggregating all of the providers’ contents in the P2P network. The top-ranked documents retrieved from this single large collection for a query were treated as the set of “relevant” documents for this query. The numbers of “relevant” documents per query were 50 for the medium-sized network and 200 for the large-sized network, which were chosen based on the average numbers of relevant documents for TREC queries 451-550 and 701-800 respectively.

Because evaluation based on the “single collection” baseline essentially measured the percentage of overlap between the documents returned by centralized search and those by federated search, we refer to the measures calculated using the “single collection” baseline as “*overlap precision*” or “*overlap recall*” to distinguish them from those using real relevance judgments. Although this methodology is not ideal, it is not unreasonable because distributed retrieval systems are not yet better than the “single collection” baseline, so the ability of a P2P network to mimic a good centralized search engine is an acceptable indicator of its performance.

The average numbers of (real or overlap) relevant providers for TREC queries 451-550 and WT10g queries in the medium-sized network were 25 and 4, and for TREC queries 701-800 and GOV queries in the large-sized network were 99 and 68, respectively. Federated search performance was affected by the degree of concentration of relevant contents.

Search efficiency was measured by the average percentage of the network reached by the query messages for each query; smaller values are preferred.

6. EXPERIMENTAL SETTINGS

The mechanism described in (Lu & Callan, 2006a) was used to conduct full-text federated search in the tested P2P networks. Here we briefly describe the search process. Each query was issued by a consumer to an initially selected hub. A Time-To-Live (TTL) field in each query message determined the maximum number of times it could be relayed (“*search radius*”). Based on the collected full-text resource descriptions, a hub that received the query used its resource selection algorithm to rank its neighboring providers by their likelihood of providing relevant content, and neighboring hubs by their likelihood of providing a path to relevant content. Each hub routed the query to the top-ranked neighboring providers selected using a learned threshold, and one top-ranked hub neighbor. A provider that received the query used its full-text document retrieval algorithm on its local collection index to generate a relevance-based ranking of its documents, and responded with a list of the 50 top-ranked documents. Each reached hub was responsible for collecting the ranked lists returned by multiple providers, using its result merging algorithm to merge them into a single, integrated ranked list, and returning it to the consumer. Finally, a consumer merged results returned by multiple hubs. See (Lu & Callan, 2006a) for details.

The network was initialized to be an overlay of empty hubs (32 for the medium-sized network and 256 for the large-sized network) randomly connecting with one another without any providers. To simplify the overlay construction process, information providers were assumed to join the network one at a time in a random order.⁶ For simplicity, each hub adjusted its hub connections after every 100 providers joined the network.

Two thresholds were used to categorize the similarity between a provider’s content and the contents covered by a sub-cluster (measured by the negative of the K-L divergence-based content distance between their full-text resource descriptions). The *high-marginal threshold* was used to distinguish between high and marginal similarity values in order to determine whether to accept the provider into an existing sub-cluster, or to accommodate the provider by initializing a new sub-

⁶ Experimental results not included in this paper show that different random initial hub connections and random orders of providers joining the network don’t have a significant impact on the properties of the constructed network overlay, although actual connections in the overlay may vary.

cluster within an existing content-based cluster. The *marginal-low threshold* was used to distinguish between marginal and low similarity values so that a provider with a low similarity to all existing sub-clusters could initiate a new content-based cluster when there was an empty hub available. The high-marginal threshold was set to -1.0 and the marginal-low threshold was -2.0.

To avoid propagating a provider’s resource description in the network indefinitely, each message carrying a provider’s description had a small TTL value so that it was only propagated within a certain radius from the initial hub. The TTL was set to 4 since most of the shortest path lengths between hubs were expected to be no more than 4. The threshold used by a hub to determine which neighboring hubs to select for further propagation of a provider’s description (based on the K-L divergence between the provider’s description and neighborhood descriptions) was 1.5.

When there was an empty hub available, a sub-cluster exceeding a certain size limit was spun off to create a new content-based cluster in representing a new emerging content area. 5 was used for the spin-off size limit of a sub-cluster. The size of the content-based cluster at each hub also had a limit for the purpose of load balance. A hub stopped accepting new providers once the limit was reached. Its value was empirically set to 275, which was roughly three times as large as the average cluster size in the networks.

Because in reality large-scale P2P networks are typically sparse, and it is quite expensive for each hub to acquire and maintain neighborhood descriptions for a large number of hub neighbors, we chose small values for various maximum connection capacities to simulate the overlay construction of a sparse network. For the medium-sized network, a hub’s maximum number of outgoing local hub connections M_{ol} was 2, its maximum number of outgoing long-range hub connections M_{og} was 3 minus its actual number of outgoing local hub connections, and its maximum number of incoming hub connections M_i was 3. For the large-sized network, the values of M_{ol} , M_{og} , and M_i were 4, 4, and 12 respectively. The exponent β for the power-law distribution was 2.0. Simple cycle detection was used to avoid cycles of length 3 in hub-hub connections to improve the efficiency of query routing.

7. EXPERIMENTAL RESULTS

In this section we present experimental results to demonstrate the effectiveness of our approach to content-based network overlay construction. Figure 7.1 (a) shows the distribution of the K-L divergence values for all connecting provider-hub pairs in the dynamically constructed content-based overlay of the medium-sized network. As a comparison, the figure also includes the within-cluster divergence distributions of a random network overlay⁷ and an overlay created by statically clustering all 2,500 information providers into 32 clusters (equal to the number of hubs in the network) using the global K-means clustering algorithm. Figure 7.1 (b) plots the within-cluster divergence distributions of content-based and random network overlays for the large-sized P2P network. Due to the computational difficulty on clustering 25,000 providers with large-sized representations, the overlay based on static clustering was not generated for the large-sized network. The figures show that the content-based overlays had smaller means in the distributions of within-cluster divergence values than the random overlays, indicating a higher degree of

⁷ For simplicity, random network overlays were constructed by using the global knowledge about the identities of all peers in the network. This is a best-case scenario for random network overlays.

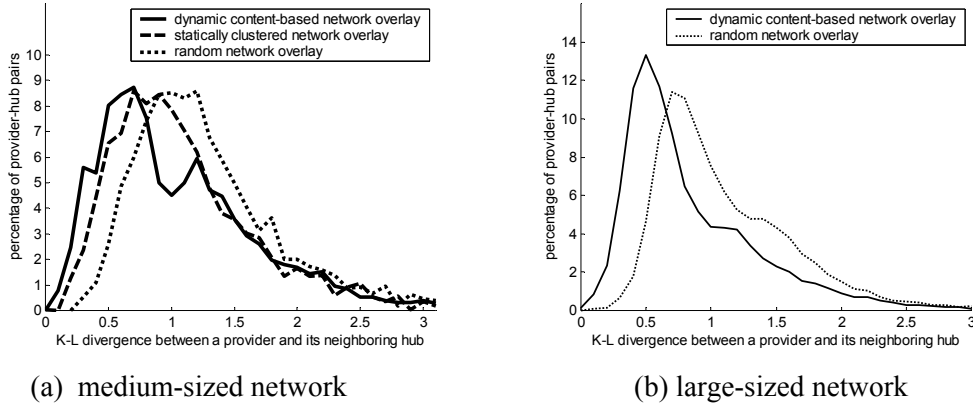


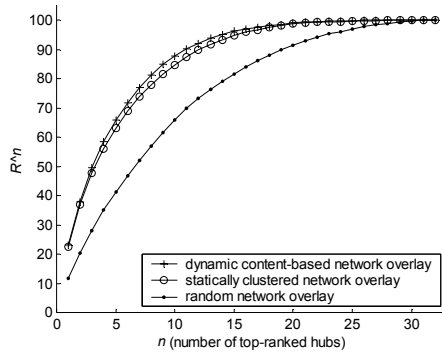
Figure 7.1: The within-cluster divergence distributions of different overlays.

cohesion for content-based clusters than random clusters. The humps between the K-L divergence values of 1.0 and 1.5 and the long tails for the distribution curves of the content-based overlays were due to the existence of provider members with marginally similar or remotely related contents in each cluster, resulting from the diverse nature of the contents and the hubs' shared responsibility for serving providers with unpopular contents. Figure 7.1 (a) also shows that the content-based clusters constructed dynamically and those created using the global K-means clustering algorithm had similar degree of cohesion, demonstrating that our incremental, adaptive clustering approach was able to generate clusters with a satisfactory degree of content-based locality.

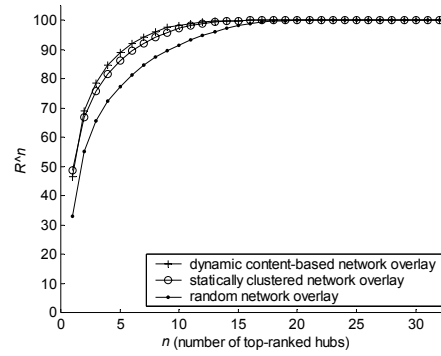
Figures 7.2 (a)-(d) depict the percentages of relevant providers accumulated via increasing numbers of hubs (R_n) for different network overlays, using different sets of queries in two network sizes. From the figures we can see that the degree of relevant content concentration in the content-based network overlay was consistently higher than that in a random overlay of the same size. The dynamically constructed and the statically clustered network overlays yielded similar relevant content concentration, again illustrating that our approach was able to generate a network overlay with desired content-based locality.

Figures 7.3 (a)-(b) show the cluster size distributions of the dynamically constructed and the statically clustered network overlays in the medium-sized network. Compared with the distribution plotted in Figure 7.3 (a) for the statically clustered network overlay, the distribution in Figure 7.3 (b) for the dynamically constructed content-based overlay was less skewed, as evident by its smaller percentages of small and large clusters, and larger percentages of medium-sized clusters.


Figure 7.3 (c) depicts the cluster size distribution of the content-based overlay in the large-sized network. As a comparison, we applied to the large-sized network the same dynamic overlay construction *without* limiting cluster sizes or transferring sub-clusters from heavily connected hubs to lightly connected ones, and the resulting cluster size distribution is shown in Figure 7.3 (d). The differences between Figure 7.3 (c) and Figure 7.3 (d) indicate that without the steps of load balancing, dynamic overlay construction based on distributed, adaptive clustering criteria may overload a few hubs with a large number of providers sharing popular contents while wasting the resources of many other hubs on serving a small number of providers with less-than-popular contents. The hump in Figure 7.3 (c) between cluster sizes of 250 and 300 was due to the cluster size limit of 275 we used. Load balancing together with distributed cultivation of sub-clusters was



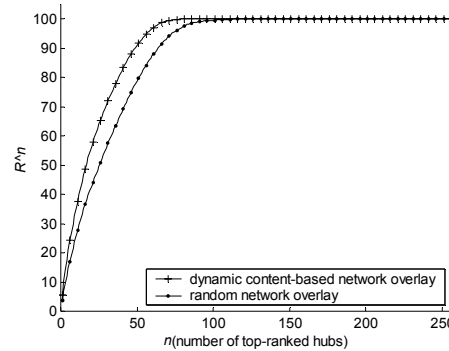
(a) TREC queries 451-550, medium-sized



(b) WT10g queries, medium-sized



(c) TREC queries 701-800, large-sized



(d) GOV queries, large-sized

Figure 7.2: The cumulative distributions of the relevant providers among hubs for different overlays.

able to spread the responsibility of serving popular contents to more hubs and reduce the number of hubs solely occupied by unpopular contents, resulting in a less skewed cluster size distribution.

Figure 7.4 plots the relations between content distances and path lengths for hub pairs. The range of content distances between hubs was much smaller in the medium-sized network due to its small number of hubs with less diverse content areas. The figure shows not only small average path lengths between hubs with similar content areas (K-L divergence no more than 1.0), but also small-to-medium path lengths between hubs with dissimilar content areas. Particularly, for the large-sized network with 256 hubs, on average 10 hub neighbors per hub, and a network density of 0.0388, the average path lengths didn't exceed 4 (the maximum path length was 5) for hub pairs with very large content distances. Therefore the content-based network overlays had locational proximity of similar content areas and short global separation of dissimilar content areas, demonstrating that they exhibited content-based small-world properties.

Figure 7.5 depicts the degree distributions at the hub level for the content-based overlays. The degree distribution of the medium-sized network shows nothing interesting due to its small scale. The degree distribution of the large-sized network illustrates that the dynamically constructed overlay for the network of 256 hubs didn't have a power-law distribution. This was due to the constraints on the minimum and maximum numbers of hub connections for each hub, which was designed on purpose to avoid overloading a few hubs with a large number of connections by distributing these connections to a larger number of hubs. When the network scales to an even larger size, the middle range of the distribution curve is expected to behave more and more like a

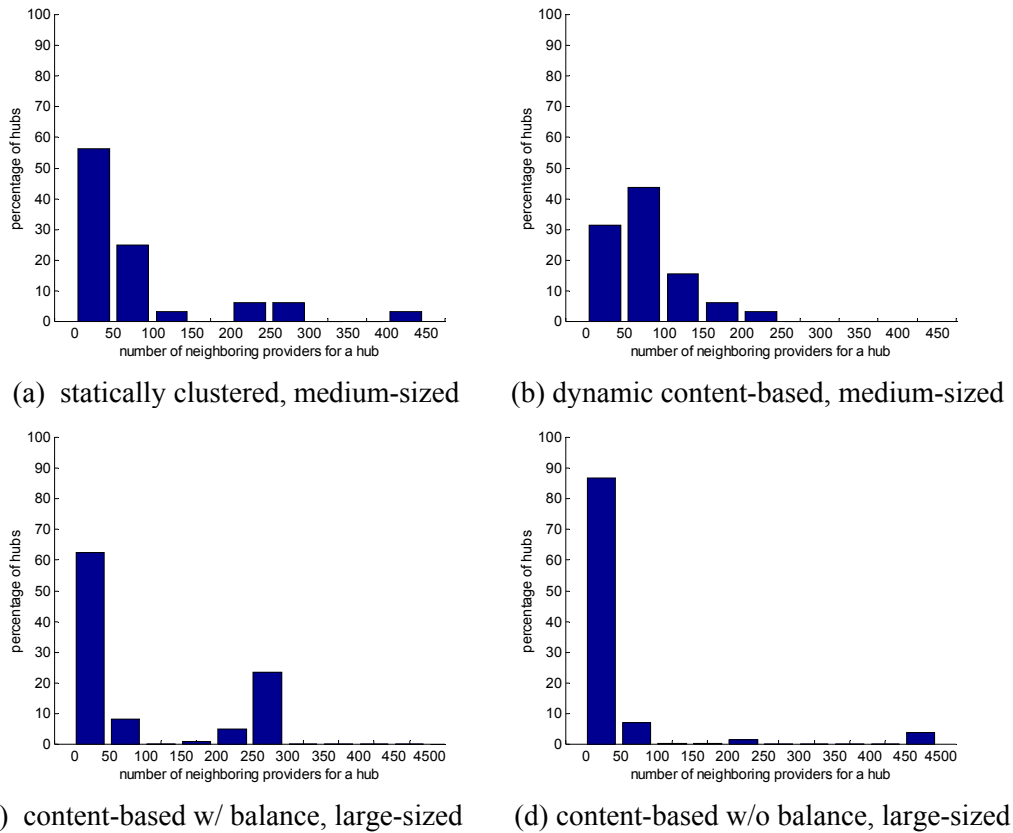


Figure 7.3: The distributions of cluster sizes for different overlays.

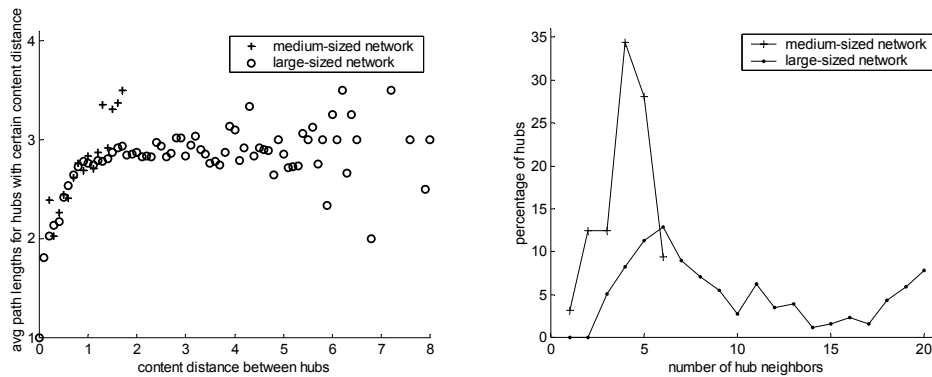






Figure 7.4: Content distances vs. path lengths. Figure 7.5: Hub-hub degree distributions.

power-law distribution, but instead of a long, diminishing tail to the right side which is typical of a power-law distribution, the curve would turn upward before finishing at the maximum allowed connection capacity. Balancing hub degrees by truncating the power-law degree distribution is not a weakness but a strength of our approach, because it avoids a few very highly connected peers at the tail of a power-law degree distribution which can easily become bottlenecks to information flow or targets of malicious attacks.

For a network overlay with content-based locality and content-based small-world properties, when an information request is initiated from a hub located in a network region containing relevant contents, most relevant documents should be covered by nearby hubs due to local clustering of



(a) precision, start from non-relevant regions (b) recall, start from non-relevant regions


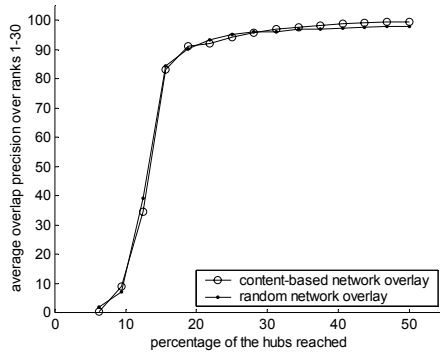


(c) precision, start from relevant regions (d) recall, start from relevant regions



Figure 7.6: The search performance for TREC queries 451-550 in the medium-sized network.

hubs with similar content areas. When an information request is initiated from a non-relevant region, the request should only need to travel along a short path to reach a relevant region thanks to the short path length between hubs of dissimilar content areas. To verify the effectiveness of the content-based network overlay in both cases, federated search performance was evaluated for queries starting from both relevant and non-relevant network regions.

Figures 7.6-7.9 show the performance of federated search in the content-based network overlays. Random network overlays were chosen as the baseline because they are most popular in operational P2P systems, and no existing approaches are applicable to constructing content-based overlays in P2P networks containing heterogeneous, open-domain contents without semantic annotations and global control. To start search in a relevant region, each query was initiated by a consumer connected to a hub which was nearest on average to relevant content. To start search in a non-relevant region, each query was issued by a consumer connected to a hub located farthest on average from relevant content. The figures indicate that when queries were initiated from relevant regions, because similar contents (and likely more relevant contents) were near to one another in the content-based network overlays but most certainly not so in the random overlays, federated search in the content-based overlays had better performance right from the beginning compared with search in the random overlays. When queries were initiated from non-relevant regions, search in the content-based overlays still outperformed that in the random overlays not only because short global separation between dissimilar content areas enabled query routing to reach relevant content regions quickly, but also because content-based locality and locational proximity of similar content



(a) overlap precision, start from non-relevant regions (b) overlap recall, start from non-relevant regions


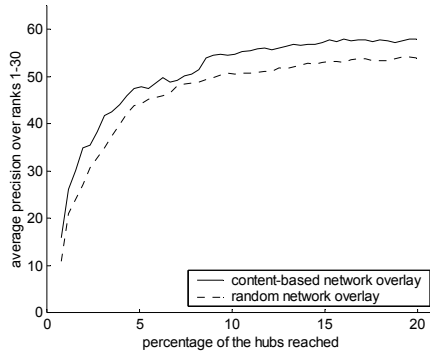


(c) overlap precision, start from relevant regions (d) overlap recall, start from relevant regions



Figure 7.7: The search performance for WT10g queries in the medium-sized network.

areas resulted in more relevant contents to be located nearby after a query arrived at a relevant content region. The analysis of the results for individual queries revealed that, compared with the random overlays, the relative performance improvement of the content-based overlays was greater when relevant documents were scattered in a large number of providers (TREC queries and GOV queries), but marginal when relevant contents were naturally concentrated at just a few providers (WT10g queries).

To focus on the navigation of queries at the hub level, the figures plot search performance against the percentage of the hubs reached by query messages. However, query messages didn't stop at hubs but continued to travel from the reached hubs to their selected neighboring providers. In consideration of efficiency, only a small number of neighboring providers were selected, which means that the relevant documents contained in the providers not selected could not be retrieved. Therefore, compared with the random overlays, although the content-based overlays enabled federated search to reach the hubs covering more relevant providers, the amount of relevant documents that were actually retrieved might have a smaller difference. This doesn't imply that the gain in performance was not worth the effort of content-based overlay construction. In fact, under the same efficiency requirement, federated search in the content-based overlays returned many fewer non-relevant documents in total, which was not reflected in the measures of average precision over ranks 1-30 and recall. Furthermore, the content-based overlays improved the robustness of full-text federated search conducted using dated resource descriptions, because with organized content distribution, the resource descriptions of hubs and neighborhoods were more resilient to changes in content due to peer arrivals and departures.



(a) precision, start from non-relevant regions (b) recall, start from non-relevant regions



(c) precision, start from relevant regions (d) recall, start from relevant regions


Figure 7.8: The search performance for TREC queries 701-800 in the large-sized network.

8. CONCLUSIONS


Current peer-to-peer network overlays have mostly been designed to support search over document names, identifiers, or keywords from a small or controlled vocabulary often with well-defined semantic meanings. In this paper we propose methods to construct a content-based P2P network overlay over heterogeneous, open-domain contents. By using local algorithms to construct a network overlay with content-based locality and content-based small-world properties, organized content distribution, good network navigability, and balanced connection degrees can be achieved to better support full-text federated search.

Our approach establishes content-based locality at the hub-provider topology level by using an adaptive clustering strategy to suit open-domain contents, and selective content propagation for efficiency and scalability. Instead of pre-determining the coverage of each content area using global semantic information, it adjusts the granularity of existing content areas and creates new content areas based on actual contents and network conditions. Rather than broadcasting to all hubs, the content information of each joining information provider is only propagated to those hubs whose content areas are likely to match the provider's content, so that hubs with dissimilar content areas are shielded from unnecessary overhead.

Our approach also constructs content-based small-world at the hub-hub topology level by adaptively adjusting hubs' local and long-range connections based on limited local views of the network. Long-range connections are determined using a power-law distribution of content



(a) overlap precision, start from non-relevant regions (b) overlap recall, start from non-relevant regions



(c) overlap precision, start from relevant regions (d) overlap recall, start from relevant regions

Figure 7.9: The search performance for GOV queries in the large-sized network.

distance, which leads to evenly distributed long-range hub connections over all “distance scales” for good navigability.

To maximally utilize the resources available in the network, hubs play a more active role in our overlay construction algorithms to minimize the burden on leaves with limited resources. Furthermore, hubs share the responsibility of serving contents and routing queries by balancing their connection degrees, which reduces overload and the impact of a single point of failure.

Compared with previous work on P2P network overlays, our approach is well-matched to distributed environments that i) contain heterogeneous, open-domain contents that may not have semantic annotations, ii) lack centralized control which makes it difficult to acquire global information, iii) involve constant changes in content and connection due to peer arrivals and departures, and iv) consist of peers varying widely in processing power and connection bandwidth.

In addition to algorithms, we also developed a new large-scale P2P testbed that includes 25,000 collections with real documents extracted from the .GOV2 dataset (Lu & Callan, 2006b), and use real queries from AOL search in evaluation. Experimental results using this new testbed and an existing P2P testbed verify the effectiveness of our approach to constructing a content-based network overlay to support full-text federated search.

ACKNOWLEDGEMENTS

This material is based on work supported by NSF grant IIS-0240334. Any opinions, findings, conclusions or recommendations expressed in this material are the authors', and do not necessarily reflect those of the sponsor.

REFERENCES

- Asvanund, A. (2004). Peer-to-peer networks: user behaviors, network effects and protocol extensions. Ph.D. thesis, the Heinz School of Public Policy and Management, Carnegie Mellon University.
- Clauset, A. & Christopher, M. (2004). How do networks become navigable? [oai:arXiv.org:cond-mat/0304563](http://oai.arXiv.org:cond-mat/0304563).
- Crespo, A. & García-Molina, H. (2002). Semantic overlay networks for P2P systems. Technical report, Computer Science Department, Stanford University.
- French, J. Powell, A. Viles, C. Emmitt, T. & Prey, K (1998). Evaluating database selection techniques: A testbed and experiment. In *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 121—129).
- Gnutella v0.4 (2000). http://www9.limewire.com/developer/gnutella_protocol_0.4.pdf.
- Khambatti, M. Ryu, K. & Dasgupta, P. (2002). Efficient discovery of implicitly formed P2P communities. *Int'l Journal of Parallel and Distributed Systems and Networks*, 5(4), 155—164.
- Kleinberg, J. (2000). The small-world phenomenon: an algorithmic perspective. In *Proceedings of 32nd ACM Symposium on Theory of Computing* (pp. 163—170).
- Li, M. Lee, W. & Sivasubramaniam, A. (2004). Semantic small world: An overlay network for peer-to-peer search. In *Proceedings of the 12th IEEE Int'l Conference on Network Protocols* (pp. 228—238).
- Löser, A. Naumann, F. Siberski, W. Nejd, W. & Thaden, U. (2003). Semantic overlay clusters within super-peer networks. In *Proceedings of the International Workshop on Databases, Information Systems and P2P Computing in Conjunction with the VLDB 2003* (pp. 33—47).
- Lu, J. & Callan, J. (2003). Peer-to-peer testbed definitions: trecwt10g-2500-bysource-v1 and trecwt10g-query-bydoc-v1. <http://www.cs.cmu.edu/~callan/Data>.
- Lu, J. & Callan, J. (2006a). Full-text federated search of text-based digital libraries in peer-to-peer networks. *Journal of Information Retrieval*, 9(4), 477—498.
- Lu, J. & Callan, J. (2006b). Peer-to-peer testbed definition: trecgov2-25000-bysource-v1. <http://www.cs.cmu.edu/~callan/Data>.
- Manku, G. Bawa, M. & Raghavan, P. (2003). Symphony: Distributed hashing in a small world. In *Proceedings of the 4th USENIX Symposium on Internet Technologies and Systems* (pp. 127—140).
- Manna, S. & Kabakcioglu, A. (2003). A Scale-free network on Euclidean space optimized by rewiring of links. *Journal of Physics A-Mathematical and General*, 36(19), 279—285.
- Merugu, S. Srinivasan, S. & Zegura, E. (2004). Adding structure to unstructured P2P networks: the use of small-world graph. *Journal of Parallel and Distributed Computing on Theoretical and Algorithmic Aspects of Sensor, Ad Hoc Wireless and P2P Networks*, 65(2), 142—153.
- Sakaryan, G. & Unger, H. (2003). Topology evolution in distributed P2P networks. In *Proceedings of Applied Informatics (AI 2003)* (pp. 378—441).
- Stenmark, D. (2005). Query expansion on a corporate intranet: Using LSI to increase relative precision in explorative search. In *Proceedings of 38th Hawaii International Conference on System Sciences*.
- Schlosser, M. Sintek, M. Decker, S. & Nejd, W. (2002). A scalable and ontology-based P2P infrastructure for semantic Web services. In *Proceedings of the 2nd IEEE International Conference on P2P Computing (P2P2002)* (pp. 104).
- Stoica, I. Morris, R. Karger, D. Kaashoek, M. & Balakrishnan, H. (2001). Chord: A scalable peer-to-peer lookup service for internet applications. In *Proceedings of the SIGCOMM'01 Conference*, (pp. 149—160).
- Watts, D. & Strogatz, S. (1998). Collective dynamics of small-world networks. *Nature*, 393, 440—442.