

GRACE: A Gradient Distance-Based Peer-to-Peer Network Supporting Efficient Content-Based Retrieval

Jianming Lv, Can Yang and Kaidong Liang

Abstract Beyond the simple file name-based search supported in peer-to-peer (P2P) file sharing networks, content-based search aims to precisely locate the files containing the desired contents. Several existing P2P systems based on structured overlays provide content-based search by indexing shared contents to numeric keys. However, these systems are usually costly when publishing and maintaining the index of contents in dynamic network churn. In this paper, we propose a novel P2P network, GRACE, which probably constructs connections among peers according to the *gradient distance* of their shared contents. GRACE can achieve high efficiency of the content-based search while gaining a significantly lower maintenance cost than those of previous efforts.

Keywords Peer-to-peer · Overlay network · Gradient distance · Content-based search

1 Introduction

Content-based search is one of the most important functions provided in peer-to-peer (P2P) content management systems and aims to precisely locate the files containing the desired contents. While P2P systems are usually large-scale,

J. Lv (✉) · C. Yang · K. Liang
South China University of Technology, Guangzhou 510006, China
e-mail: jmlv@scut.edu.cn

C. Yang
e-mail: cscyang@scut.edu.cn

K. Liang
e-mail: liangkaidong@gmail.com

heterogeneous, highly distributed, and dynamically changed, it is challenging to perform efficient and robust content-based search.

Recently, some content-based search algorithms [1–3] over structured P2P overlays are proposed by mapping the shared contents to low-dimensional keys. The mapping operations adopted in above systems inflict high cost to publish and maintain the key-based indexes, especially in the dynamic networks where peers join and depart frequently.

In this paper, we present a probably constructed P2P network, GRACE, to support efficient, resilient content-based search with low maintenance cost. In the GRACE network, peers build the neighboring connections probably according to the *gradient distance* between their shared documents. Compared with traditional P2P content-based search algorithms, GRACE can support efficient content-based search of text documents with relatively low maintenance cost. GRACE needs no overhead of publishing and maintenance of key-based indexes, which is necessary in traditional search algorithms [1–3] based on structured overlays. The maintenance cost of GRACE is less than 1 % of those.

The rest of this paper is organized as follows. Section 2 introduces related works. Section 3 presents the model of the GRACE network. The experimental results are discussed in Sect. 4. Finally, Sect. 5 summarizes the paper.

2 Related Work

Structured overlay networks [4–6] have a kind of distributed infrastructures to support $O(\log N)$ key-based search, which are usually named as the distributed hash table (DHT). The topology of the network is constructed based on the global unique ID assigned to each peer. Given any numeric key, the structured overlay can locate the peer having the closest ID to the key within $O(\log N)$ hops.

Several recent research efforts [1–3, 7] build content-based search algorithms over DHT. Specifically, the researches [1, 2] map each keyword of a shared text document into a unique numeric key by a uniform hash function to support P2P full-text search. Muller et al. [3] propose a content-based image retrieval algorithm by mapping the color histogram vector of each image into a fixed set of numeric keys. Tang et al. [7] map the high-dimensional document vectors into low-dimensional numeric keys by using the latent semantic index (LSI) method. All of them need to publish index items constructed as $\langle key, element \rangle$ in the networks to support efficient search, which can cause the following latent problems:

- High communication cost is required to maintain the consistence of the index in the dynamic network.
- Uneven distribution of the shared data elements can cause load imbalance in these systems. In full-text search systems [1, 2], the Zipf-distributed keyword popularity can make the load of the peers maintaining the index items of hot keywords extremely heavy.

3 The GRACE Network

3.1 Basic Model

The GRACE model builds up connection probably between any pair of peers according to the *gradient distance* between their shared documents. The *gradient distance* is a kind of document distance with the following *gradient property*: Given any document v in a document dataset, the number of remaining documents within x distance of v is proportional to x^β , $x \in (1, +\infty)$, where β is a constant bigger than zero.

In the GRACE model, the *gradient distance* between any pair of documents is defined as follows:

$$d(v_i, v_j) = 1 / \left(\frac{v_i \cdot v_j}{\|v_i\| \|v_j\|} + \theta \right) \quad (0 < \theta < 1) \quad (1)$$

Here, θ is a constant to prevent the denominator to be zero. v_i and v_j are the content vectors of the two documents. For a text document, the content vector is defined as the term vector [8], whose elements indicate the frequencies of the terms appearing in the document.

Based on the gradient distance defined above, any peer P_i in the GRACE network selects another peer P_j as its neighbor with the following probability:

$$l(P_i, P_j) = \alpha |H(P_i)|^{-1} \sum_{v_m \in H(P_i)} \sum_{v_n \in H(P_j)} d(v_m, v_n)^{-\beta} f(v_n)^{-1} \quad (2)$$

where $H(P_i)$ and $H(P_j)$ denote the document sets shared by P_i and P_j . $d(v_m, v_n)$ is the gradient distance between any two documents v_m and v_n shared by P_i and P_j , respectively. $f(v_n)$ means the ratio of peers in the network sharing the document v_n , and the parameter β is the constant defined in the gradient property of the distance. $\alpha \in (0, +\infty)$ is a constant used to normalize the probability.

3.2 Construction of the GRACE Network

The main task of constructing a GRACE network is to build up the neighborhood of online peers. A simple solution is to select the neighbors of each peer according to the connection probability with all existing peers in the network calculated as Eq. (2). However, this global selection strategy is impractical for its huge communication cost. A more efficient two-phase selection mechanism is adopted in GRACE as follows:

Phase I: Each peer selects several peers randomly from the network to form its candidate set.

Phase II: The peer selects peers from the candidate set as its neighbors. The probability for a candidate to be selected is calculated according to Eq. (2).

To implement the above phases in a distributed manner, each peer maintains four tables: *Host Table*, *Random Seed Table*, *Candidate Table*, and *Neighbor Table*.

The Host Table is used to store the contact information of the hosts appearing in the peer's interaction history. For a peer joining the GRACE network for the first time, the Host Table is initialized to contain some stable online peers like what Emule [9] do.

The Random Seed Table consists of peers randomly selected from the Host Table and used for candidate selection in Phase I. Each host in the table is named as a *random seed*. The Candidate Table records the selected candidate peers.

The procedure to construct the Neighbor Table is illustrated in Algorithm 1. Step 1 and step 2 of Algorithm 1 are executed to construct the Random Seed Table from the Host Table. Steps 3–9 illustrate how to construct the Candidate Table by interacting with the random seeds. The candidates selected here can be viewed as a random sample of all peers in the network. Step 10 and step 11 are to select the GRACE neighbors from the candidates according to the probability defined in Eq. (2). The document collections $H(P_i)$ transferred in step 4 and step 5 are presented as content vectors and can be compressed to a small size in practical systems.

Algorithm 1: The joining procedure of a peer

Input: P_0 : The peer to join the network.

Output: The Neighbor Table of P_0 .

Notations:

Hosts (P_i): The Host Table of the peer P_i .

Random_Seeds (P_i): The Random Seed Table of P_i .

Candidates (P_i): The Candidate Table of P_i .

Neighbors (P_i): The Neighbor Table of P_i .

Method:

//construct P_0 's Random Seed Table.

1. **for** $\forall(P_i) \in \text{Hosts}(P_0)$ **do**

2. P_0 inserts P_i to *Random_Seeds* (P_0) with probability c .

//construct P_0 's Candidate Table.

3. **for** $\forall(P_i) \in \text{Random_Seeds}(P_0)$ **do**

//exchange the shared document vectors

4. P_0 sends $H(P_0)$ to P_i

5. P_i sends $H(P_i)$ to P_0

//select all Random Seeds of P_i as candidates

6. **for** $\forall(P_j) \in \text{Random_Seeds}(P_i)$ **do**

7. $Pr_j \leftarrow l(P_0, P_j)$ //Eq. (2)

8. P_i send (P_j, Pr_j) to P_0

9. P_0 insert (P_j, Pr_j) into *candidates*(P_0)

//select the GRACE neighbors form candidates

10. **for** $(P_j, Pr_j) \in \text{candidates}(P_0)$ **do**

11. P_0 insert P_j to *Neighbors* (P_0) with probability Pr_j

3.3 Content-Based Search in GRACE

A simple greedy algorithm is adopted to perform content-based search in GRACE. While a peer forwards a query in the network, it selects another peer from its Neighbor Table as the next hop, which shares the document with minimal distance to the query. Each searched peer independently checks its local stored documents and returns back the records of the top k documents closest to the query. Then, the initiator of the query gathers all returned records, ranks the documents according to the distance to the query, and composes the search result as the top k documents.

4 Experiments

4.1 Data Collection in Experiments

We evaluate the experimental systems on the dataset TREC WT10G [11], which contains 1,692,096 Web documents from 11,680 Web sites. We deem each Web site to be a peer and associate the documents on a site to the corresponding peer as shared contents. We also construct the query set containing 100 queries by randomly sampling as [12]. Each query is obtained by first randomly choosing a document from the dataset and then randomly choosing 5 terms from the document as the searched keywords of the query. The search task is to search the top k documents most related to the query.

4.2 Comparison with Other Overlays

We compare GRACE with two algorithms described as follows:

- *K-Chord* [1, 2]. K-Chord is based on Chord [5] and supports keyword-based search.
- *Rand*. Connections are randomly built among peers [10] and random-walker mechanism is adopted to route search requests.

4.3 Search Latency Analysis

To understand how the performance of GRACE is affected by the network scale, we randomly select peers from the whole datasets to form the subsets with different scales. Then, we build the P2P networks over each subset and test the search efficiency.

Fig. 1 Search latency in different networks

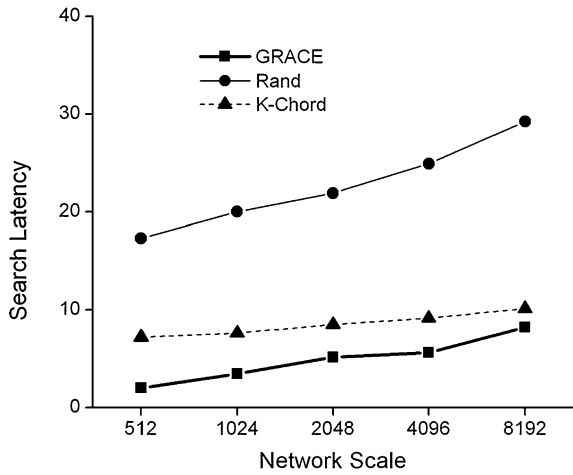


Fig. 2 Communication cost of network maintenance. Here, *Join* means the cost for a new peer to join the overlay. *PTM* means periodical topology maintenance cost. *PIM* means periodical index maintenance cost. *Leave* means the cost for a peer to leave the overlay

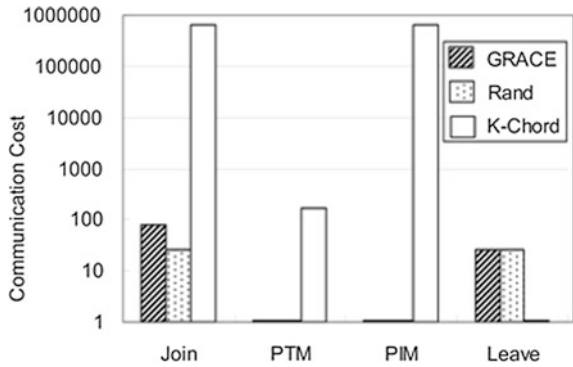


Figure 1 shows the average search latency of GRACE, Rand, and K-Chord under different network scales in the WT10G dataset. We can see that GRACE is more efficient than K-Chord and Rand to locate the top k documents most similar to the query.

4.4 Communication Cost

Figure 2 illustrates the communication cost of different systems run on the WT10G dataset. The y-axis of the figure is the communication cost measured by the number of messages. We can see that the joining cost of K-Chord is much higher than that of GRACE and Rand. Moreover, in K-Chord, the cost for joining and indexing maintenance makes up the largest percentage. In fact, this overhead is required in most of structured overlay-based systems for publishing and

maintaining data indexes. Compared with these systems, the advantage of GRACE is to support efficient search with no need of index publishing, so its maintenance cost can be less than 1 ‰ of K-Chord.

5 Conclusions

The GRACE presented in this paper probably constructs connections among peers according to the gradient distance of their shared contents. GRACE provides efficient content-based search according to the similarity of shared contents without any additional key-based index. Experiments validate the search performance of GRACE and show that its maintenance cost is less than 1 ‰ of structured overlay-based search systems.

Acknowledgments The work described in this paper was supported by the grants from National Natural Science Foundation of China (No. 61300221), the Comprehensive Strategic Cooperation Project of Guangdong Province and Chinese Academy of Sciences (No. 2012B090400016), the Technology Planning Project of Guangdong Province (No. 2012A011100005), and the Fundamental Research Funds for the Central Universities (Project No. 2011ZM0069).

References

1. Reynolds, P., Vahdat, A.: Efficient peer-to-peer keyword searching. In: Proceedings of 4th ACM/IFIP/USENIX International Middleware Conference, pp. 21–40. (2003)
2. Yang, Y., Dunlap, R., Rexroad, M., Cooper, B.F.: Performance of full text search in structured and unstructured peer-to-peer systems. In: Proceedings of 25th IEEE International Conference on Computer Communication, pp. 1–12. (2006)
3. Muller, W., Boykin, P.O., Sarshar, N., Roychowdhury, V.P.: Comparison of image similarity queries in P2P systems. In: Proceedings of 6th IEEE Conference on Peer-to-Peer Computing, pp. 98–105. (2006)
4. Ratnasamy, S., Francis, P., Handley, M., Karp, R., Shenker, S.: A scalable content-addressable network. In: Proceedings of ACM SIGCOMM, pp. 161–172. (2001)
5. Stoica, I., Morris, R., Karger, D., Kaashoek, F., Balakrishnan, H.: Chord: a scalable peer-to-peer lookup service for internet applications. In: Proceedings of ACM SIGCOMM, pp. 149–160. (2001)
6. Zhao, B.Y., Huang, L., Stribling, J., Rhea, S.C., Joseph, A.D., Kubiawicz, J.D.: Tapestry: a resilient global-scale overlay for service deployment. *IEEE J. Sel. Areas Commun.* **22**(1), 41–53 (2004). doi:[10.1109/JSAC.2003.818784](https://doi.org/10.1109/JSAC.2003.818784)
7. Tang, C., Xu, Z., Dwarkadas, S.: Peer-to-peer information retrieval using self-organizing semantic overlay networks. In: Proceedings of ACM SIGCOMM, pp. 175–186. (2003)
8. Baeza-Yates, R., Ribeiro-Neto, B.: *Modern Information Retrieval*, pp. 41–44. ACM Press, New York (1999)
9. Emule.: Official emule site. <http://www.emule.org/>

10. Lv, C., Cao, P., Cohen, E., Li, K., Shenker, S.: Search and replication in unstructured peer-to-peer networks. In: Proceedings of 16th ACM International Conference on Supercomputing, pp. 84–95. (2002)
11. TREC.: TREC web dataset: WT10g. http://ir.dcs.gla.ac.uk/test_collections/wt10g.html. (2003)
12. Bawa, M., Manku G.S., Raghavan, P.: SETS: search enhanced by topic-segmentation. In: Proceedings of ACM SIGIR, pp. 306–313. (2003)