Optimization of Content Placement Scheme for Social Media on Distributed Content Clouds

Qian Zhang[1,2]. Runzhi Li[2]. Yusong Lin[2]. Zongmin Wang[2]

Abstract. Social networking applications have been more and more popular. The huge demands of bandwidth and storage for social media have brought great challenges to the network engineering. The recently emerged content clouds shed light on this dilemma. Towards the content placement scheme on clouds, partitioning the social contents on cloud servers has draw significant interests from the literatures. The existing studies demonstrate that YouTube videos exhibit strong geographic locality of interest. Yet the previous works on partitioning social contents on cloud servers overlooked the geographic popularity of videos. In this paper, we present a geographic-aware content placement scheme on distributed content clouds which takes the geographic popularity of videos into account .Aiming at reduce the cross-boundary traffic and realize the load-balance on cloud servers as well as preserve social relationship. Simulations with data traces from YouTube demonstrate the efficiency of the proposed algorithm.

Keywords: Social media. Content placement scheme. Clouds. Geographic popularity .

1 Introduction

In recent years, social networked applications and services have been dominating the Web 2.0 world. The most popular social media include YouTube for video sharing, Facebook or RenRen for online social networking, and Twitter for microblogging. Besides these representatives, many other user-generated content (UGC) applications have emerged and been developing extremely fast. The provision of

[1,2]Q.Zhang (✉),[2] R.Li, [2]Y.Lin, [2]Z.Wang

[1] School of Information Engineering, Zhengzhou University, Zhengzhou

[2] Henan Provincial Key Lab on Information Networking, Zhengzhou University, Zhengzhou, PRC. e-mail: qzhang@zzu.edu.cn

resource is a great challenge for deploying these applications and social media faces a much greater challenge because of the huge resource demands of bandwidth and storage. However, the recently emerged cloud service sheds light on this dilemma. The advent of cloud computing systems like Amazon EC2 [1] has streamlined horizontal scaling by providing the ability to lease virtual machines(VMs) dynamically from the cloud. A content placement scheme for social media on distributed content clouds is essential for the benefit of the social media application's development and competition.

Towards the content placement scheme of social media on clouds, one of the important steps is to partition the social media contents and assign them into a number of cloud servers. There have been some works on solving this problem, e.g., SNAP [2] and SPAR [3] which both are effective tools to partition the social network and aiming at preserving the social relationship. Considering that the unbalanced load on cloud servers when the partitioning algorithm only takes the social relationship into account, Xu [4] proposed a Weighted Partitioning Around Medoids(wPAM) algorithm to partition the social networked video repository and focused on load-balance on cloud servers in terms of user access.

However, the recent studies on YouTube show that the videos exhibit strong geographic locality of interest and the significant characteristic is largely overlooked by the previous partitioning schemes. In this paper, we propose a geographic-aware content placement scheme which takes the geographic popularity of videos into account. The problem can be formulated as a constrained *k*-medoids clustering problem which under the constraint of minimizing the cross-boundary traffic and the imbalanced weight on cloud servers. Compared with the previous scheme which overlooked the geographic popularity of interest, the proposed one effectively reduce the cross-boundary traffic and realize the load-balance on cloud servers as well as preserve social relationship.

The reminder of this paper is organized as follows: Section 2 reviews the related work. The motivation our work is described in Section 3. The problem statement and our proposed algorithm are given in Section 4 and 5. The efficiency of the proposed algorithm is examined in Section 6. Finally, Section 7 concludes the paper.

2 Related Work

The recent years have witnessed an explosion of social media streaming as a new killer Internet application. YouTube, the most representative online social media, has draw significant interests from literature. There have been numerous studies on YouTube-like social media systems and most of them focus on the social structure, user behavior and network usage. Towards the migration scheme of social media to clouds, [5] proposed a cluster-based social media system with cloud-assisted and realized the capacity migration to clouds, which is primarily aimed at minimizing the lease cost from cloud servers. The work in [4] realized the content migration of social media to clouds, in which a Weighted Partitioning Around Medoids(wPAM) algorithm is proposed to partition the social networked video repository and focused on load-balance on cloud servers in terms of user access. SPAR [3] also considers the cloud scenario, which is a social partitioning and replication middle-ware that replicates linked nodes in the same server and tries to minimize the replications. SNAP [2] is a tool for analyzing and partitioning small-world network, and it introduces a series of algorithms trying to maximize the modularity of the graph in parallel manner.

Different from these works, our work is motivated by the latest study which demonstrates that YouTube videos exhibit strong geographic locality of interest [6]. Therefore, in the proposed content placement scheme of social media on distributed content clouds, we consider the geographic popularity of videos which is an important factor that is largely overlooked in these previous works.

3 Motivation

Anders Brodersen et.al [6] studied the geographic popularity of YouTube videos and demonstrated that YouTube videos enjoy a strong geographic locality of interest. The study in [6] showed that the videos tend to become popular in locally confined area, rather than in a globally wide region and there are about 40% of YouTube videos that enjoy at least 80% of their views in a single region. It also demonstrated that the view focus, which is the highest fraction of views that a video has received in a single region over its entire lifetime, decreases linearly as the order of magnitude of the number of views grows. The videos with more than 1,000 views exhibit a steady value of view focus. In other words, the regional popularity of videos can be inferred by its total view number. As the strong geographic locality of interest of videos which is described above, a content placement scheme for social media overlooked this characteristic would lead to unbalanced placement. In particularly, considering the cross-boundary traffic, it is essential to partition the social media content in a geographic-aware way.

The view focus F_i is the highest fraction of views that video i has received in a single region over its entire lifetime, which is demonstrated as $F_i = 1/w_i \times \max w_{ik}$. w_{ik} is the views that video i has received in region k. From [6], when the video with more than 1000 views, the view focus exhibits a steady value of about 0.4. When the video with less than 1000 views, the view focus can be deduced by the total views, $Fi=0.8-0.1\times lgw_i$.

4 Problem statement

Consider a social graph with N nodes $n_1, n_2, \dots, n_N$, each node corresponds to a video in social media system and n_i has a weight w_i which is the total number of view for the corresponding video. There are k cloud servers $k_1, k_2, \dots, k_K$ which are located in k different areas. The popularity in area $j(1 \leqslant j \leqslant K)$ of video j is w_{ij}. Suppose the total view number of video i is w_i.

We try to partition the nodes into k clusters(cloud servers in k different areas), $C_1, C_2, \dots, C_K$. Each cluster C_j has a weight W_j, which is the summation of popularity of nodes in cluster C_j and it can be calculated as $W_j = \sum_{n_i \in C_j} w_{ij}$

The local request ratio R_j of each cloud server j is corresponding to the local request ratio in each cluster C_j and it is defined as $R_j = \sum_{n_i \in C_j} w_{ij} / \sum_{n_i \in C_j} w_i$

We suppose there is a representative node o_j in each cluster C_j. $d(n_s, n_t)$ is the distance metric between the two node n_s and n_t and is also the similarity between n_s and n_t. There are several metrics for measuring social networks are used, such as betweenness [9], condutance [8], modularity [2], and number of replicas [3]. We use the similarity as described in [4]. The geographic-aware partitioning problem is to find a partition which minimizes $E = \sum_{j=1}^{k} \sum_{n_i \in C_j} d(n_i, o_j)$ under the constraint of

$$\left| W_{j1} - W_{j2} \right| < \Delta_1,\ j_1, j_2 = 1, \dots, k \quad (1)$$

$$R_j > \Delta_2,\ j = 1, \dots, k \quad (2)$$

The first constraint is to realize the load balance among the cloud servers by controlling the weight difference in Δ_1.To mitigate the cross-boundary traffic, the second constraint is required to constrain the local request ratio in each cluster j should be larger than Δ_2 .

From the problem described above, the geographic-aware partitioning problem can be considered as a k-medodis clustering problem under the constraint of minimizing the cross-boundary traffic and the imbalanced weight on cloud servers.

5 Solution

To solve the problem, a geographic-aware partitioning algorithm is proposed in this section. The k-medoids clustering problem has already been proven to be NPhard, and a variety of heuristic algorithms have been proposed. As given in [7], PAM is a typical algorithm for *k*-medoids clustering problem. We develop our algorithm based on PAM and the following is the pseudo-code for the proposed geographic-aware partitioning algorithm.

Algorithm 1 Geographic-aware partitioning algorithm

INPUT: N nodes $n_1,..n_N$ with weight $w_1,\ldots,w_N,$ cluster number K, weight constraint threshold Δ_1, local request ratio threshold Δ_2.

Method: calculate the geographic popularity of each node I in different K areas, which is $w_{i1,\ldots}w_{iK;}$ arbitrarily choose K nodes as the initial representative nodes; for each representative node o_j;

repeat assign each remaining node n_j with weight w_j, to the first nearest cluster i with weight W_i, if satisfying $W_i + w_{ji} < \overline{W} + \Delta_1$ and $R_i > \Delta_2$;

for each non-representative node o_r; compute the total cost, S= E_r –E, of swapping representative node, o_j, with o_r; Choose the lowest cost S, if S<0 then swap o_j with o_r which has the lowest cost to form the new set of K representative nodes;

until no change;

OUTPUT: a set of K clusters satisfying the weight and local request ratio constraint.

The distinct feature of the proposed geographic-aware algorithm is that the local request ratio is considered and the video's geographic popularity can be calculated by the total popularity which is described as in Sec. 3.1. The node is assigned to the nearest cluster to maintain the social relationship. Considering the load-balance on the cloud servers, the weight of the cluster adding on the weight of the new node in this cluster should not be larger than $\overline{W}+\Delta_1$, where $\overline{W}$ is the average weight of all clusters and $\sum_{n_i \in C_j} d(n_i, o)$ is the weight difference constraint. To mitigate the cross-boundary traffic, the local request ratio of the cluster when a new node adds in should be larger than Δ_2. E_r is the distance between the non-representative node and other nodes in the cluster, which is $\sum_{n_i \in C_j} d(n_i, o_r)$. E is the distance between the original representative node and other nodes in the cluster, which is $\sum_{n_i \in C_j} d(n_i, o)$. Choose the lowest cost S, if S is negative, then o_j is replaced or swapped with o_r since actual E would be reduced; if $E_r - E$ is positive, the current representative node o_j is acceptable and nothing is changed in the iteration.

PAM does not scale well for large dataset, CLARA(Clustering LARge Applications) [7] was introduced to deal with larger dataset. Instead of finding representative objects for the entire data set, CLARA draws a sample of the data set, applies

PAM on the sample and finds the medoids of the sample. Experiments reported in [7] indicate that 5 samples of size 40+2k give satisfactory results. We utilize CLARA based on the proposed geographic-aware partitioning algorithm without any further modification to deal with larger dataset.

6 Evaluation

In this section, we evaluate the efficiency of the proposed geographic-partition algorithm. There are two evaluation criteria: 1) The local request ratio in each area, which can be calculated as the local request divided by the total request in one area. 2) The percentage of transitions. A transition is defined as two consecutive videos are held by two different servers in different areas, and the less the number of transitions is, the better the social relationship is preserved. We compare our proposed algorithm with WPAM which is only considering the load-balance on cloud servers and test efficiency of our proposed algorithm when the weight constraint and local request constraint changes.

6.1 Evaluation parameters

The videos' popularity and the social relationship between the videos are based on the crawled data from 3000 YouTube videos. In our evaluation, we divided them into 10 distributed cloud servers in different 10 areas. The view focus of each video is calculated as described in Sec.3. The fraction of videos' focus in different 10 areas is shown as Table 1. We apply the proposed geographic-aware content placement scheme to place these videos into 10 different on cloud servers in different 10 areas. When the video's popularity larger than 1000, the view focus of this video is almost steady at 0.4. Therefore the local request ratio of each cluster is around 0.4 is acceptable. Unless stated, the local request constraint Δ_2 is supposed to 0.34 and the weight constraint Δ_1 is supposed to $0.01\times\overline{W}$.

Area	area1	area 2	area 3	area4	area5	area6	area7	area8	area9	area10
Videos	37.7%	6.6%	4.9%	4.1%	3.2%	2.9%	2.5%	2.5%	2.5%	2.4%

Table 1 : The fraction of videos' focus in different 10 areas

6.2 Results

We compare our proposed scheme to WPAM which is not considering the geographic distributed cloud servers. Fig.1 shows the local request ratio in each cluster based on WPAM and our proposed algorithm. From Fig.1, the local request ratio in each cluster which is partitioned based on PAM is around 0.1 and the local request ratio in our proposed algorithm is around 0.34. Then we test the efficiency of the proposed algorithm when the weight constraint and local request ratio constraint changes. Fig.2 shows the percentage of transitions against different constraint. From Fig.2, we can see that when we tight the constraint the percentage of transition will be slightly increase which also means the social relationship is partly broken. Fig. 2 also illustrates that the impact of the constraint on the percentage of transition is slight.

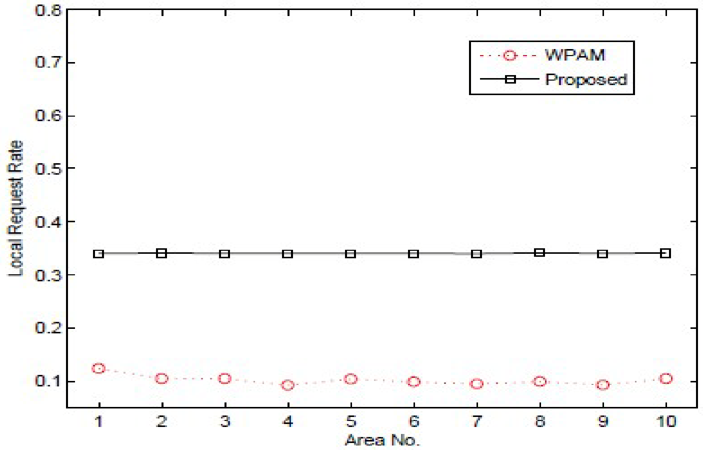

Fig.1 Compare of local request ratio

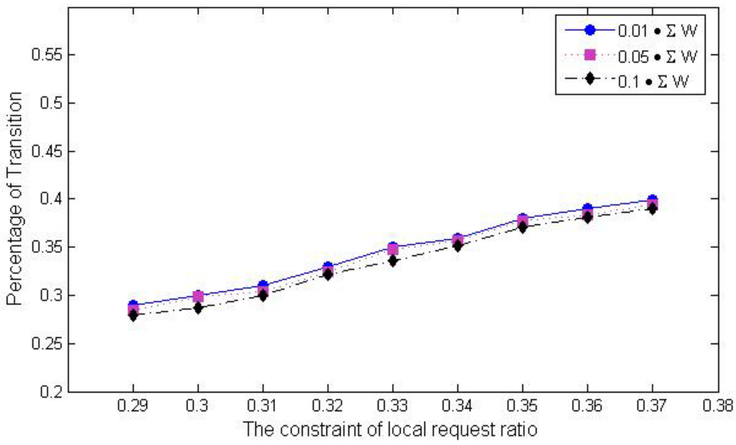

Fig.2 Comparison of percentage of transitions

7 Conclusion

In this paper, we take the geographic distributed cloud servers and the geographic popularity of YouTube-like videos into account in the problem of content placement of social media in the cloud scenario. The recent studies demonstrated that YouTube videos enjoy a strong geographic locality of interest and we found that this characteristic should not be overlooked especially in the geographic distributed cloud scenario. Then we proposed a geographic-aware content placement scheme which takes the geographic popularity of videos into account. Aiming at reduce the cross-boundary traffic and realize the load-balance on cloud servers as well as preserve social relationship. Simulations with data traces from YouTube demonstrated the efficiency of the proposed algorithm.

8 Acknowledgment

This paper is sponsored by Key Program on Science & Technology Research OF He Nan Province under Grant No. 13A520562 and Innovation Project of He Nan Province under Grant No. 2011HASTIT003 with Zhengzhou University.

9 References

[1] Amazon Elastic Compute Cloud, http://aws.amazon.com/ec2

[2] D. A. Bader and K. Madduri(2008). SNAP, Small-world Network Analysis and Partitioning: An open-source parallel graph framework for the exploration of large-scale networks. In Proc. of IPDPS.

[3] J. M. Pujol, V. Erramilli, G. Siganos, X. Yang, N. Laoutaris, P. Chhabra, and P. Rodriguez(2010). The Little Engines(s) That Cloud: Scaling Online Social Networks. In Proc. of SIGCOMM.

[4] X. Cheng and J.C.Liu(2011). Load-Balanced Migration of Social Media to Content Clouds. In Proc. of NOSSDAV.

[5] Q. Zhang, Y.S.Lin and Z.M.Wang(2013). Cost-effective capacity migration of Peer-to-Peer social media to clouds. Peer-to-Peer Networking and Applications:Volume 6, Issue 3 , pp 247-256

[6] A. Brodersen, S. Scellato and M. Wattenhofer(2012). YouTube Around the World: Geographic Popularity of Videos. In Proc. of WWW

[7] L. Kaufman and P. J. Rousseeuw(1990). Finding Groups in Data: An Introduction to Cluster Analysis.Wiley.

[8] N. Mishra, R. Schreiber, I. Stanton, and R. Tarjan(2007). Clustering Social Networks. In Algorithms and Models for the Web-Graph, volume 4863 of Lecture Notes in Computer Science, Pages 56-67. Springer Berlin/Heidelber

[9] M. E. J. Newman and M. Grivan(2004). Finding and Evaluating Community Structure and Networks. Physical Review E, 69(2):026113