Journal of Computational Science xxx (2017) xxx–xxx

Contents lists available at ScienceDirect

Journal of Computational Science

journal homepage: www.elsevier.com/locate/jocs

Optimizing cost for geo-distributed storage systems in online social networks

Jingya Zhou[a,b,*], Jianxi Fan[a], Juncheng Jia[a], Baolei Cheng[a], Zhao Liu[a]

[a] School of Computer Science and Technology, Soochow University, Suzhou 215006, China
[b] Collaborative Innovation Center of Novel Software Technology and Industrialization, Nanjing 210093, China

ARTICLE INFO

Article history:
Received 30 November 2016
Received in revised form 2 July 2017
Accepted 2 August 2017
Available online xxx

Keywords:
Online social networks
Cost optimization
Inter-data center traffic
Storage
Geo-distributed storage systems

ABSTRACT

Globally distributed data centers provide an opportunity to deploy geo-distributed Online Social Networks (OSNs). For so big data generated by users, how to store them among those data centers is a key issue in the geo-distributed storage system. Today's popular OSN providers store users' data in each deployed data center, so as to guarantee access latency. However, the full replication manner brings relatively high storage cost and traffic cost, which extremely increases the economic expenditure of OSN providers. Data placement based on social graph partitioning is an efficient way to minimize cost, but it requires the information of entire social graph and cannot fully guarantee latency. Recently, accomplished by partitioning replication is proposed to optimize cost as well as guarantee latency, but it has two drawbacks: (1) the separated manners of optimization cannot efficiently reduce the cost; (2) the placement of master replicas and slave replicas influence each other, and eventually reduces the optimization effects. In this paper, we explore an integrated manner of optimizing partitioning and replication simultaneously without distinguishing replica's role. We propose a lightweight replica placement (LRP) scheme, which conducts optimizations in a distributed manner and is well adapted to dynamic scenarios. Evaluations with two datasets from Twitter and Facebook show that LRP significantly reduces the cost compared with state-of-the-art schemes.

1. Introduction

Nowadays, Online Social Networks (OSNs) are undoubtedly one of the most popular Internet services. Billions of people connect, communicate and share content via their social accounts as OSNs make communication more convenient and efficient than ever before. For example, Facebook and Twitter had 1.59 billion [1] and 332 million [2] monthly active users as of January 2016, respectively. Those users are often globally distributed and they come from various continents and nations with different ages, races, etc. According to a recent report about the statistics of user locations collected from two famous OSNs, Facebook and Twitter [3], we draw a pair of pie charts as shown in Fig. 1, where the top five countries for Facebook users are the United States (14%), India (9%), Brazil (7%), Indonesia (5%) and Mexico (4%), and more than 60% of Facebook users come from the other countries. Similar to Facebook, people who are located in other countries outside of the United States make up 77% of Twitter users. As a matter of fact, most of users are located in many other countries all over the world rather than a single country such as the United States.

For OSN service providers, it really matters to the user experienced performance how to efficiently store massive amounts of users' data [4,5]. If users' data are stored in the United States, the access latency experienced by users from distant countries cannot be well guaranteed. For example, Facebook plans to build its sixth data center in Ireland. Combined with the data center in Sweden, Facebook will have two data centers in Europe, which could significantly enhance its service levels for European users. The cloud computing paradigm offers more choices for service providers to deploy their services [6]. Different from the conventional approaches that use dedicated data center owned by the service providers, geo-distributed clouds provide a much more economic solution [7]. DLS [8] presents a highly efficient caching and prefetching mechanism tailored to reduce metadata access latency. But DLS does not exploit the economic strengths of cloud. Facing such big OSN data [9,10], in order to realize the economic potentials of cloud computing, a practical problem need to be addressed: How to store users' data across those geo-distributed data centers, such that the cost could be optimized while the access latency is

---

[*] Corresponding author at: School of Computer Science and Technology, Soochow University, Suzhou 215006, China.
E-mail address: jy_zhou@suda.edu.cn (J. Zhou).

http://dx.doi.org/10.1016/j.jocs.2017.08.001
1877-7503/© 2017 Elsevier B.V. All rights reserved.

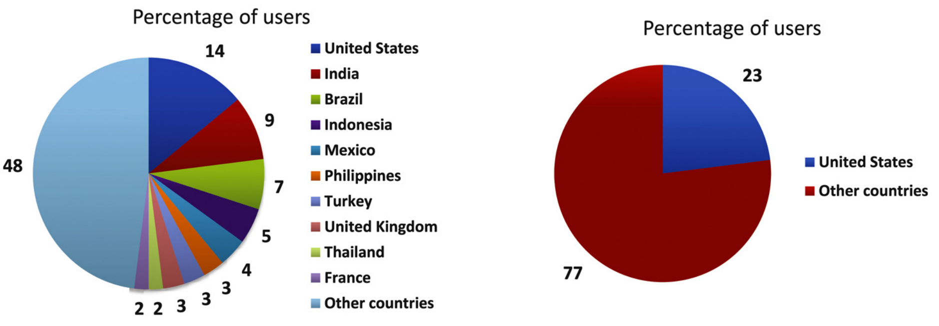

(a) Facebook

(b) Twitter

Fig. 1. The statistics of OSNs user locations (http://sproutsocial.com/insights/new-social-media-demographics/).

guaranteed? An intuitive approach requires to store the requested data closer to the users. Currently, many popular OSN providers store users' data in multiple data centers in a full replication manner, and it requires that each data center has complete replicas of all users' data [11]. By considering cost, the following two issues arise:

- High storage cost, which results from the full replication of increasing number of users' data. It is reported that Facebook has a 14% increase of users every year. Let us assume that the average size of each user's data is 1 MB, the increased data volume is more than 2 PB, and the volume will grow continuously every year.
- High communication cost. Single-master multi-slave paradigm [12] as a common replication model is widely used to maintain data consistency. In this model, a slave data center forwards an update to the master data center, which then pushes the update to all data centers. During frequent interactions across users, updates will yield high communication cost due to high inter-data center traffic.

It is challenging to realize an efficient data placement that is able to minimize cost as well as guarantee access latency. The best solution is that instead of storing all users' data replicas each data center stores data of users who are geographically proximate to it [13]. However, different from other Internet services, OSNs have to deal with highly interactive operations [14], such as browsing some friends' profile pages or posting comments, which requires fetching the data of friends. Several works have explored content replication via a Content Delivery Network (CDN) or cache network [15,16]. Scellato et al. [15] exploit geo-social properties of OSN users to improve the efficiency of caching in different CDN sites, and then propose a location-aware cache replacement policy to ensure that the relevant content replica is placed close to the users who may access to it. The contributions of [15] are limited in social cascades rather than general interactions, and more importantly it does not consider data migration across different sites.

It would be better to have friends' data co-located, since a user does not have to go to more than one data center for fetching her friends' data, and save the additional latency and inter-data center traffic. Social locality is often used to evaluate how extent the data of both a user and her friends are co-located in the same data center. A substantial body of work tries to optimize OSN data placement via maximally preserving social locality [12,17–19]. S-PUT [17] implements a socially-aware data placement approach based on graph partitioning technology. The original problem is modeled as a min-cut social graph partitioning problem that requires each user and her friends being co-located in one partition, so as to minimize cut weight, which is equivalent to keep social locality. Since the social graph is connected, perfect social locality can never be achieved by partitioning. SPAR [18] co-locates the data of users' every friend in the same site by means of replication technology. Although SPAR can preserve social locality perfectly, excessive replication leads to more inter-data center synchronization traffic. Liu et al. [12] propose to select user data with high read rate and low write rate to create replicas, so that social locality is improved without increasing inter-data traffic. Optimizing social locality solely may not achieve the desired objectives. Jiao et al. [20] exploit geographical locality by greedy algorithms to resolve iteratively master and slave replica placement. However, most of existing studies do not realize the optimization of social locality and geographical locality. Furthermore, in practice, different data centers may charge varied prices for data storage and OSNs usually are of dynamic, which makes problem more challenging.

In this paper, we explore the cost-optimizing placement of massive amounts of geo-dispersed user data across multiple data centers. For OSN providers, we carefully model the cost-optimizing data placement problem, and propose a Lightweight Replica Placement (LRP) scheme, which optimizes the choice of data center for data storage. The detailed contributions are summarized as follows:

- A new modeling framework for cost estimation in the storage system is proposed for online social networks, and the model identifies four types of cost including storage cost, inter-data center traffic cost, user-data center traffic cost and migration cost. We analyze how placement operation affects the cost, and formulate a cost-minimizing problem of data placement.
- We propose an efficient online scheme LRP to resolve the problem. The motivation of LRP is to optimize social locality as well as geographical locality. Through optimizing social locality we aim to minimize the cost incurred by storage and communications, while through optimizing geographical locality we could provide users with latency-guaranteed services. More importantly, we combine data placement and data replication together. To realize optimization together in an integrated manner rather than the separated manner, we place data replicas based on the principles that the latency is guaranteed and the cost can be reduced. In specific, we do not distinguish master from slave, which is able to avoid the influence between master replication and slave replication.
- We analyze the dynamic scenario in OSNs, and conclude four cases including the addition of both new user and new relation, the removal of existing relation, and changes of request rates. A dynamic algorithm with reference to these cases is presented to handle OSN dynamics in practical environments.
- Performance of LRP is evaluated via groups of experiments upon real-world datasets. We extensively compare the performance of our scheme to other schemes such as SPAR [18], SD^3 [12], cosplay [21], etc. The evaluation results show that LRP enables

Please cite this article in press as: J. Zhou, et al., Optimizing cost for geo-distributed storage systems in online social networks, J. Comput. Sci. (2017), http://dx.doi.org/10.1016/j.jocs.2017.08.001