Computer Networks 81 (2015) 164–177

Contents lists available at ScienceDirect

Computer Networks

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

Minimizing data redundancy for high reliable cloud storage systems

Zhen Huang [a,1], Jinbang Chen [b,*,1], Yisong Lin [c], Pengfei You [a], Yuxing Peng [a]

[a] National University of Defense Technology, China
[b] East China Normal University, China
[c] Institute of GLD, China

ARTICLE INFO

Article history:
Received 8 October 2014
Received in revised form 30 January 2015
Accepted 17 February 2015
Available online 24 February 2015

Keywords:
Cloud storage system
Redundancy
Reliability
Generating function

ABSTRACT

Cloud storage system provides reliable service to users by widely deploying redundancy schemes in its system – which brings high reliability to the data storage, but inversely introduces significant overhead to the system, consisting of storage cost and energy consumption. The core behind this issue is how to leverage the relationship between data redundancy and data reliability. To optimize both concurrently is apparently difficult. As such, to fix one as a constraint and then to reach another one becomes the consensus. We aim in the paper to pursue a storage allocation scheme that minimizes the data redundancy while achieving a given (high) data reliability. For this purpose, we have provided a novel model based on generating function. With this model, we have proposed a practical and efficient storage allocation scheme, which is proved to be able to minimize the data redundancy. We analytically demonstrate that the suggested solution brings several advantages, in particular the reduction of the search space and the acceleration to the computation. We also assess the improvement on the savings of data redundancy experimentally by adopting availability traces collected from real world – which encouragingly shows that the reduction of data redundancy by our solution can reach up to more than 30% as compared to the heuristic method recently proposed in the research community.

1. Introduction

Cloud computing, with its promise to provide reliable service to users in an efficient and cheap manner, has attracted significant interests from both industry and academia [1]. Different from the supercomputing systems, inexpensive commodities are commonly used in cloud systems due to the consideration of scalability [2]. The reliability issue in such systems is thus particularly important. To ensure the data reliability, redundancy scheme is a basic solution and has been extensively deployed [3]. With this scheme, the intuitive idea is to store copies of data objects over a set of network nodes for the successful recovery in case of loss or failure. As such, one of the key issues is how to allocate redundant data over a given set of network nodes, which is conventionally referred to as the storage allocation, aiming at achieving maximum reliability by using minimum redundancy.

In general, two important factors need to be taken into account when discussing the storage allocation problem: data redundancy ("*RY*" for short) and data reliability

* Corresponding author at: Room 604, Information Building, 500 Dongchuan Road, 200241 Shanghai, China. Tel.: +86 (0)21 54 34 51 88.

E-mail addresses: maths_www@163.com (Z. Huang), jbchen@cs.ecnu.edu.cn (J. Chen).

[1] This study was carried out with National Laboratory of Parallel and Distributed Processing – Department of Computer – National University of Defense Technology – China, and Shanghai Key Laboratory of Multidimensional Information Processing – Department of Computer Science and Technology – East China Normal University – China.

http://dx.doi.org/10.1016/j.comnet.2015.02.013
1389-1286/

("*DR*" for short). Data redundancy *RY* is the answer to how much redundant data should be deployed in the system. Data reliability *DR* tells the probability that data is safely stored in the system for a certain duration. On one hand, increasing the data redundancy can improve the data reliability and at the same time disperse the I/O overhead in a distributed way. More importantly, it on the other hand increases the storage overhead and the energy consumption – remind that the power becomes an important concern today when designing the system. According to the reports [4,5] from IDC (International Data Corporation), currently there exists an evident gap between the demand on data creation and the storage capacity of hardware devices. And the gap will grow rapidly in the next coming decade. Furthermore, the storage systems consume around 37–40% of the energy of all IT components [6]. Therefore, minimizing data redundancy can significantly reduce the per-byte cost of a storage system. In this sense, minimizing *RY* is crucial to a cloud storage system.

It is clear to see that low data redundancy and high data reliability are in some sense contradictory, meaning that to find the optimal solution with minimum data redundancy and maximum data reliability is extremely hard. Thus, there exists a tradeoff between them. The usual way to cope with is to fix one and then try to find the optimal for the other one. For this purpose, to figure out their relations is necessary. However, it is non-trivial to formulate the relations for all these factors and further to derive the related properties. Still, storage allocation is an open and challenging issue in the research community.[2] In fact, optimizing data redundancy makes more sense in real cloud systems as: (i) besides redundancy scheme, such systems usually adopt many other mechanisms to guarantee the data reliability; (ii) minimizing data redundancy brings significant decrease on the cost of the distributed systems as presented before. For those considerations, we aim in this paper to pursue a storage allocation scheme that minimize data redundancy while achieving a given (high) data reliability – this objective is in line with what Pamies-Juarez et al. work for in [7,8]. In their method, Monte Carlo approximation is used to measure the data reliability, and a heuristic optimization algorithm is then followed to find the best assignment. In heterogeneous settings, their results have shown that data redundancy can be reduced up to 70% as compared to the traditional equal – allocation schemes. Although appealing, their method is descriptive (using a Monte Carlo approximation) and cannot be used for generic cases. Our contributions are as follows: we have reported in this paper new interesting results in the form of a new and novel model built based on a variant of generating function [9] and suggested for storage allocation in cloud storage systems. The advantage of the new model is the achievement of minimization of redundancy for a given (high) reliability. We have demonstrated that, our suggested method outperforms the existing solutions, in particular the heuristic method recently proposed in the research community. The solution we have proposed is practical – with this model, we are able to quantitatively analyze the properties of optimal storage allocation – which allows us to reduce search space and accelerate the computation. In the end, we not only analytically evaluate the reduction of the search space, but also experimentally assess the improvement on the savings of data redundancy using particular sets of data collected from real world.

The remaining of this paper is organized as follows: Section 2 illustrates the redundancy scheme and reviews the previous work on storage allocation. An efficient solution we propose for the storage allocation are then given in Section 3. In Section 4, we evaluate our method through analysis and experiments by adopting availability traces collected from real world. We finally conclude the paper with a summary and future work in Section 5.

2. Background and related work

2.1. Redundancy scheme

In cloud storage system, using redundancy scheme to achieve data reliability is straightforward. To fight against failures, it distributes copies of data objects to a set of storage nodes. Redundancy schemes have been widely studied in the research community [10–12]. Generally speaking, redundancy scheme can be classified into two types: replication and erasure code. Replication scheme is simple and intuitional, which replicates each data block into n copies and then distributes them to different network nodes. By contrast, erasure code (*e.g.*, Reed–Solomon [13]) encodes k data blocks into $(n-k)$ coded blocks, resulting in n blocks in total. These n blocks are then distributed into different nodes [14,15].

Conventionally, the data redundancy of erasure code is defined as $RY = n/k$. It is clear to see that when $k = 1$, erasure code turns into replication. Thus, replication can be regarded as the special case of erasure code. When $k > 1$, erasure code consumes less storage space as compared to replication [16], meaning that erasure code can improve the reliability with the same consumption of storage. Due to its flexibility, erasure code becomes a promising solution for the next-generation cloud storage systems and have been explosively discussed more recently [17–19]. Some practical cloud storage systems have started to partly deploy erasure codes, such as GFS [2] and HDFS [20]. Some other cloud backup systems have fully deployed erasure cods, including Wuala [21] and CleverSafe [22]. Remind that we mainly focus on erasure code in this paper when discussing the cloud storage allocation. However, our method can be easily extended to the case of replication scheme.

2.2. Storage allocation problem

Storage allocation problems arise from the emergence of coding schemes that store more than one redundant block into the same storage node, such as the Regenerating Code [23] and ER-Hierarchical Code [24], where the repair cost for lost data can be reduced. In this case, the computation for data reliability can be very

[2] see http://storagewiki.ece.utexas.edu/doku.php?id=wiki:open_problems.