Approach to Big Data Applications on Cloud System

Jinhong Kim¹, Sung-Tae Hwang²,

¹ College of Information and Communication Engineering, SungKyunKwan University, Suwon, Korea ² College of Engineering, Hansung University, Seoul, Korea {jinhkm, cosmosth}@gmail.com

Abstract. Consistency management within cloud storage systems is of high importance. The Consistency–Performance tradeoff is, arguably, the main tradeoff. Many static consistency solutions fail in reaching an efficient equilibrium between consistency and performance for dynamic cloud workloads. In this context, more opportunistic, adaptive consistency models are needed in order to meet the requirements of Big Data Applications (BDAs). However, most of the existing adaptive policies either lack automation or fail to apprehend and include the specific consistency requirements of the application outside its access pattern. Hereafter, in this paper, we tackle this specific issue of consistency management for dynamic workloads in the cloud. Accordingly, we provide an adaptive model that tunes the consistency level at runtime in order to provide consistency when needed and performance when possible.

Keywords: Big Data Applications (BDAs), Cloud System, ACM

1 Introduction

Recently, data sizes have been growing exponentially within many organizations. Data is everywhere and comes from multiple sources: social media, smart phones, sensors etc. Big Data introduces numerous complications to the different aspects of data storage and management. These complications are due to the overwhelming sizes, but also the velocity required and the complexity of data coming from different sources with different requirements at high load variability [1],[2]. In order to deal with the related challenges, many Big Data systems rely on large and novel infrastructures, as well as new platforms and programming models. In this context, the emerging paradigm of Cloud Computing offers excellent means for Big Data. Thereby, corporations can acquire the resources needed for their Big Data applications at a low cost when needed. Meanwhile, they avoid large investments on physical infrastructures that need huge efforts for building and maintaining them, which, in addition, requires a high level of expertise. Within cloud storage, replication is a very important feature for Big Data. At wide area cloud scales, data is replicated

² Prof. Sung-Tae Hwang is Corresponding Author

Advanced Science and Technology Letters Vol.98 (CES-CUBE 2015)

across multiple data centers in order to deal with fast response and local availability requirements. Therefore, clients can request data locally from a replica within the closest datacenter and get a fast response [3]. Moreover, geographical replication provides data durability, fault tolerance and disaster recovery by duplicating redundant data in different geographical areas [4]. However, one issue that arises with replication is guaranteeing data consistency across replicas. In this context, insuring strong consistency requires huge synchronization efforts across different locations and thus, exposes the users to high network latencies [5]. This affects the performance and the availability of cloud storage solutions. One particular alternative that has become very popular is eventual consistency. Eventual consistency may tolerate inconsistency at some points in time but guarantees the convergence of all replicas to the same state at a future time. The management of consistency heavily impacts storage systems.

2 Related Works

2.1 Big Data

Big Data is more than growing sizes of datasets, but rather the complexity that such a growth generates with regard to different aspects of data handling. Big Data refers to datasets that are of a big volume, need big velocity, or exhibit big variety. **Big Volume.** Nowadays, data sizes are exponentially increasing. Multiple companies and organizations are experiencing Data Deluge phenomenon due to multiple factors. **Big Velocity**. For many organizations, the most challenging aspect of Big Data is not exclusively the large volume. It is rather, how fast to process data to meet demands. A wide range of applications require fast data processing in near real-time manner, no matter how overwhelming the size of data. **Big Variety.** Data comes in various formats: from text data to video and audio, from traditional DBMS formats to semi-structured data, to large binary objects. The primary challenge is to integrate all these types of data, and manage them in an efficient way [6].

2.2 Big Data Application in the Cloud

The alarming growth rate of Big Data introduces unprecedented challenges and problems to IT experts. One of the first and major problems was of infrastructure nature. Companies and organizations face an overwhelming data tsunami that exceeds their infrastructure capacities. The emergence of the Cloud Computing paradigm provided a safety net as it provides an infinite amount of resources on-demand over Internet at a fair price [7]. Moreover, major cloud vendors that possess the right tools and experience to deal with Big Data, provide storage solutions and Big Data platforms to

Advanced Science and Technology Letters Vol.98 (CES-CUBE 2015)

3 ACM: Adaptive Consistency Model

The way consistency is handled has a big impact on performance. Traditional synchronous replication dictates that an update must be propagated to all the replicas before returning a success. In contrast, eventual consistency by means of asynchronous quorum replication propagates data lazily to other replicas. Here the consistency level is commonly chosen on a per-operation basis and is represented by the number of replicas in a subset of all the replicas. Data accesses and updates are performed to all replicas in the quorum. Thus, using this level for both read and write operations guarantees that the intersection of replicas involved in both operations contains at least one replica with the latest update [8]. A partial quorum has a smaller subset of replicas, hence returning the most recent data when read is issued, is not guaranteed. In the following, the term consistency level refers to the number of replicas involved in the access operation. Additionally, in the case of distributed data replication, network latency may be high and thus, a performance-defining factor.

4 Implementation and Evaluation

We use two clusters in the Sophia site with a total of 84 nodes and 496 cores. All nodes are equipped with x86_64 CPUs and 4 GB of memory. The nodes are interconnected with Gigabit Ethernet. All nodes from the first cluster have two hard disks with combined capacity of 600 GB per node. As for the second cluster, the nodes are all equipped with hard disks of 250 GB. We have first studied the impact of the workload access patterns, the number of clients, and network latency on the stale reads estimation. Accordingly, we use two workloads: workload-A, which has a heavy read-update access pattern, and workload-B, which has a heavy read access pattern with a small portion of writes representing approximately 5% of the total number of operations as shown in Figure 1.



Fig. 1. Workloads and Number of Clients impact on stale read probability(Left), Network latency impact on stale read probability(Right)

We can see that for workload-A, the probability of stale reads gradually decreases with the number of threads, because increasing the thread number increases the throughput and thus increases the reads and writes rate. Also, we notice the probability reduction gap is big during the transition (changing the number of threads).

5 Conclusions

Eventual consistency was employed in cloud storage system as an alternative to traditional strong consistency to achieve scalable and high-performance services. Many commercial cloud storage systems have already adopted the eventual consistency approach. In this paper, we presented Cloud System, a novel approach that handles data consistency in cloud storage adaptively by choosing the most appropriate consistency level dynamically at run time. Our approach is designed to be completely tunable to provide the system or the application administrator with the possibility of controlling the degree of compromise between performance and consistency.

References

- 1. Rajkumar Buyya et al. "Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility". In: Future Gener. Comput. Syst. 25.6 (2009), pp. 599–616.
- 2. Brad Calder et al. "Windows Azure Storage: a highly available cloud storage service with strong consistency". In: Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles. SOSP '11. Cascais, Portugal: ACM, 2011, pp. 143–157.
- 3. Jeffrey Dean and Sanjay Ghemawat. "MapReduce: simplified data processing on large clusters". In: Commun. ACM 51.1 (Jan. 2008), pp. 107–113.
- 4. Martin Ester et al. "A density-based algorithm for discovering clusters in large spatial databases with noise". In: Second International Conference on Knowledge Discovery and Data Mining. AAAI Press, 1996, pp. 226–231.
- 5. Seth Gilbert and Nancy Lynch. "Brewer's Conjecture and the Feasibility of Consistent Available Partition-Tolerant Web Services". In: ACM SIGACT News. 2002.
- 6. Shadi Ibrahim et al. "Adaptive Disk I/O Scheduling for MapReduce in Virtualized Environment". In: Proceedings of the 2011 International Conference on Parallel Processing (ICPP'11). ICPP '11. Taipei, Taiwan, 2011, pp. 335–344.
- Rini T. Kaushik and Milind Bhandarkar. "GreenHDFS: towards an energyconserving, storage-efficient, hybrid Hadoop compute cluster". In: Proceedings of the 2010 international conference on Power aware computing and systems. HotPower'10. Vancouver, BC, Canada: USENIX Association, 2010, pp. 1–9.
- Cheng Li et al. "Making geo-replicated systems fast as possible, consistent when necessary". In: Proceedings of the 10th USENIX conference on Operating Systems Design and Implementation. OSDI'12. Hollywood, CA, USA: USENIX Association, 2012, pp. 265– 278.