ISSN 2616-5767 Vol.7, Issue 6: 123-127, DOI: 10.25236/AJETS.2024.070618

Efficient Storage and Access Management System for Network Databases Driven by Cloud Computing Technology

Yiru Zhang

Consumer Electronics Technology, Amazon, New York, NY 10044, USA
yiruz315@163.com

Abstract: *This article focuses on data storage and computing issues in cloud environments. With the advent of the big data era, cloud computing has become a key support for data storage and intelligent computing. This article analyzes the key challenges in cloud storage and computing, including data reliability, availability, and security, as well as the resource coordination scheduling problem between GPU and CPU in intelligent computing. In response to these challenges, this article proposes multiple innovative solutions: firstly, the DMcache extension architecture MapperX is designed, which improves the availability of storage systems and shortens disk recovery time through adaptive metadata bittrees; Secondly, a SwornDisk encrypted storage method combining LSM architecture and remote update encryption mechanism was proposed to enhance data security; Thirdly, the Elastic Scheduler framework has been developed to achieve elastic scheduling of GPU and CPU collaborative computing through local gradient accumulation algorithm; The fourth is the design of ParaX acceleration method, which optimizes the deep learning computing performance of multi-core CPUs through the "single instance single core" strategy. These methods have achieved significant performance improvements in experiments. Finally, this article looks forward to future research directions, including intelligent ABT highly available collaborative storage mechanisms, highly readable SwornDisk technology, and GPU-CPU collaborative elastic computing methods combining ES and ParaX.*

Keywords: *Cloud storage and computing, DMcache, SwornDisk encrypted storage, Elastic Scheduler, ParaX multi-core CPU acceleration method*

1. Introduction

With the rapid development of the Internet, the Internet of Things and artificial intelligence technology, the global big data industry is entering a golden age of rapid growth. Through in-depth mining of data value, various fields have driven innovation and progress in medical, education, commerce, industry, agriculture and other fields. In this context, data storage and computation, as the core of big data processing, face the dual challenges of data explosion and diversified types, and traditional technologies are no longer able to meet the demands. Cloud computing, with its powerful storage, computing, and analytical capabilities, has become a key force in addressing these challenges. However, although the IaaS layer of public cloud services serves as the foundational support, it is often constrained by issues such as low resource utilization, high task latency, insufficient business resilience, and data security risks. To this end, cloud storage needs to enhance the reliability, availability, and security of data, and improve computing efficiency and resource scheduling flexibility. This paper focuses on the efficiency, reliability, and dynamic challenges of data storage and intelligent computing in cloud environments. By proposing the MapperX high availability collaborative storage mechanism, SwornDisk high security log structure storage solution, Elastic Scheduler deep learning computing elastic scheduling framework, and ParaX multi-core CPU memory access acceleration method, it not only provides innovative solutions to current technical problems, but also lays a solid foundation for the development of high-performance computing and big data processing technologies in the future. This series of studies not only promotes technological integration and innovation, but also drives the transformation of society towards intelligence and sustainability. It requires us to constantly learn new skills, adjust our ways of thinking, and respond to challenges in the rapidly changing wave of technology, in order to maintain our competitive advantage industry.

ISSN 2616-5767 Vol.7, Issue 6: 123-127, DOI: 10.25236/AJETS.2024.070618

2. Correlation Theory

2.1 Cloud Data Storage

Data storage in cloud environments is a top priority in cloud computing research, and significant progress has been made in two directions in recent years: optimizing the internal storage architecture of single machine nodes and expanding the distributed storage architecture. On the one hand, for single machine nodes, researchers are committed to improving storage efficiency and availability through cost-effective storage solutions, such as adopting HDD-SSD collaborative storage architecture [1], accelerating I/O performance by using fast SSDs as cache for HDDs, while maintaining the low-cost advantage of HDDs. However, such solutions face issues such as long recovery times after crashes and complex maintenance of data consistency. On the other hand, in the face of the massive data storage demands of the big data era, distributed storage architecture has become the mainstream trend. Distributed storage not only expands the capacity of standalone storage, but also provides flexible logical storage interfaces through virtualization technology, such as HDFS and Ceph cloud storage systems. These systems ensure data reliability and availability through redundant replica technology, and use block storage mechanisms to optimize storage load balancing. In addition, log structured storage (LSM)[2], as an innovation at the software level, effectively improves write performance and reduces space write amplification issues through sequential write optimization and merging mechanisms, and is widely used in the design of databases and file systems. In summary, data storage technology in cloud environments is continuously optimizing the efficiency of single node storage and expanding distributed storage architectures to meet the high performance, high availability, and low-cost requirements of big data processing.

2.2 Cloud Security Field

In cloud environments, the risk of data leakage and tampering significantly increases due to the elimination of physical boundaries of user data and the sharing of computing resources. In order to address this challenge, the field of cloud security has extensively researched various technical means, among which encryption algorithms, Trusted Execution Environment (TEE) [3], and storage encryption are particularly critical. Encryption algorithms such as AES and RSA ensure data security and integrity through symmetric and asymmetric methods, respectively, but face the complexity of cloud key management. TEE technology, such as Intel SGX, AMD SEV, etc., effectively resists external attacks by isolating the trusted program running environment, although existing solutions have trade-off issues between security and performance. In terms of storage encryption, Intel SGX's PFS system achieves data confidentiality, integrity, and freshness protection through a variant MHT structure, but its shortcomings include unencrypted metadata and low efficiency in random write IO. Although other secure file systems such as Oblivate and SecureFS have made breakthroughs in security, performance overhead is still an issue that cannot be ignored. Therefore, the continuous innovation and optimization of cloud security technology aim to balance security and efficiency, ensuring worry free security of cloud user data and applications.

3. Research Method

3.1 Secure and Fast Storage Encryption System

Cloud security faces severe challenges in the field of cloud computing, especially when user data is physically blurred in public cloud environments, making it vulnerable to attacks such as eavesdropping, theft, and tampering. Although Trusted Execution Environment (TEE) technologies such as Intel SGX, AMD SEV, and ARM TrustZone enhance the security of cloud server memory data through encryption and strong isolation of user memory data, secure storage of application data on disk remains a major challenge, placing higher demands on data confidentiality, integrity, and freshness. Existing TEE technologies such as Intel SGX's PFS system provide strong security based on MHT architecture, but the O (logN) I/O overhead of read and write operations and unencrypted metadata increase performance burden and side channel attack risk.To overcome these limitations, this chapter innovatively proposes a secure and fast storage encryption system SwornDisk based on the Log Structure Merge Tree (LSM) architecture. The system adopts a remote update encryption mechanism, effectively avoiding the complexity and inefficiency of traditional local update modes. It not only solves the problem of rollback attacks, but also ensures the confidentiality, integrity, and anonymity of data. Its core lies in the introduction of Merkle Hash B-Tree (MHBT) storage encryption structure, combined with the security of MHT and the fast lookup capability of B-Tree, to achieve O (1) complexity write mechanism and O (logN) read complexity, significantly improving I/O performance. In addition, SwornDisk provides a low-level block device interface, enhancing the flexibility and application scope of TEE storage. Experimental results have shown that SwornDisk significantly improves write performance while ensuring security, especially in random write testing where performance is increased by up to 116 times, providing an efficient and reliable solution for data security in cloud environments.

3.2 Cloud Storage and Intelligent Computing

This article has achieved significant research results and innovative breakthroughs in the field of cloud computing, especially in cloud storage and intelligent computing. A new extension MapperX for DMcache hybrid architecture is proposed to address the availability issue of hybrid storage architecture, which utilizes adaptive meta information bit tree (ABT) technology to finely manage dirty bit data[4], significantly improving the availability and recovery efficiency of storage systems. At the same time, to address the security challenges of encrypted storage in cloud environments, SwornDisk was designed, which combines LSM architecture with remote update encryption mechanism, effectively ensuring multiple security attributes of data and optimizing IO performance. In the field of intelligent computing, the Elastic Scheduler (ES) framework is innovatively proposed to address the dynamic elastic scheduling problem of GPU/CPU collaborative computing. The framework utilizes local gradient accumulation algorithm to solve the speed mismatch and momentum compensation problems, achieving efficient collaborative computing. In addition, the ParaX method was introduced to address the issue of memory contention in deep learning with multi-core CPUs. Through the "single instance single core" strategy and optimized core communication mechanism, the processing speed of CPUs in deep learning tasks was significantly improved.

Looking ahead to the future, the research work in this article will be further deepened, including designing an adaptive meta information bit tree structure for intelligent decision-making to optimize MapperX performance, optimizing SwornDisk's caching strategy to enhance read performance, integrating ES with ParaX to build a complete GPU-CPU collaborative elastic deep learning system, and introducing load balancing and fault-tolerant mechanisms for container management to comprehensively improve the flexibility and reliability of cloud computing systems. Especially in the field of deep learning, the proposal of the ES framework breaks the traditional boundaries between CPU and GPU in collaborative computing. By dynamically adjusting batch sizes and introducing momentum compensation mechanisms, it effectively solves the synchronization problem caused by speed mismatch between CPU and GPU, significantly improving the efficiency and accuracy of model training in dynamic training environments. The experimental results show that ES has significantly improved the training and inference performance of DNN applications such as ResNet50 and MobileNet-v1 on the MXNet framework, and performs well in dynamic training scenarios, demonstrating its high reliability and robustness. This series of innovations not only provides new ideas for the development of cloud computing technology, but also lays a solid foundation for deep learning to achieve efficient and flexible training in complex and ever-changing cloud environments.

4. Results and Discussion

4.1 Research Background

With the increasing demand for hardware computing power in deep learning (DL), more and more DL tasks are being deployed to cloud environments, utilizing the vast computing resources in the cloud. Although GPUs perform well in floating-point operations and memory bandwidth, CPUs are widely used in DL training and inference tasks due to their low latency characteristics and the presence of a large number of idle CPUs during off peak hours in cloud environments. Especially for multi-core CPUs with NUMA (Non Consistent Memory Access) architecture, integrating multi slot CPUs can improve overall processor performance and help accelerate the processing speed of DL tasks. However, as DNN models become increasingly complex, the demand for memory bandwidth significantly increases, and CPUs face the problem of insufficient bandwidth utilization when processing these models.The current DL platforms mostly adopt the "single CPU, single instance" mode, where each GPU acts as a worker and is assigned an instance. This mode is effective for GPUs with high memory bandwidth, but when directly applied to CPUs with low memory bandwidth, it can lead to competition between cores for limited memory bandwidth, resulting in core starvation and a serious decrease in parallelism. In addition, communication issues across NUMA nodes within the CPU are often overlooked, and traditional distributed communication frameworks generate additional overhead within a single machine, limiting performance improvement.

To address these issues, this article analyzes the main reasons for the low efficiency of multi-core CPUs in DL platforms and proposes the ParaX system. ParaX uses the "single core single instance" method to divide input data into multiple mini sub batches and allocate an instance to each CPU core, thereby removing the synchronization barrier of executing each layer of the DNN model. Through a super light scheduling strategy, it utilizes the randomness of the execution time of each layer to fully overlap the memory intensive network layer and the computation intensive network layer, in order to improve the overall bandwidth utilization. Meanwhile, ParaX has designed a gradient server communication mechanism based on shared memory [5], which reduces the synchronization overhead within NUMA and improves the overall training speed. The experimental results show that ParaX significantly improves the execution speed of training and inference on dual NUMA Intel 8280 CPUs, achieving training speed improvements of 1.73 times to 2.93 times and inference speed improvements of 2.08 times to 2.11 times, respectively.In summary, ParaX effectively solves the inefficiency problem of multi-core CPUs in DL platforms by optimizing CPU bandwidth utilization and communication efficiency across NUMA nodes, providing a new solution for efficient execution of DL in cloud environments.

4.2 Systems Design

The ParaX system is ingeniously designed to address the challenges of tight memory bandwidth in multi-core CPU environments [6] and inefficient communication across NUMA architectures, significantly improving the efficiency of deep learning training through three core innovative strategies. Firstly, the single core single instance method is adopted, where each CPU core is independently assigned a DNN model instance, achieving bandwidth borrowing between computationally intensive and memory intensive network layers, effectively alleviating memory bandwidth pressure and eliminating synchronization barriers. Secondly, the introduction of gradient server communication mechanism utilizes shared memory to directly access weights, avoiding the additional overhead of traditional parameter update mechanisms, reducing remote memory access requirements, and significantly improving the efficiency of parameter updates. Finally, for the distributed training environment, an x-core single instance method was designed to allocate instance and core resources reasonably, while maximizing the computational potential of multi-core CPUs while maintaining high-precision training. This series of innovative designs together form the ParaX system, paving the way for efficient training of deep learning on multi-core CPU platforms.

4.3 Actual Testing Verification

In this paper, the function and performance of ParaX system on MXNet platform are tested comprehensively, and the single-core single-instance mode, multi-core single-instance mode, overlapping scheduling and gradient server mechanism are tested deeply. The experimental environment is a server equipped with dual socket Intel Platinum 8280 CPUs, with 28 2.70GHz cores per socket, totaling 56 cores, supplemented by 39424KB L3 cache, 192GB DDR4-2933 memory, and Intel 750 PCIE 400GB NVMe Express SSD. We have selected multiple DNN models, including image recognition (ResNet MobileNet, Inception-BN) Natural language processing [7] (LSTM, GNMT) and recommendation systems (NCF) were tested using FP32 data type. For the performance of ParaX, we have set a reasonable batch size to avoid memory overflow and ensure training accuracy, while evaluating its performance in different applications.

The experimental results show that ParaX has significantly improved training and inference performance compared to the original MXNet, especially in single CPU single instance mode. By overlapping computation intensive and memory access intensive network layers and utilizing gradient server mechanism, ParaX achieves training acceleration from 1.73 × to 2.93 × and inference acceleration from 2.08 × to 2.11 ×. For LSTM, due to its sensitivity to batch size, we also tested its performance under different batches, and the results showed that ParaX approached GPU performance in terms of inference speed after quantization optimization. Furthermore, ParaX also performs well in CPU utilization and memory bandwidth utilization, reaching 94.8% and 88.6% respectively, significantly better than traditional methods.

In the evaluation of the gradient server mechanism, we compared it with the state-of-the-art PS and RAR mechanisms. Experiments have shown that as the number of instances increases, the GS mechanism outperforms PS and RAR in both throughput and update latency, thanks to its shared memory mechanism that reduces memory copy overhead and fully utilizes all cores. Meanwhile, the

GS mechanism also exhibits lower data transmission burden in terms of total data access.

In addition, we also explored the impact of batch size on ParaX performance. The experiment shows that after the total batch reaches a certain scale, the growth of throughput tends to flatten out, while the memory footprint increases with the increase of batch size, but it does not become a performance bottleneck. This result validates the efficiency of ParaX in memory bandwidth and computing resource management, providing strong support for efficient training of deep learning in multi-core CPU environments.

5. Conclusion

This article conducts in-depth research and innovation in the field of cloud computing, especially in cloud storage and intelligent computing, and proposes effective solutions to several key problems that currently exist. Firstly, in response to the availability issue of hybrid storage architecture in cloud environments, a new extension MapperX of DMcache hybrid architecture is proposed, which achieves fine management and optimized recovery of dirty bit data through adaptive meta information bit tree (ABT), significantly improving the overall availability of storage. Secondly, in response to the security issues of encrypted storage in cloud environments, SwornDisk was designed based on LSM architecture and remote update encryption mechanism, effectively ensuring the confidentiality, integrity, freshness, and anonymity of data while improving IO performance. Thirdly, in response to the dynamic elastic scheduling problem of GPU/CPU collaborative computing, the Elastic Scheduler (ES) framework[8] was proposed, which solved the speed mismatch and momentum compensation problems in dynamic computing through local gradient accumulation algorithm, achieving efficient collaborative computing. Finally, the ParaX method is proposed to address the issue of memory contention in deep learning computation with multi-core CPUs. By implementing a "single instance, single core" execution strategy and optimized core communication mechanism, the method significantly improves the processing speed of CPUs in deep learning tasks.

Looking ahead to the future, the work of this article can be further deepened and expanded in the following aspects: firstly, designing an adaptive meta information bit tree structure for intelligent decision-making to further enhance the availability and recovery efficiency of MapperX; Secondly, optimize SwornDisk's caching strategy and support Cache line capability to improve its read performance; The third is to combine ES with ParaX to implement a complete GPU-CPU collaborative elastic deep learning method, and introduce container based load balancing and fault-tolerant mechanisms to further enhance the flexibility and reliability of cloud computing systems. These future research directions will provide new ideas and solutions for the further development of cloud computing technology.

References

[1] Wang S, Lu Z, Cao Q, et al. Exploration and Exploitation for Buffer-Controlled HDD-Writes for SSD-HDD Hybrid Storage Server [J]. Transactions on Storage (TOS), 2022.

[2] Spillane R P, Wang W, Gao J, et al. Using an LSM tree file structure for the on-disk format of an object storage platform:US201816213714[P]. US11093472B2 [2024-09-13].

[3] Karame G, Soriente C. Method And System For Performing Remote Attestation With A Gateway In The Context Of A Trusted Execution Environment (TEE):US201917435712 [P]. US2022156390A1 [2024-09-13].

[4] Yu Y, He W, Jin J, et al. iDNA-ABT: advanced deep learning model for detecting DNA methylation with adaptive features and transductive information maximization [J]. Bioinformatics, 2021. DOI:10.1093/bioinformatics/btab677.

[5] Chung M T, Weidendorfer J, Kranzlmueller F D. From reactive to proactive load balancing for task-based parallel applications in distributedmemory machines [J]. Concurrency and computation: practice and experience, 2023, 35(24): 7828-7829. DOI:10.1002/cpe.7828.

[6] Schwartz J, Kurniawati H, Hutter M. Combining a Meta-Policy and Monte-Carlo Planning for Scalable Type-Based Reasoning in Partially Observable Environments[J]. ArXiv, 2023, 6067. DOI:10.48550/arXiv.2306.06067.

[7] Albraikan A A, Maray M, Alotaibi, Faiz Abdullah Alnfiai, Mrim M. Kumar, ArunSayed, Ahmed. Bio-Inspired Artificial Intelligence with Natural Language Processing Based on Deceptive Content Detection in Social Networking [J]. Biomimetics, 2023, 8(6).

[8] Phan L X, Chishti M, Miller R J, et al. Systems and methods for varying elastic modulusappliances:US202117383286[P]. US2021347103A1 [2024-09-13]. DOI: US6964564 B2.