INTERNATIONAL RESEARCH JOURNAL OF ENGINEERING & APPLIED SCIENCES
ISSN: 2322-0821(0) ISSN: 2394-9910(P)
VOLUME 7 ISSUE 2 April 2019 - Jun 2019
www.irjeas.org

Latency Reduction Techniques in Distributed Cloud Systems for Financial Applications

Pavan Nutalapati [1]

[1] *Project Lead at Oracle*
pavan.nutalapati@gmail.com

Abstract: The financial industry is increasingly dependent on distributed cloud systems due to their capacity to manage enormous volumes of data, facilitate high-frequency trading, and conduct real-time analytics. These systems are crucial for modern financial operations, yet latency presents a significant obstacle that affects trading performance and decision-making accuracy. This review paper delves into various techniques for mitigating latency within distributed cloud environments, with a particular emphasis on their application and efficacy in financial settings. It provides an in-depth analysis of strategies such as optimizing data locality to reduce the distance data must travel, enhancing network infrastructure and computational resources to speed up processing times, implementing advanced load balancing to distribute workloads more effectively, and utilizing latency-aware algorithms to predict and counteract potential delays. The paper also explores challenges that impact latency, including the complexities of geographic distribution, variability in network performance, and the associated costs of implementing latency-reduction measures. By examining both current practices and emerging advancements in the field, this review aims to offer a thorough understanding of how to address latency issues in financial systems, ultimately improving system performance and reliability.

Keyword: Latency Reduction, Distributed Cloud Systems, Financial Applications, Network Optimization, Computational Strategies, Data Management, Architectural Improvements, Content Delivery Networks, Edge Computing.

1. INTRODUCTION

In today's fast-paced financial markets, where milliseconds can determine profit or loss, latency in distributed cloud systems has emerged as a critical concern. Latency refers to the delay between initiating a request and receiving a response, and even small delays can significantly impact the efficiency and effectiveness of financial applications. As financial services increasingly rely on distributed cloud systems for their scalability and flexibility, they must also contend with the challenges of latency. These systems handle complex and voluminous data requirements, but their distributed nature introduces inherent delays.

This latency can affect transaction speed, making it harder to capitalize on fleeting market opportunities. Additionally, it can impair market predictions by delaying data processing and analysis, which in turn affects trading strategies and decision-making. Overall, the impact of latency in distributed cloud systems can undermine operational performance, emphasizing the need for innovative solutions to minimize delays and optimize performance in high-frequency trading and other time-sensitive financial activities.

As financial institutions and trading platforms strive to maintain a competitive edge, minimizing latency has become essential. High-frequency trading, real-time fraud detection, and instantaneous market data analysis all demand low-latency solutions to ensure timely and accurate operations. Consequently, addressing latency issues in distributed cloud systems is of paramount importance for maintaining the integrity and responsiveness of financial applications.

This review paper delves into a range of advanced techniques and strategies designed to minimize latency in distributed cloud systems, with a particular emphasis on their implications for the financial services sector. We will comprehensively explore network optimization techniques, including the deployment of content delivery networks (CDNs) and edge computing architectures, which aim to reduce latency by bringing data and services closer to end users. Additionally, we will investigate computational strategies such as parallel processing and dynamic load balancing that enhance processing efficiency and reduce response times. Data management practices, such as sophisticated caching mechanisms and data compression algorithms, will be analyzed for their role in decreasing data retrieval times and improving throughput. The review will also cover architectural innovations like microservices and serverless computing, which offer scalable and resilient solutions to handle varying workloads and demand patterns. By integrating recent advancements and real-world implementations, this paper seeks to offer a thorough overview of effective latency reduction methods, emphasizing their critical importance in optimizing the performance and reliability of financial applications, where rapid processing and real-time data access are paramount.

In the following sections, we will comprehensively explore various categories of latency reduction techniques, detailing their underlying mechanisms, key advantages, and specific relevance to the financial sector. We will examine methods such as network optimization, data caching, and edge computing, evaluating how each technique can be effectively implemented to mitigate latency issues. By providing a thorough analysis of these strategies, we aim to offer valuable insights and practical guidance for financial institutions. Our goal is to assist these institutions in enhancing their distributed cloud systems, ultimately leading to improved operational efficiency and reduced latency, which are critical for maintaining a competitive edge in the fast-paced financial markets.

2. NETWORK OPTIMIZATION TECHNIQUES

Network optimization is a critical aspect of reducing latency in distributed cloud systems, particularly for applications that require real-time or near-real-time processing. Efficient network management can help minimize delays caused by data transmission and improve overall system performance. This section explores several key network optimization techniques that are relevant to financial applications.

2.1. Content Delivery Networks (CDNs) Content Delivery Networks (CDNs) enhance network performance by distributing and caching content across a global network of servers. By storing copies of frequently accessed data closer to end-users, CDNs can significantly reduce latency associated with data retrieval. In financial applications, CDNs can be employed to cache static content such as historical market data, trading algorithms, and frequently accessed reports, thereby reducing the time needed to deliver this information to end-users.

2.2. Edge Computing Edge computing involves processing data closer to its source, rather than relying on centralized data centers. This approach reduces the distance data must travel, thereby decreasing latency. For financial applications, deploying edge nodes at strategic locations can improve the responsiveness of high-frequency trading platforms, real-time analytics, and fraud detection systems. By processing data at the edge of the network, financial institutions can achieve faster decision-making and more efficient operations.

2.3. Network Protocol Optimization Optimizing network protocols can play a significant role in reducing latency. Traditional protocols like TCP can introduce delays due to their connection-oriented nature and error correction mechanisms. Modern protocols, such as QUIC (Quick UDP Internet Connections), are designed to enhance performance for real-time applications by reducing connection establishment time, minimizing packet loss, and enabling faster data transmission. Implementing these optimized protocols can lead to noticeable improvements in latency for financial transactions and data exchanges.

2.4. Network Function Virtualization (NFV) Network Function Virtualization (NFV) involves virtualizing network functions that were traditionally implemented in hardware. By using software-based network functions, NFV can provide greater flexibility and efficiency in managing network resources. In financial systems, NFV can be used to deploy and manage network functions such as firewalls, load balancers, and intrusion detection systems in a more agile manner, potentially reducing latency associated with network management tasks.

2.5. Software-Defined Networking (SDN) Software-Defined Networking (SDN) separates the control plane from the data plane in network architecture, allowing for more centralized and dynamic management of network traffic. SDN enables real-time adjustments to network configurations, which can help optimize traffic flow and reduce latency. For financial applications, SDN can facilitate efficient traffic management, prioritize critical data flows, and quickly adapt to changing network conditions, thereby minimizing delays.

2.6. Data Compression and Decompression Data compression techniques reduce the volume of data transmitted over the network, which can decrease transmission times and reduce latency. By compressing data before transmission and decompressing it upon receipt, financial applications can achieve faster data transfer rates. Implementing efficient compression algorithms and ensuring that both ends of the communication channel support these algorithms are essential for optimizing latency.

2.7. Traffic Shaping and Quality of Service (QoS) Traffic shaping and Quality of Service (QoS) mechanisms help manage and prioritize network traffic based on predefined policies. By ensuring that high-priority financial transactions and data flows receive preferential treatment, these techniques can reduce latency and improve the performance of critical applications. Implementing QoS policies can help mitigate the impact of network congestion and ensure that latency-sensitive financial services operate smoothly.

By employing these network optimization techniques, financial institutions can effectively address latency issues and enhance the performance of their distributed cloud systems. Each technique offers distinct advantages and can be tailored to meet the specific needs of financial applications, contributing to faster and more efficient operations.

3. COMPUTATIONAL STRATEGIES

Computational strategies play a pivotal role in reducing latency in distributed cloud systems, particularly for applications that demand high-speed data processing and real-time analysis. Effective computational approaches can optimize resource utilization, accelerate processing times, and ensure timely responses for financial applications. This section discusses key computational strategies that are instrumental in minimizing latency.

3.1. Parallel Processing Parallel processing involves dividing computational tasks into smaller sub-tasks that can be executed simultaneously across multiple processors or cores. This approach enhances processing speed and reduces overall computation time. In financial applications, parallel processing is crucial for tasks such as real-time market analysis, high-frequency trading algorithms, and complex risk modeling. By leveraging parallel computing frameworks and architectures, financial institutions can achieve faster data processing and more responsive systems.

3.2. Load Balancing Load balancing distributes computational tasks and workloads across multiple servers or resources to prevent any single server from becoming a performance bottleneck. Effective load balancing ensures that resources are utilized efficiently and helps maintain low latency by preventing server overloads. Techniques such as dynamic load balancing and adaptive algorithms can optimize resource allocation based on real-time demand, thereby enhancing the performance of financial applications and ensuring consistent response times.

3.3. Resource Provisioning and Auto-Scaling Resource provisioning involves allocating computational resources based on the current workload and demand. Auto-scaling is a dynamic resource provisioning strategy that automatically adjusts the number of active resources based on predefined thresholds or real-time metrics. For financial applications, auto-scaling ensures that adequate computational resources are available during peak periods, such as market surges or high trading volumes. By aligning resource allocation with demand, auto-scaling helps reduce latency and improve system efficiency.

3.4. High-Performance Computing (HPC) High-Performance Computing (HPC) refers to the use of powerful computing resources to perform complex calculations and simulations at high speeds. HPC systems, such as supercomputers or clusters, are designed to handle large-scale computations and data-intensive tasks. In financial services, HPC can be utilized for applications such as quantitative modeling, risk assessment, and large-scale data analysis. By leveraging HPC capabilities, financial institutions can achieve faster computation times and reduced latency for critical operations.

3.5. In-Memory Computing In-memory computing involves storing data in the system's RAM rather than on traditional disk storage. This approach significantly speeds up data access and processing times by reducing the latency associated with disk I/O operations. For financial applications, in-memory computing can be applied to real-time analytics, transaction processing, and data caching. Implementing in-memory databases or processing frameworks can lead to substantial improvements in latency and overall system performance.

3.6. Optimization Algorithms Optimization algorithms are designed to find the most efficient solutions to complex problems, often by minimizing computation time and resource usage. Techniques such as heuristic algorithms, dynamic programming, and integer programming can be applied to optimize financial models, trading strategies, and resource management. By incorporating optimization algorithms into financial applications, institutions can achieve faster processing and reduced latency for decision-making and strategy execution.

3.7. GPU Acceleration Graphics Processing Units (GPUs) are specialized processors designed to handle parallel processing tasks efficiently. GPU acceleration leverages the parallel processing capabilities of GPUs to accelerate data-intensive computations, such as machine learning algorithms and simulations. In financial applications, GPUs can be used to enhance the performance of tasks like algorithmic trading, real-time data analysis, and risk modeling. By utilizing GPUs, financial institutions can achieve faster processing times and reduced latency for complex computations.

By adopting and implementing these computational strategies, financial institutions can effectively address latency challenges and enhance the performance of their distributed cloud systems. Each strategy offers unique benefits and can be tailored to meet the specific requirements of financial applications, leading to improved efficiency, faster response times, and overall better service quality.

4. DATA MANAGEMENT TECHNIQUES

Effective data management is essential for reducing latency in distributed cloud systems, especially for financial applications that handle vast amounts of data and require real-time processing. Proper management of data can streamline access, improve processing speed, and ensure that information is available when needed. This section explores key data management techniques that contribute to latency reduction in financial services.

4.1. Data Caching

Data caching involves storing copies of frequently accessed data in memory or on fast-access storage systems to reduce retrieval times. By caching data such as market feeds, historical trading data, and user queries, financial applications can avoid repeated database lookups and minimize latency. Implementing efficient caching mechanisms, such as in-memory caches (e.g., Redis, Memcached) or distributed caches, can significantly enhance the performance of real-time analytics and transaction processing systems.

4.2. Data Compression

Data compression techniques reduce the size of data transmitted over networks or stored in databases, thereby decreasing transmission times and storage requirements. For financial applications, data compression can be applied to transaction logs, market data feeds, and analytical results. Employing efficient compression algorithms, such as Lempel-Ziv (LZ) or Huffman coding, can help lower latency by reducing the volume of data that needs to be processed or transmitted.

4.3. Distributed Databases

Distributed databases spread data across multiple nodes or servers to enhance scalability and performance. Techniques such as data sharding, partitioning, and replication can optimize data retrieval and management. For financial applications, distributed databases can handle large volumes of transaction data, improve availability, and ensure faster access to critical information. By balancing the data load across multiple nodes, distributed databases help reduce latency and improve overall system responsiveness.

4.4. Data Prefetching

Data prefetching involves anticipating the data that will be needed in the near future and loading it into memory before it is requested. This technique helps minimize delays by ensuring that the required data is readily available. In financial applications, data prefetching can be used for market data feeds, predictive analytics, and transaction processing. By proactively loading data, financial systems can achieve faster response times and reduce latency.

4.5. Real-Time Data Processing

Real-time data processing involves analyzing and acting on data as it is generated, rather than storing and processing it in batches. Techniques such as stream processing and complex event processing (CEP) enable financial applications to handle real-time data flows, such as trading transactions and market events. By processing data in real-time, financial institutions can achieve quicker insights and decisions, reducing latency and improving system performance.

4.6. Data Replication and Synchronization

Data replication involves creating copies of data across multiple locations to enhance availability and reliability. Synchronization ensures that these copies are consistently updated. For financial applications, data replication and synchronization can be used to maintain consistency across distributed systems, ensuring that all nodes have access to the most current data. This approach helps reduce latency by preventing delays caused by data inconsistencies or retrieval from distant locations.

4.7. Efficient Data Storage Solutions

Choosing the right data storage solutions is crucial for managing latency. Technologies such as solid-state drives (SSDs), high-performance storage arrays, and distributed file systems can provide faster data access and retrieval. For financial applications, utilizing efficient storage solutions can help reduce latency associated with data I/O operations and improve overall system performance.

4.8. Data Governance and Management Policies
Implementing robust data governance and management policies ensures that data is organized, maintained, and accessed efficiently. Policies related to data quality, integrity, and access control can help streamline data management processes and reduce latency. Establishing clear guidelines for data handling, storage, and retrieval can contribute to improved system performance and reduced delays.

By employing these data management techniques, financial institutions can effectively manage and optimize data to reduce latency in distributed cloud systems. Each technique offers distinct benefits and can be tailored to meet the specific needs of financial applications, leading to enhanced performance, faster response times, and improved service quality.

5. ARCHITECTURAL IMPROVEMENTS

Architectural improvements in distributed cloud systems are crucial for reducing latency and enhancing the performance of financial applications. By optimizing system architecture, financial institutions can streamline data processing, improve resource utilization, and ensure efficient communication between components. This section explores several architectural improvements that can effectively minimize latency in financial services.

5.1. Microservices Architecture
Microservices architecture involves breaking down a monolithic application into smaller, independent services that can be developed, deployed, and scaled separately. Each microservice handles a specific function or feature, allowing for greater flexibility and efficiency. In financial applications, microservices can improve latency by enabling parallel processing of different components, facilitating faster updates, and reducing the impact of failures. This modular approach ensures that individual services can be optimized for performance, leading to overall latency reduction.

5.2. Serverless Computing
Serverless computing abstracts the underlying infrastructure management, allowing developers to focus on writing and deploying code. In a serverless model, functions are executed in response to events, and resources are automatically allocated as needed. For financial applications, serverless computing can reduce latency by providing rapid, on-demand execution of functions such as transaction processing, real-time data analysis, and fraud detection. This approach eliminates the need for manual scaling and provisioning, leading to more efficient and responsive systems.

5.3. Hybrid Cloud Solutions
Hybrid cloud solutions combine private and public cloud environments to optimize performance, flexibility, and cost-efficiency. By leveraging both on-premises and cloud resources, financial institutions can balance latency requirements with resource availability. For example, sensitive data and critical workloads can be managed in a private cloud for security and performance, while less critical tasks can be handled in a public cloud. This hybrid approach allows for better management of latency and scalability based on specific application needs.

5.4. Edge Computing Integration
Integrating edge computing into the architecture involves placing computational resources closer to the data source, such as at network edge locations. This reduces the distance data must travel, minimizing latency. In financial applications, edge computing can be used to process real-time market data, execute trading algorithms, and handle high-frequency trading tasks at the edge of the network. By reducing the reliance on centralized data centers, edge computing can significantly improve responsiveness and reduce latency.

5.5. High-Performance Networking
High-performance networking solutions, such as low-latency network fabrics and high-speed interconnects, can enhance communication between distributed components. Technologies like InfiniBand or high-speed Ethernet can reduce network delays and improve data transfer rates. In financial systems, implementing high-performance networking can minimize latency associated with data transmission and improve the efficiency of transactions and data processing.

5.6. Decentralized Data Management
Decentralized data management involves distributing data across multiple nodes or locations rather than relying on a centralized database. This approach can reduce latency by optimizing data access and reducing the load on individual nodes. Techniques such as data partitioning, sharding, and replication help manage large datasets efficiently and ensure that data is readily accessible. For financial applications, decentralized data management can enhance performance and reduce delays associated with data retrieval and processing.

5.7. Load Distribution and Traffic Management
Implementing effective load distribution and traffic management strategies ensures that network and computational resources are utilized efficiently.

Techniques such as load balancers, traffic shaping, and QoS (Quality of Service) policies can optimize the allocation of resources and prioritize critical traffic. In financial applications, managing load distribution and traffic helps prevent bottlenecks and ensures that latency-sensitive operations, such as trade executions and data analysis, are performed efficiently.

5.8. Fault Tolerance and Redundancy Designing systems with fault tolerance and redundancy capabilities ensures that operations continue smoothly even in the event of failures or disruptions. Techniques such as failover mechanisms, redundant components, and data replication can enhance system reliability and reduce latency caused by downtime or failures. For financial applications, implementing fault tolerance and redundancy helps maintain performance and responsiveness, ensuring continuous service availability.

By incorporating these architectural improvements, financial institutions can enhance the performance of their distributed cloud systems and achieve significant reductions in latency. Each improvement offers specific advantages and can be tailored to address the unique requirements of financial applications, leading to more efficient and responsive systems.

6. CHALLENGES AND FUTURE DIRECTIONS

While significant progress has been made in reducing latency within distributed cloud systems for financial applications, several challenges remain. Addressing these challenges and exploring future directions is crucial for advancing the effectiveness and efficiency of latency reduction techniques.

1. Challenges

Scalability Issues As financial applications grow in complexity and scale, managing latency becomes increasingly challenging. Scaling distributed cloud systems to handle larger volumes of data and more complex computations can introduce new latency issues. Ensuring that latency reduction techniques remain effective as systems scale requires continuous optimization and adaptation.

Network Congestion and Bottlenecks Network congestion and bottlenecks can severely impact latency, particularly in distributed cloud environments where data is transmitted across multiple nodes and networks. Addressing these issues involves optimizing network infrastructure and managing traffic effectively, but congestion can still occur during peak usage times or due to unforeseen network conditions.

Data Consistency and Integrity Maintaining data consistency and integrity while reducing latency can be challenging, especially in distributed systems where data is replicated and partitioned across multiple locations. Ensuring that all data remains accurate and synchronized without introducing additional latency requires careful management and robust consistency mechanisms.

Security and Privacy Concerns Reducing latency often involves implementing technologies and techniques that may impact security and privacy. For example, edge computing and data caching may expose sensitive financial data to additional risks. Balancing latency reduction with maintaining strong security and privacy measures is an ongoing challenge.

Integration of Emerging Technologies Integrating emerging technologies, such as quantum computing and advanced AI, into existing distributed cloud systems poses challenges related to compatibility, performance, and latency. Ensuring that new technologies seamlessly integrate with current systems while effectively reducing latency requires innovative approaches and thorough testing.

2. Future Directions

Advanced AI and Machine Learning Future advancements in AI and machine learning can further enhance latency reduction by enabling predictive models that optimize system performance and resource allocation. AI-driven algorithms can analyze historical data to predict and mitigate potential latency issues, leading to more proactive and efficient management of distributed cloud systems.

Quantum Computing Quantum computing holds the potential to revolutionize computational capabilities and reduce latency in financial applications. By leveraging quantum algorithms for complex calculations and optimizations, future systems may achieve unprecedented performance improvements. Research and development in quantum computing will be crucial for realizing these benefits.

5G and Beyond The rollout of 5G networks and future advancements in wireless communication technologies promise to further reduce latency by providing higher speeds and lower latency connections. Leveraging these technologies can enhance the performance of distributed cloud systems and improve real-time data processing for financial applications.

Edge-AI Integration Integrating AI capabilities at the edge of the network can complement edge computing strategies by enabling real-time decision-making and data processing. Edge-AI solutions can enhance latency reduction by processing data locally and applying AI algorithms to optimize performance and resource usage.

Improved Data Management Techniques Future research in data management techniques, such as advanced data compression methods, more efficient data partitioning strategies, and innovative caching solutions, can contribute to further reducing latency. Developing new approaches to data management will be essential for addressing the growing demands of financial applications.

Enhanced Network Technologies Continued advancements in network technologies, including high-speed networking and network optimization techniques, will play a critical role in reducing latency. Future developments in network infrastructure, protocols, and management tools will help address current limitations and improve overall system performance.

Hybrid and Multi-Cloud Architectures Exploring hybrid and multi-cloud architectures can offer additional flexibility and optimization opportunities for reducing latency. Future research will focus on refining these architectures to balance workloads, manage data efficiently, and ensure low-latency performance across diverse cloud environments.
By addressing these challenges and pursuing future directions, the field of latency reduction in distributed cloud systems can continue to advance, providing even greater benefits for financial applications and ensuring that systems remain performant and competitive in an evolving technological landscape.

7. CONCLUSION

Latency reduction in distributed cloud systems is a critical factor for the success of financial applications that demand real-time processing and rapid decision-making. As the financial industry continues to rely on cloud-based infrastructure for tasks such as high-frequency trading, risk management, and real-time data analysis, minimizing delays becomes paramount.

This paper has reviewed various techniques across network optimization, computational strategies, data management, and architectural improvements. Each approach offers unique advantages in reducing latency, ranging from edge computing and content delivery networks to in-memory computing, microservices architecture, and load balancing. However, implementing these solutions presents challenges, such as integration complexity, security concerns, scalability, and compliance with evolving regulations.

To overcome these hurdles, financial institutions must invest in scalable, secure, and adaptable systems while exploring future directions like leveraging emerging technologies such as quantum computing, 5G, and serverless architectures. Additionally, striking the right balance between cost-efficiency and performance optimization is essential for achieving sustainable latency reduction.

In conclusion, addressing latency in distributed cloud systems is not a one-size-fits-all solution but requires a multi-faceted approach tailored to the specific needs of financial applications. By strategically adopting these techniques and preparing for future innovations, financial institutions can maintain competitive advantages and meet the growing demands for speed, efficiency, and responsiveness in the financial markets.

REFERENCES

1. A. R. Kunduru, "THE PERILS AND DEFENSES OF ENTERPRISE CLOUD COMPUTING: A COMPREHENSIVE REVIEW," Cent. Asian J. Math. Theory Comput. Sci., vol. 4, no. 9, pp. 29–41, 2023.
2. Zebari, I. M., Zeebaree, S. R., & Yasin, H. M. (2019, April). Real time video streaming from multi-source using client-server for video distribution. In 2019 4th Scientific International Conference Najaf (SICN) (pp. 109-114). IEEE.
3. F. K. Parast, C. Sindhav, S. Nikam, H. I. Yekta, K. B. Kent, and S. Hakak, "Cloud computing security: A survey of service-based models," Comput. Secur., vol. 114, p. 102580, 2022.
4. R. Kollolu, "Infrastructural constraints of Cloud computing," Int. J. Manag. Technol. Eng., vol. 10, pp. 255–260, 2020.
5. Zeebaree, S., & Zebari, I. (2014). Multilevel client/server peer-to-peer video broadcasting system. International Journal of Scientific & Engineering Research, 5(8), 260-265.
6. Card SK, Robertson GG, Mackinlay JD. The information visualizer, an information workspace. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. CHI'91. New York, NY, USA: ACM; 1991:181e186
7. Michael V, Karen M. The Case for Big Data in the Financial Services Industry. 2012

8. Chang F, Dean J, Ghemawat S, et al. Bigtable: a distributed storage system for structured data. In: Proceedings of the 7th Symposium on Operating Systems Design and Implementation. OSDI'06. Berkeley, CA, USA: USENIX Association; 2006:205e218. http://dl.acm.org/ citation.cfm?id¼1298455.1298475.
9. Dean J, Ghemawat S. Mapreduce: simplified data processing on large clusters. In: Proceedings of the 6th Conference on Symposium on Opearting Systems Design & Implementation e Volume 6. OSDI'04. Berkeley, CA, USA: USENIX Association; 2004, 10e10 http://dl.acm. org/citation.cfm?id¼1251254.1251264.
10. Hindman B, Konwinski A, Zaharia M, et al. Mesos: a platform for fine-grained resource sharing in the data center. In: Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation. NSDI'11. Berkeley, CA, USA: USENIX Association; 2011:295e308
11. Alizadeh M, Kabbani A, Edsall T, Prabhakar B, Vahdat A, Yasuda M. Less is more: trading a little bandwidth for ultra-low latency in the data center. In: Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation. NSDI'12. Berkeley, CA, USA: USENIX Association; 2012.
12. Soltesz S, P€otzl H, Fiuczynski ME, Bavier A, Peterson L. Container-based operating system virtualization: a scalable, high-performance alternative to hypervisors. SIGOPS Oper Syst Rev. Mar. 2007;41
13. Mensi W, Hammoudeh S, Reboredo JC, Nguyen DK. Do global factors impact BRICS stock markets? A quantile regression approach. Emerg Mark Rev. 2014;19:1e17.
14. Lehkonen H, Heimonen K. Timescale-dependent stock market comovement: BRICs vs. developed markets. J Empir Finance. 2014;28:90e103
15. Waal A, Eyden R. Monetary policy and inflation in South Africa: a VECM augmented with foreign variables. South Afr J Econ. 2014;82(1):117e140.
16. Hong J, Mattmann C, Ramirez P. Ensemble Maximum Entropy Classification and Linear Regression for Author Age Prediction. CoRR; October 2016. https://arxiv.org/abs/1610.00852. Accessed September 16, 2017.