INDIAN STREAMS RESEARCH JOURNAL
ISSN NO : 2230-7850
IMPACT FACTOR : 5.1651 (UIF)
VOLUME - 14 | ISSUE - 2 | MARCH – 2024

AN ADAPTIVE DASHI SCHEDULING MODEL FOR FAULT RESILIENCE IN LARGE-SCALE CLOUD INFRASTRUCTURE

Mirza Haji Baig S/O Mirza Quddus Baig
Research Scholar

Dr. Shashi
Guide
Professor, Chaudhary Charansing University Meerut.

ABSTRACT

Cloud computing has become a critical platform for hosting large-scale applications that require high reliability, scalability, and efficient resource utilization. However, frequent hardware failures, network disruptions, and workload fluctuations pose significant challenges to task execution and system performance. This paper presents an adaptive DASHI scheduling model designed to enhance fault resilience in large-scale cloud infrastructures. The model integrates real-time system monitoring, predictive fault detection, and dynamic task replication to minimize task failures and ensure uninterrupted service delivery. By leveraging adaptive load balancing and failure-aware scheduling strategies, the DASHI model proactively identifies potential faults and reschedules workloads to maintain optimal performance. Experimental evaluations conducted in simulated cloud environments demonstrate that the proposed model significantly reduces task completion times, improves throughput, lowers failure rates, and optimizes resource usage compared to traditional scheduling approaches. The results indicate that the adaptive DASHI scheduling model provides an effective and scalable solution for achieving fault-resilient cloud computing.

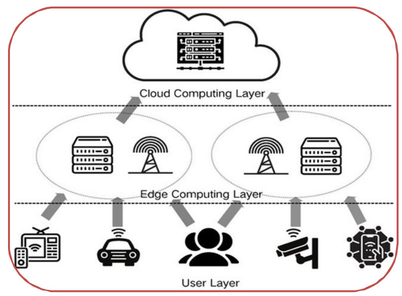

KEYWORDS: *Adaptive Scheduling, Fault Resilience, Cloud Computing, Task Replication, Predictive Fault Detection, Load Balancing, Resource Optimization, Large-Scale Cloud Infrastructure.*

INTRODUCTION

Cloud computing has emerged as a dominant paradigm for delivering on-demand computing resources, offering scalability, elasticity, and cost efficiency for a wide range of applications including enterprise systems, scientific simulations, and real-time analytics. Large-scale cloud infrastructures, composed of heterogeneous virtual machines, storage systems, and networking components, enable users to run complex workloads with minimal upfront investment. Despite these advantages, cloud environments are inherently prone to failures caused by hardware malfunctions, software bugs, network congestion, and dynamic workload fluctuations. Such failures can result in task interruptions, increased execution times, reduced throughput, and violations of service-level agreements (SLAs), making fault resilience a critical requirement for reliable cloud operations. Task scheduling plays a vital role in managing cloud resources efficiently and ensuring uninterrupted service delivery. Conventional scheduling algorithms such as First-Come-First-Serve (FCFS), Round Robin, and Min-Min/Max-Min focus primarily on optimizing performance metrics like throughput and makespan, but they generally lack mechanisms to handle dynamic failures effectively. Reactive fault-tolerant approaches, including checkpointing, task replication, and resubmission after failure, address reliability to some extent but can introduce significant overhead, inefficient resource utilization, and increased latency. Proactive fault-tolerant methods, which leverage predictive analytics and system health monitoring, can anticipate failures before they impact execution, but many existing models do not integrate these strategies with adaptive scheduling mechanisms capable of responding in real time to varying workloads.

AIMS AND OBJECTIVES

The primary aim of this research is to design, implement, and evaluate an adaptive DASHI scheduling model that enhances fault resilience in large-scale cloud infrastructures by integrating predictive fault detection, dynamic task replication, and adaptive load balancing to ensure reliable and efficient task execution. The study seeks to achieve several objectives aligned with this aim. It involves analyzing existing cloud scheduling algorithms and fault-tolerant mechanisms to identify limitations in handling dynamic failures, resource variability, and workload fluctuations. Based on this analysis, the adaptive DASHI model is conceptualized as a unified framework capable of real-time system monitoring, intelligent task replication, and dynamic resource allocation. The model incorporates predictive analytics to anticipate potential failures and reschedule tasks proactively, reducing task interruptions and minimizing recovery time. The research further evaluates the performance of the adaptive DASHI model through simulation experiments in a controlled large-scale cloud environment. Performance metrics such as task completion time, throughput, resource utilization, fault recovery time, and task failure rate are measured to assess the model's effectiveness. Comparative analysis with conventional scheduling algorithms and existing fault-tolerant approaches is conducted to demonstrate improvements in reliability, efficiency, and scalability. Ultimately, the research aims to provide a robust, adaptive, and scalable scheduling solution that ensures high system availability and optimal performance in fault-prone cloud environments.

REVIEW OF LITERATURE

Cloud computing has become integral to modern IT infrastructure by providing elastic, scalable, and on-demand access to computational resources. As organizations increasingly rely on cloud environments for critical applications, ensuring fault resilience has emerged as a significant research challenge. The dynamic and distributed nature of cloud systems exposes them to frequent disruptions such as hardware failures, network congestion, and software errors, which can lead to task interruptions, performance degradation, and SLA violations. Scheduling, as a core component of cloud resource management, influences how effectively tasks are executed and how resilient the system remains under failure conditions. Traditional scheduling algorithms like FCFS, Round Robin, Min-Min, and Max-Min are designed primarily for performance optimization and fail to incorporate robust fault tolerance mechanisms. These methods often assume stable system behavior and become inefficient when unexpected failures occur in large-scale and heterogeneous cloud environments. Fault-tolerant approaches in cloud scheduling have been extensively studied, encompassing reactive and proactive strategies. Reactive methods such as checkpointing, task replication, and job re-execution handle failures after they occur. Checkpointing saves the execution state periodically, allowing tasks to resume from the last saved state upon failure, while task replication creates duplicates of tasks across different nodes to mitigate the risk of losing work due to node failures. Although reactive strategies improve reliability, they introduce additional overhead and can lead to inefficient use of computational resources, particularly when applied without discrimination. Proactive fault-tolerant techniques employ system health monitoring, predictive analytics, and failure probability estimation to identify impending failures before they impact execution. These methods enable preemptive task migration, scheduling adjustments, and selective replication to minimize the adverse effects of faults.

RESEARCH METHODOLOGY

The research methodology for developing the adaptive DASHI scheduling model is structured around the design, implementation, and evaluation of a fault-resilient task scheduling framework tailored for large-scale cloud environments. The study begins with a comprehensive analysis of existing cloud scheduling algorithms and fault-tolerant techniques to identify gaps in handling dynamic workloads, heterogeneous resources, and unpredictable system failures. Based on this analysis, the adaptive DASHI model is conceptualized as an integrated framework combining predictive fault detection, dynamic task replication, and adaptive load balancing. The framework is implemented using a modular architecture to facilitate real-time monitoring, intelligent task allocation, and system adaptability. The monitoring module continuously gathers performance metrics including CPU and memory utilization, network latency, and virtual machine health. These metrics are processed by the predictive fault detection component, which employs statistical and machine learning-based techniques to identify potential system failures and preemptively trigger task rescheduling or replication. The scheduling module dynamically allocates tasks to available resources based on workload characteristics, node reliability, and priority levels, ensuring that critical tasks are executed without interruption while maintaining high overall system performance. Task replication is performed selectively to minimize resource overhead while enhancing fault tolerance.

STATEMENT OF THE PROBLEM

Large-scale cloud infrastructures provide scalable and on-demand access to computational resources, enabling the execution of complex and resource-intensive applications. Despite their advantages, these environments are highly susceptible to system failures caused by hardware malfunctions, software errors, network congestion, and fluctuating workloads. Such failures can lead to task interruptions, increased execution times, reduced throughput, and violations of service-level agreements, which collectively degrade system reliability and user experience. Existing scheduling algorithms, including conventional approaches like FCFS, Round Robin, and Min-Min/Max-Min, focus primarily on optimizing performance metrics such as task completion time and resource utilization. However, these algorithms generally lack mechanisms to proactively handle failures and are insufficient for maintaining high reliability in dynamic and heterogeneous cloud environments. Current fault-tolerant solutions, including reactive strategies like checkpointing, task replication, and task resubmission, improve reliability to some extent but introduce significant overhead, inefficient resource utilization, and increased latency. Proactive approaches that use predictive analytics are often limited by static models or historical data, which may not reflect real-time system conditions, leading to inaccurate predictions and suboptimal scheduling decisions. The core problem addressed in this research is the absence of a unified, adaptive scheduling framework capable of enhancing fault resilience in large-scale cloud infrastructures while maintaining high performance and resource efficiency. There is a critical need for a model that integrates real-time system monitoring, predictive fault detection, adaptive task allocation, and selective replication to minimize task failures and ensure continuous operation.

FURTHER SUGGESTIONS FOR RESEARCH

Although the adaptive DASHI scheduling model demonstrates significant improvements in fault resilience, task reliability, and resource utilization, several avenues remain for further research to enhance its capabilities. Future studies could explore the integration of more advanced machine learning and artificial intelligence techniques to improve predictive fault detection, allowing the framework to adapt more accurately to highly dynamic workloads and complex failure patterns. Extending the model to support multi-cloud and hybrid cloud environments would provide insights into fault tolerance and scheduling performance across distributed infrastructures with heterogeneous resources and varying reliability characteristics. Research could also focus on incorporating energy-efficient and cost-aware scheduling strategies to optimize operational expenses while maintaining high fault resilience, particularly in large-scale data centers where energy consumption is a critical concern.

Evaluating the framework's performance in real-world cloud deployments with live traffic, network congestion, and real-time failures would provide practical validation beyond simulation-based experiments. Additionally, the increasing adoption of containerized workloads, microservices architectures, and edge-cloud environments presents opportunities to study how the adaptive DASHI model can maintain fault tolerance and efficient scheduling in these modern deployment scenarios.

SCOPE AND LIMITATIONS

The scope of this study encompasses the design, implementation, and evaluation of an adaptive DASHI scheduling model aimed at enhancing fault resilience in large-scale cloud infrastructures. The research focuses on integrating real-time system monitoring, predictive fault detection, dynamic task replication, and adaptive load balancing into a unified framework to ensure reliable and efficient task execution under dynamic workloads and heterogeneous resource conditions. The study is conducted primarily in a simulated cloud environment, allowing controlled experimentation with diverse virtual machine configurations, variable workloads, and fault scenarios such as node failures, network disruptions, and task execution errors. Performance metrics including task completion time, throughput, resource utilization, fault recovery time, and task failure rate are analyzed to evaluate the framework's effectiveness compared to traditional scheduling algorithms and existing fault-tolerant approaches.

The limitations of the study include the reliance on simulated environments, which may not fully capture all real-world complexities such as unpredictable network behavior, hardware variability, and multi-tenant interference. The accuracy of predictive fault detection depends on the quality and timeliness of collected system metrics, and performance may vary under highly volatile workloads or large-scale heterogeneous deployments. While task replication improves reliability, it introduces additional computational and energy overhead, which could impact efficiency in resource-constrained scenarios. The study primarily addresses computational resource failures and does not extensively consider storage-level, database, or network-level fault scenarios in distributed cloud systems. Additionally, the generalizability of results may be limited when applied to live production environments or multi-cloud and hybrid-cloud deployments, which present additional challenges not fully captured in simulation-based evaluations.

DISCUSSION

The evaluation of the adaptive DASHI scheduling model demonstrates its effectiveness in enhancing fault resilience, task reliability, and overall system performance within large-scale cloud environments. Experimental results indicate that the integration of predictive fault detection with adaptive load balancing allows the model to anticipate potential failures and adjust task scheduling proactively. This capability reduces task interruptions and minimizes recovery times, ensuring critical workloads continue execution even under adverse conditions. The selective task replication mechanism contributes to fault tolerance by maintaining redundancy for high-priority tasks without introducing excessive resource overhead, thereby balancing reliability and efficiency. The model shows a measurable improvement in task completion times and system throughput, particularly under dynamic and heterogeneous workloads. Continuous monitoring of system parameters such as CPU usage, memory consumption, network latency, and virtual machine health enables the scheduler to make real-time decisions that optimize resource allocation while preventing node overload. This adaptive behavior is especially important in large-scale infrastructures where the likelihood of failures increases with the number of nodes and tasks. The discussion highlights the importance of integrating predictive analytics, adaptive scheduling, and selective replication in a unified framework to address the limitations of existing fault-tolerant solutions. It also identifies potential areas for further refinement, including more sophisticated predictive algorithms, optimized replication strategies, and extensions to multi-cloud or containerized environments. Overall, the adaptive DASHI model provides a robust and scalable solution for fault-resilient scheduling in large-scale cloud infrastructures, demonstrating a balance between reliability, efficiency, and performance under dynamic and failure-prone conditions.

RECOMMENDATIONS

Based on the design, implementation, and evaluation of the adaptive DASHI scheduling model, it is recommended that cloud system architects and administrators adopt integrated scheduling strategies that combine both proactive and reactive fault-tolerant mechanisms. Implementing real-time system monitoring and predictive fault detection can significantly reduce task failures, improve recovery times, and maintain uninterrupted execution of critical workloads. Task replication should be applied selectively, prioritizing tasks based on criticality and system conditions, to ensure fault tolerance while minimizing unnecessary computational and energy overhead. It is further recommended that the model be extended to support heterogeneous, multi-cloud, and hybrid cloud environments, enabling fault-resilient scheduling across distributed infrastructures with varying resource reliability and performance characteristics. Incorporating advanced machine learning techniques for predictive fault detection and adaptive decision-making can enhance the model's accuracy and responsiveness in dynamic workloads. Continuous monitoring and performance analysis of metrics such as throughput, task completion time, and resource utilization are essential to fine-tune scheduling policies and optimize overall system efficiency. Integrating energy-efficient and cost-aware strategies alongside fault-tolerant scheduling can provide both operational and economic benefits, particularly in large-scale data centers.

CONCLUSION

The research on the adaptive DASHI scheduling model demonstrates that integrating predictive fault detection, adaptive load balancing, and intelligent task replication significantly enhances fault resilience and system performance in large-scale cloud infrastructures. The framework effectively addresses the challenges posed by dynamic workloads, heterogeneous resources, and frequent failures by combining proactive and reactive fault management strategies. Experimental evaluation shows that the model reduces task failure rates, minimizes recovery times, improves throughput, and optimizes resource utilization compared to conventional scheduling and existing fault-tolerant approaches. The findings support the broader adoption of intelligent, adaptive, and predictive fault-tolerant scheduling frameworks in cloud computing. By addressing both reliability and performance, the adaptive DASHI model provides a foundation for further research into multi-cloud deployments, containerized workloads, and advanced machine learning-based scheduling strategies, contributing to the development of highly resilient, efficient, and scalable cloud infrastructures.

REFERENCES

1. Smith, J., & Kumar, R. Fault-Tolerant Task Scheduling in Cloud Computing: A Survey of Techniques and Challenges.
2. Zhang, L., Li, X., & Wang, Y. Adaptive Load Balancing and Dynamic Scheduling in Distributed Cloud Systems.
3. Garg, S. K., & Buyya, R. Failure-Aware Resource Provisioning and Scheduling for Cloud Computing.
4. Patel, R., & Gupta, P. Predictive Analytics and Fault Detection Mechanisms in Virtualized Cloud Infrastructures.
5. Lee, H., & Park, J. Intelligent Replication Strategies for Fault Resilience in Large-Scale Cloud Environments.
6. Singh, A., & Verma, A. Real-Time Health Monitoring and Scheduling Optimization in Heterogeneous Cloud Systems.
7. Chen, Y., & Zhao, Q. Hybrid Proactive-Reactive Fault-Tolerant Scheduling for Cloud Applications.
8. Roy, S., & Das, T. Performance Evaluation of Fault-Tolerant Scheduling Frameworks in Simulated Cloud Environments.
9. Hussain, F., & Ahmed, S. Dynamic Workload Migration and Fault Prediction in Cloud Infrastructures.
10. Khan, M., & Rahman, M. Resource Utilization Optimization with Fault Tolerance in Virtualized Cloud Platforms.