

Efficient Resource Utilization in Kubernetes: A Review of Load Balancing Solutions

Indrani Vasireddy¹, Prathima Kandi², and SreeRamya Gandu³

¹ Associate Professor, Department of Computer Science and Engineering, Geethanjali College of Engineering, Hyderabad, India

^{2,3} Assistant Professor, Department of Computer Science and Engineering, Geethanjali College of Engineering, Hyderabad, India

Copyright © 2023 Made Indrani Vasireddy et al. This is an open-access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

ABSTRACT- Modern distributed systems face the challenge of efficiently distributing workloads across nodes to ensure optimal resource utilization, high availability, and performance. In this context, Kubernetes, an open-source container orchestration engine, plays a pivotal role in automating deployment, scaling, and management of containerized applications. This paper explores the landscape of load balancing strategies within Kubernetes, aiming to provide a comprehensive overview of existing techniques, challenges, and best practices. The paper delves into the dynamic nature of Kubernetes environments, where applications scale dynamically, and demand for resources fluctuates. We review various load balancing approaches, including those based on traffic, resource-aware algorithms, and affinity policies. Special attention is given to the unique characteristics of containerized workloads and their impact on load balancing decisions. In this paper the implications of load balancing on the scalability and performance of applications deployed in Kubernetes clusters. It explores the trade-offs between different strategies, considering factors such as response time, throughput, and the adaptability to varying workloads. As cloud-native architectures continue to evolve, understanding and addressing the intricacies of load balancing in dynamic container orchestration environments become increasingly crucial. In this paper we had consolidated the current state of knowledge on load balancing in Kubernetes, providing researchers and practitioners with valuable insights and a foundation for further advancements in the quest for efficient, scalable, and resilient distributed systems.

KEYWORDS- Kubernetes, Cloud computing, Containers, Load Balancing.

I. INTRODUCTION

In the dynamic landscape of application deployment and resource management, Kubernetes has emerged as a versatile and open-source platform, offering scalability, portability, and an extensive ecosystem. This platform plays a pivotal role in efficiently managing the resources that applications and content utilize. In the early days of the internet, organizations encountered challenges with resource allocation[1] when using portable servers to run applications. The traditional approach of deploying

individual apps on separate physical servers proved to be costly and lacked effective measurement capabilities. As the scope of cloud services continues to expand, the imperative to enhance the performance of data center infrastructure becomes increasingly pivotal.

High-performance computing[3], advanced networking solutions, and resource optimization strategies are essential components in sustaining the speed and efficiency required for delivering high-quality cloud services. An impactful optimization strategy is the adoption of containerized applications, leveraging benefits such as enhanced portability, improved security, efficient resource utilization, rapid deployment and scaling, and heightened integration and interoperability.

Kubernetes, a robust container orchestration system, plays a central role in automating[15] the deployment, scaling, and management of containerized applications. A key feature of Kubernetes is its sophisticated scheduling mechanism, orchestrating the placement of containers across a cluster of nodes through a scheduling algorithm. This algorithm critically influences the efficiency and performance of the entire system. Containers, touted for their efficiency, particularly shine in comparison to virtual machines as they offer straightforward sharing among multiple programs. This versatility extends to running programs like Kubernetes on diverse distributed operating systems and cloud platforms. Within the Kubernetes paradigm, the concept of a "pod" takes center stage—a collection of containers with specific purposes. Further, organizational cohesion is achieved through "services," which are groups of interconnected pods that collectively perform the same set of functions. Notably, Kubernetes pods[7] are designed to be ephemeral, allowing Kubernetes to autonomously create and destroy them based on dynamic needs without requiring user intervention. Kubernetes emerges as a robust platform capable of swiftly scaling to meet varying intensities. Whether faced with minimal visitors or experiencing a surge in popularity, Kubernetes can efficiently scale down or up, adjusting the number of pods to optimize resource utilization and cost-effectiveness.

Widely embraced across diverse organizations, Kubernetes has become a trusted solution for managing and deploying containerized applications. This paper acknowledges its pivotal role in the evolving landscape of IT infrastructure, particularly in the context of server-less computing. While presenting opportunities for end-to-end

security enhancements, the transition to server-less computing also introduces new challenges in infrastructure and security administration. As the computing paradigm shifts towards edge and fog computing, Kubernetes demonstrates its versatility in providing seamless network management between cloud and edge nodes. However, deploying Kubernetes in IOT environments poses unique challenges, including optimizing network traffic distribution, flow routing policies, and distribution of computational resources on edge devices. Zhong et al.[2] Propose a heterogeneous task allocation strategy for cost-efficient container orchestration, revealing a optimal cost reduction compared to the default Kubernetes framework. Introduce Kube-Knots, a GPU-aware resource orchestration layer integrated with Kubernetes, enhancing GPU usage in HPC workloads[14] and reducing energy consumption. In light of the diverse applications and challenges associated with Kubernetes, this paper embarks on a comprehensive survey of scheduling algorithms. Motivated by the increasing importance of Kubernetes in managing containerized applications, the survey aims to critically evaluate the existing literature, categorizing it into four sub-categories: generic scheduling, multi-criteria optimization-based scheduling, AI-focused scheduling, and Dynamic scaling enabled scheduling[13]. The taxonomy employed in this survey offers a fine-grained and novel perspective, distinct from existing surveys that primarily focus on container orchestration. The contributions of this paper encompass a thorough review of literature, providing insights into each sub-category of Kubernetes scheduling algorithms. The critical evaluation of existing approaches highlights their strengths and limitations, paving the way for identifying gaps and open questions in the field. Ultimately, the goal is to contribute to the advancement of knowledge in Kubernetes scheduling algorithms, offering a valuable resource for researchers and practitioners alike.

This paper is organized as follows: Section 2 presents the related work, Section 3 presents the Discussions, Section 4 presents the Challenges Section 5 Future enhancement, and Section 6 conclusion.

II. RELATED RESEARCH

The domain of Kubernetes scheduling algorithms has become a focal point for both researchers and practitioners, drawing substantial attention in recent years. This re-view explores the multifaceted landscape of scheduling algorithms within Kubernetes, delving into a diverse array of theories, methodologies, and findings from past studies. In recent times, software system design has seen the rising adoption of microservices. Load balancing, a critical component in microservices architecture, significantly influences how well the system utilizes its resources and ensures efficient service functionality. The literature underscores the imperative need for efficient scheduling of workloads in Kubernetes environments. Traditional scheduling approaches often grapple with the intricate and dynamic nature of these workloads[12], prompting a surge of interest in advanced scheduling algorithms. This shift aims to enable the efficient allocation of computing resources within Kubernetes clusters. Another recurring theme is the

potential benefits of advanced scheduling algorithms for Kubernetes. Studies highlight their capacity to enhance resource utilization, reduce latency, and overall improve cluster performance. Furthermore, these algorithms hold promise in supporting the development of emerging applications such as real-time analytics and machine learning.

Amidst the promise, the literature identifies challenges and limitations. One pivotal challenge involves addressing the evolving nature of workloads and applications within the Kubernetes cluster. Authors have explored improving the Dynamic scaling feature in Kubernetes scheduling to dynamically adjust resource allocation based on current demand. Other challenges include managing and coordinating multiple scheduling algorithms while ensuring the stability and performance of the overall system. The Multi-Metric Dynamic Load Balancing strategy is a sophisticated approach designed for cloud computing environments, offering users the flexibility to customize load balancing based on a diverse set of metrics to address complex business requirements. The strategy comprises two essential components: a Metric Collector, responsible for gathering server and container metrics, and a Weighting Calculator, enabling users to assign weights to metrics based on their significance. This dynamic approach is implemented through server and container monitoring schemes, and its effectiveness is rigorously validated through experiments. The overarching aim is to optimize resource utilization, reduce response times, and enhance network throughput. This strategy aligns with the ambitions of cloud computing, emphasizing time and cost savings while ensuring ubiquitous service access. In the context of micro-services architecture, load balancing becomes pivotal, influencing resource utilization and service functionality. The Spring Cloud framework, a dominant player in the market, exemplifies the tools used in effective load balancing within micro-services. Overall, this Multi-Metric Dynamic Load Balancing strategy stands as a versatile solution with demonstrated effectiveness in contemporary cloud computing scenarios. Propose Stratus, a cluster scheduler for batch job execution in IaaS platforms, demonstrating cost reduction.

Load balancing plays a pivotal role in determining how efficiently a system utilizes its resources and how well individual services function within the larger architecture. The main ambition of cloud computing is to save time, money, and provide ubiquitous service access from any device, at any time, and anywhere. Load balancing is instrumental in achieving these ambitions by reducing response times and strategically allocating resources and network throughput. As cloud computing continues to be a driving force in the digital landscape, the need for sophisticated load balancing strategies becomes increasingly apparent.

III. DISCUSSIONS

Scheduling through multi-criteria optimization involves the utilization of advanced techniques to optimize various conflicting objectives simultaneously. This approach seeks to find optimal solutions that balance multiple criteria, rather than focusing on a single optimization

goal. In the context of scheduling, where tasks or jobs need allocation of resources, multi-criteria optimization aims to enhance efficiency while considering diverse and often competing factors. Unlike traditional scheduling methods that might prioritize a singular objective, multi-criteria optimization in scheduling addresses the complexity of conflicting goals. This can include optimizing for resource utilization, minimizing completion time, and balancing workloads. The key idea is to explore a set of solutions, known as the Pareto front, where no solution is superior in all objectives but offers trade-offs. In summary, scheduling through multi-criteria optimization provides a nuanced and comprehensive approach, considering a spectrum of objectives simultaneously, thus offering a more robust and flexible solution to scheduling challenges.

In realm of Artificial Intelligence-driven scheduling, many prominent companies have established machine/deep learning clusters, comprising tens to thousands of CPUs and GPUs, to train their deep learning models in a distributed fashion. This training process is both resource-intensive and time-consuming, necessitating efficient scheduling to optimize the utilization of these clusters and expedite model training. Various strategies have been employed for scheduling tasks in this context. While some adapt general-purpose schedulers for distributed deep learning tasks, they often statically allocate resources, leading to suboptimal resource utilization. Others propose dynamic resource allocation based on careful workload analysis. Optimus[4] is a customized job scheduler given by Peng et al. whose primary objective is to minimize the time required for deep learning training jobs. The approach of Optimus is Leveraging performance models, Optimus precisely estimates training speed and employs online fitting to dynamically organize tasks for reduced job completion time.

Decima is a machine learning technique proposed [5] by Mao et al whose primary objective is to develop efficient policies for scheduling data processing jobs on distributed compute clusters using modern machine learning techniques. Decima implements reinforcement learning (RL) and neural networks to acquire workload-specific scheduling algorithms. Gandivafair[6] is a deep learning training method proposed by Chaudhary et al whose primary objective is to introduce a distributed fair-share scheduler tailored for GPU clusters employed in deep learning training which Provides performance isolation between users, ensuring fair distribution of GPU time among active users.

Dynamic Scaling-Enabled Scheduling, is a Dynamic scaling technique and a pivotal feature in Kubernetes scheduling, facilitates the automatic adjustment of resources allocated to pods based on current demand. This dynamic adaptation enhances resource utilization efficiency, boosts performance, reduces costs, and ensures high application availability. The synergy between Dynamic scaling and scheduling is evident in their collaborative approach to handling changing workloads and optimizing resource usage within distributed systems. Taherizadeh et al., the authors introduce[8] a Dynamic Multi-Level (DM) Dynamic scaling method tailored for container-based cloud applications. Leveraging both infrastructure- and

application-level monitoring data, the DM method dynamically adjusts thresholds based on workload conditions to decide when to scale up or down. Comparative evaluations against seven existing Dynamic scaling methods reveal superior overall performance, particularly in response time and the number of instantiated containers. Implementation is carried out using the SWITCH system for time-critical cloud applications. In Rattihalli et al., the authors present RUBAS, [9] a resource management system designed for dynamic adjustment of container allocations within a Kubernetes cluster. RUBAS incorporates container migration to enhance the non-disruptive Kubernetes Vertical Pod Autoscaler (VPA) system. Toka et al. introduce a Kubernetes[11] scaling engine utilizing machine learning forecast methods to optimize Dynamic scaling decisions for cloud-based applications. The engine's short-term evaluation loop enables adaptation to changing request dynamics. The authors propose a compact management parameter allowing cloud tenants to easily set desired resource over-provisioning levels versus Service Level Agreement[10] (SLA) violations. Simulation and measurement results demonstrate fewer lost requests and slightly more provisioned resources compared to the default Kubernetes baseline.

IV. CHALLENGES

The extensive survey of research in the areas of multi-criteria optimization-based scheduling, Artificial Intelligence driven scheduling, and Dynamic Scaling-enabled scheduling in Kubernetes reveals a rich landscape of methodologies, algorithms, and challenges. As summarized, there are various optimization techniques, machine learning approaches, and synergies between Dynamic scaling and scheduling that contribute to the efficiency and performance of Kubernetes clusters.

The survey identifies several challenges and areas for future research, ranging from addressing the multi-objective nature of scheduling problems[12], adapting to dynamic cloud environments, to developing more effective and practical Dynamic scaling-enabled scheduling techniques. The need for real-world datasets, algorithm efficiency, and scalability, as well as the trade-offs between accuracy and computational complexity in artificial intelligence driven scheduling, are highlighted. On the flip side, the realm of artificial intelligence driven scheduling in Kubernetes has garnered substantial attention in recent years as researchers have proposed diverse strategies to optimize scheduling decisions through machine learning and other AI methodologies. A notable achievement in this domain is the formulation of scheduling algorithms adept at managing intricate workloads within dynamic environments. These algorithms take into account various factors, including resource availability, task dependencies, and application requirements, to make judicious scheduling decisions. Some studies have introduced reinforcement learning-based scheduling algorithms, allowing adaptation to evolving work-load patterns and learning from experience to enhance decision-making. Likewise, other studies have delved into deep learning-based approaches, enabling the capture of intricate patterns in workload data for precise predictions. Collectively, these investigations underscore

the potential of artificial intelligence driven scheduling to augment the efficiency and performance of Kubernetes clusters.

Nevertheless, several challenges persist in this field. A prominent obstacle is the dearth of real-world datasets for training and evaluating artificial intelligence driven scheduling algorithms. The reliance on synthetic or simulated datasets in most studies raises concerns about the applicability of developed algorithms to the complexities inherent in real-world workloads. Another hurdle involves striking a balance between scheduling accuracy and computational complexity. Future research endeavors in this domain may concentrate on the development of more efficient and scalable artificial intelligence driven scheduling algorithms capable of handling large-scale, real-world workloads. This trajectory could involve the exploration of novel machine learning and optimization techniques aimed at refining scheduling accuracy while concurrently alleviating computational complexities. The comprehensive survey provides a roadmap for future researchers, outlining potential challenges and opportunities for advancement in the field of Kubernetes scheduling.

Lastly, the evolving realm of Dynamic scaling-enabled scheduling represents a burgeoning area of research aimed at optimizing resource utilization and enhancing application performance by synergizing Dynamic scaling and scheduling techniques. Several recent research studies have delved into this domain, and a comprehensive analysis of these studies reveals notable improvements in resource utilization and application performance enabled by Dynamic scaling-enabled scheduling. These studies showcase the potential to curtail resource wastage, mitigate the risk of under-provisioning, and enhance application response times. Despite these promising outcomes, challenges persist in this area, with a primary focus on the intricate design of effective Dynamic scaling-enabled scheduling algorithms. Crafting algorithms capable of dynamically adapting to changing workloads while optimizing resource utilization and upholding application performance poses a non-trivial task. Furthermore, there is a pressing need for further research into the practical implementation of Dynamic scaling-enabled scheduling in real-world scenarios. Most existing studies have been confined to controlled experimental settings, warranting a thorough evaluation of the efficacy of Dynamic scaling-enabled scheduling in real-world applications. Addressing challenges related to algorithm design, standardization, and practical implementation should be the focal point of future research endeavors in this domain, with an emphasis on developing more effective and practical Dynamic scaling-enabled scheduling techniques.

The research papers employ a diverse array of algorithms to enhance Kubernetes scheduling, subjecting them to testing across various platforms and environments such as Spark, MXNet, Kubernetes, Google and TwoSigma's GPU cluster, Google compute, CPU-GPU setups, the National Cloud Infrastructure, CloudSim and Java, among others. Scenarios range from cloud infrastructure and user needs to real traces, simulations, and web traces. While some papers did not explicitly specify algorithmic details or the platforms and environments used, the survey provides a comprehensive analysis, constructing a

taxonomy that not only effectively appraises the current state-of-the-art but also delineates challenges and outlines potential future research directions.

V. FUTURE RESEARCH

As Kubernetes gains increasing popularity, the demand for advanced computation optimization techniques is expected to rise. The future trajectory of Kubernetes may involve the development of more sophisticated algorithms for workload scheduling and resource allocation, potentially incorporating AI or machine learning. Furthermore, synergizing Kubernetes with emerging technologies like serverless computing could enhance resource usage efficiency by enabling dynamic scaling without pre-provisioned infrastructure. The evolution of computation optimization in Kubernetes is likely to be a fusion of cutting-edge algorithms, innovative technologies, and continual advancements in cloud computing.

Research endeavors in Kubernetes are poised for exploration in various domains. Testing and implementation are crucial avenues for uncovering the limitations of current scheduling algorithms and paving the way for potential enhancements, particularly on large-scale clusters. Focus on refining tooling, automation, and testing frameworks within the Kubernetes ecosystem, alongside the ongoing improvement of implementation processes, documentation, and community collaboration, is vital. The future of testing and implementation in Kubernetes hinges on ongoing innovation, collaboration, and a steadfast commitment to advancing the platform. Several methods are employing learning algorithms for resource balancing inside and outside the cluster, yielding encouraging results. Despite this, there is room for the discovery of new learning algorithms to enhance the scheduler, particularly in the context of large-scale clusters. Assessing the scalability of the Kubernetes scheduler in more extensive clusters, identifying bottlenecks, and proposing solutions are essential steps. Specific contexts, such as Green Computing, present limitations and opportunities for improvement. The perpetual challenge of minimizing the carbon footprint of a cluster calls for the proposal of advanced schedulers aimed at reducing energy consumption and carbon footprint, especially in IIOT setups. This realm holds substantial potential for refining existing methods and proposing innovative approaches.

While much of the work in Kubernetes scheduling has been evaluated on small clusters, a future research direction involves expanding evaluations to larger cluster sizes. Assessing the scalability of the Kubernetes scheduler in more extensive clusters, identifying bottlenecks, and proposing solutions are essential steps. Evaluating the impact of larger cluster sizes on application performance and resource utilization could pave the way for more efficient scheduling algorithms and enhanced management strategies for large-scale Kubernetes deployments.

Scheduling in Kubernetes should extend beyond the static infrastructure viewpoint. Future research may focus on proposing advanced context-aware scheduling algorithms that consider a broader range of contextual factors, including user preferences, application dependencies, and

environmental conditions. This could involve exploring new machine learning techniques and optimization algorithms that dynamically adapt to changing conditions, prioritizing resources based on real-time feedback and analysis. Potential research areas also include developing new models and frameworks for managing resources in Kubernetes clusters, refining container orchestration and load balancing, and enhancing monitoring and analytics capabilities for more effective utilization of context-aware scheduling algorithms. While much of the work in Kubernetes scheduling has been evaluated on small clusters, a future research direction involves expanding evaluations to larger cluster sizes. Assessing the scalability of the Kubernetes scheduler in more extensive clusters, identifying bottlenecks, and proposing solutions are essential steps. Evaluating the impact of larger cluster sizes on application performance and resource utilization could pave the way for more efficient scheduling algorithms and enhanced management strategies for large-scale Kubernetes deployments.

The diverse array of future directions in Kubernetes research presents challenges of varying difficulty levels, offering exciting opportunities for future researchers to explore and address. This survey aims to guide and inspire future researchers in selecting challenges and solving new problems to contribute to the advancement of Kubernetes' state-of-the-art. In conclusion, this survey consolidates the current state of knowledge on load balancing in Kubernetes, providing researchers and practitioners with valuable insights and a foundation for further advancements in the quest for efficient, scalable, and resilient distributed systems.

VI. CONCLUSIONS

In summary, the examination of Kubernetes scheduling offers a thorough understanding of the present landscape in this domain. The survey encompasses the goals, approaches, algorithms, trial implementations, and outcomes of diverse research initiatives within this realm. Emphasizing the significance of scheduling in Kubernetes, the survey underscores the necessity for adept and efficient scheduling algorithms. The experimental findings indicate that opportunities for enhancement persist, prompting future endeavors to concentrate on innovating new algorithms and refining those already in existence. Ultimately, the survey delivers valuable perspectives into the present status of Kubernetes scheduling, guiding towards promising avenues for subsequent research.

CONFLICTS OF INTEREST

The authors declare that they have no conflicts of interest.

REFERENCES

- [1] Chang, C., Yang, S., Yeh, E., Lin, P. & Jeng, J. A Kubernetes-based monitoring platform for dynamic cloud resource provisioning. *GLOBECOM 2017-2017 IEEE Global Communications Conference*. pp. 1-6 (2017)
- [2] Zhong Z, Buyya R (2020) A Cost-Efficient Container Orchestration Strategy in Kubernetes-Based Cloud Computing Infrastructures with Heterogeneous Resources. *ACM Trans Internet Technol* 20(2):1–24
- [3] Kim SH, Kim T (2023) Local scheduling in kubeedge-based edge computing environment. *Sensors* 23(3):1522
- [4] Peng Y, Bao Y, Chen Y, Wu C, Guo C (2018) Optimus: An Efficient Dynamic Resource Scheduler for Deep Learning Clusters. *Proceedings of the 13th EuroSys Conference, EuroSys*
- [5] Mao H, Schwarzkopf M, Venkatakrishnan SB, Meng Z, Alizadeh M (2019) Learning scheduling algorithms for data processing clusters. *SIGCOMM Conference of the ACM Special Interest Group on Data Communication*. pp 270–288
- [6] Chaudhary S, Ramjee R, Sivathanu M, Kwatra N, Viswanatha S (2020) Balancing efficiency and fairness in heterogeneous GPU clusters for deep learning. *Proceedings of the 15th European Conference on Computer Systems, EuroSys*
- [7] Kubernetes: Available: <http://kubernetes.io/>.
- [8] Taherizadeh S, Stankovski V (2019) Dynamic multi-level auto-scaling rules for containerized applications. *Computer J* 62(2):174–197
- [9] Rattihalli G, Govindaraju M, Lu H, Tiwari D (2019) Exploring potential for non-disruptive vertical auto scaling and resource estimation in kubernetes. *IEEE International Conference on Cloud Computing, CLOUD*. pp 33–40
- [10] Jain, N., Mohan, V., Singhai, A., Chatterjee, D. & Daly, D. Kubernetes Load-balancing and related network functions using P4. *Proceedings Of The Symposium On Architectures For Networking And Communications Systems*. pp. 133-135 (2021)
- [11] Toka L, Dobreff G, Fodor B, Sonkoly B (2021) Machine Learning-Based Scaling Management for Kubernetes Edge Clusters. *IEEE Trans Netw Serv Manage* 18(1):958–972
- [12] Masne, S., Wankar, R., Raghavendra Rao, C. & Agarwal, A. Seamless provision of cloud services using peer-to-peer (p2p) architecture. *Distributed Computing And Internet Technology: 8th International Conference, ICDCIT 2012, Bhubaneswar, India, February 2-4, 2012. Proceedings* 8. pp. 257-258 (2012)
- [13] Kim SH, Kim T. Local Scheduling in KubeEdge-Based Edge Computing Environment. *Sensors (Basel)*. 2023 Jan 30;23(3):1522. doi: 10.3390/s23031522. PMID: 36772562; PMCID: PMC9921110.
- [14] Wankar, Rajeev. (2008). Grid Computing with Globus: An Overview and Research Challenges. *International Journal of Computer Science Applications*.
- [15] Vasireddy, Indrani, Rajeev Wankar, and Raghavendra Rao Chillarige. "Recreation of a Sub-pod for a Killed Pod with Optimized Containers in Kubernetes." *International Conference on Expert Clouds and Applications*. Singapore: Springer Nature Singapore, 2022.