

Resource Management Strategies in Heterogeneous Distributed Systems

Anila Gogineni*

Citation: Gogineni A. Chaos Engineering in the Cloud-Native Era: Evaluating Distributed AI Model Resilience on Kubernetes. *J Artif Intell Mach Learn & Data Sci* 2024, 2(2), 2182-2189. DOI: doi.org/10.51219/JAIMLD/anila-gogineni/478

Received: 02 September, 2024; **Accepted:** 18 September, 2024; **Published:** 20 September, 2024

***Corresponding author:** Anila Gogineni, Independent Researcher, USA, E-mail: anila.ssn@gmail.com

Copyright: © 2024 Gogineni A., This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

ABSTRACT

This research targets workload partitioning for GPUs, TPUs and CPUs in heterogeneous distributed systems. This paper focuses on the issues connected with the heterogeneity of the hardware, balanced loads, latency of data transfer and energy consumption. It highlights the characteristics of proactive scheduling algorithms for resources, which enhance the system's efficiency and capability. These applications are selected in the areas of machine learning and artificial intelligence, big data and cloud computing and scientific computing and modelling to demonstrate how systems with heterogeneity can improve certain computational workloads. Finally, the study provides information on fault tolerance and cost-effectiveness as a way of finding efficient and effective ways of managing available resources for mixed-distributed environments.

Keywords: Resource management, Heterogeneous systems, Load balancing, GPU, TPU, CPU, Distributed computing, Performance optimization, Scalability, Task scheduling

1. Introduction

The efficient management of computational resources in heterogeneous distributed systems is critical in optimizing performance and minimizing operational costs. The systems consist of these heterogeneous processing units-Graphics Processing Units (GPUs), Tensor Processing Units (TPUs) and Central Processing Units (CPUs), which perform particularly well at specialized computation tasks. GPUs are designed for parallel workloads, like deep learning and scientific simulations, where things are parallel; TPUs are optimized for neural network inference and training; CPUs offer versatility, good for things that aren't parallel or have branching (instructions that cause them to go to a different path). These heterogeneous units must be scheduled, resources must be allocated and loads must be balanced so that they do not incur bottlenecks and so that the resources are used optimally. In this research, we present an analysis of the interactions between hardware architecture, workload characteristics and dynamic scheduling mechanisms.

The technical frameworks discussed include containerized environments and resource orchestration, workload profiling and runtime optimization strategies. This research takes an important step in the direction of improving the productivity of these mixed-distributed systems by providing solutions for such issues as contention, thermal management and power consumption.

1.1. Resource Scheduling in Heterogeneous Systems

Workload distribution between computing units of different types, GPUs, TPUs and CPUs, requires resource scheduling in heterogeneous distributed systems. Traditional resource management approaches are designed to work for homogeneous systems and although capable, they do not fully exploit dedicated hardware, such as GPUs (graphics processing units) and TPUs (tensor processing units)¹. As heterogeneous systems arise, we must develop algorithms for dynamically placing tasks across diverse processing units sensitive to item-level workload, available processing power and user-specified energy utilization.

A few recently published papers at the Reinforcement Learning and Deep Discrete Optimization Workshop showed how recent scheduling algorithm advances can be used to predict where workloads should be placed to optimize performance metrics based on historical data using deep learning techniques. Furthermore, the problem's inherent complexity is addressed through scheduling in heterogeneous environments using techniques like Reinforcement Learning (RL) and Multi Objectives Optimization (MOO). These methods aim to optimize power efficiency while minimizing response time with respect to throughput balancing². In environments of distributed workload, it is necessary to avoid overload of particular processing units or to use impactfully resources by local and functional schedulers in a system. Thus, in addition, container orchestration systems like Kubernetes have been extended to support heterogeneous hardware by integrating native device drivers for GPUs and TPUs to provide dynamic resource allocation on the basis of task demand.

1.2. Workload partitioning and distribution strategies

Workload partitioning and distribution strategies that are efficient for a distributed environment, in which workloads and hardware accelerators are distributed, are important for balancing multiple modes in diverse hardware accelerators. In systems with hardware heterogeneity, tasks also need to be decomposed in a way that best takes advantage of the performance of each hardware unit (CPUs generally are better for doing low parallel tasks, GPUs and TPUs are better at highly parallel tasks like matrix multiplications or training deep learning models).

Task cloning was studied in³ as one method of automatic workload partitioning: complex tasks are split into several subtasks and these are assigned to the best hardware unit in terms of their computational capacity. In recent years, hybrid parallelism, having separate execution streams within a GPU or CPU, has been used to improve performance for batched operations in deep learning models ranging through frameworks such as TensorFlow and PyTorch. To implement these systems, we rely on both model parallelism and data parallelism strategies for efficient splitting of computational workloads. Also, we have a workload distribution scheme that includes load balancing and moving tasks between devices in order to avoid bottlenecks and devices that get utilized to maximum potential without under-utilisation or over-utilisation.

Different models implement predictive analytics to forecast workloads and preemptively allocate resources in order to increase the accuracy of task scheduling. Further, in making the allocation of resources to tasks more intelligent and adaptive, machine learning-based methods are employed to optimize task mapping as a function of historical performance data. To minimize all kinds of hardware conflicts while maximizing total system performance, the challenge lies in balancing heterogeneous workloads.

1.3. Energy efficiency and power-aware resource management

One central concern in heterogeneous distributed systems is energy consumption because a dedicated hardware accelerator requires a high amount of power. The goal of power-aware resource management is to minimize energy consumption that is subject to a certain level of system performance and independence from energy grids. Hardware-level power measurement tools are utilized in these systems to monitor and dynamically adjust the

use of the resource based on workload demands.

Power-aware scheduling has been proposed by a variety of techniques, such as dynamic voltage and frequency scaling (DVFS) and energy-efficient job placement algorithms. DVFS lowers the frequency used by the processing units to fit the workload and vice versa. The goal of energy-aware job placement is to place tasks into processing units that have capabilities to provide good performance and energy efficiency for the device. In addition to such means, there are some applications that incorporate machine learning models that predict the workloads' energy consumption and, consequently, perform scheduling optimization.

1.4. Fault tolerance and reliability in heterogeneous systems

In this environment, system dependability is more significant because failure characteristics and rates of heterogeneous computing platforms often differ - the heterogeneous array may contain, for instance, GPUs, TPUs and CPUs. In any of these processing units, some failures may cause a system to become unstable and may also lead to inefficiency. They have tackled these challenges by implementing other fault-tolerant mechanisms of the resource management system, which work in real-time to discover and rectify the faults.

Such methods as redundancy and checkpointing are more familiar when used in fault-tolerance systems. Redundancy is where instead of keeping one device that does a certain job, two or more such machines are employed; in case one fails, the other is there to take over without anyone noticing. Checkpointing saves the state periodically and so in the event of failure, it becomes easier to restore the system to the most recent previous stable state⁴. Furthermore, passive monitoring is commonly used to predict hardware failures before they occur since possible actions are to transfer resources to other locations in an attempt to avoid potential breakdowns.

Other mechanisms, including Paxos and Raft for the distributed consensus, are further discussed to maintain reliability in heterogeneous systems. These protocols aid in keeping data duplicated and synchronized so that computations can go on even with some nodes out of reach. Managed effectively, the inclusion of fault tolerance within the resource management strategy means that the system stays resilient in the face of hardware dysfunction, including high availability with minimal outages.

1.5. Scalability and adaptation to dynamic environments

In heterogeneous distributed systems, where the resource like CPU, GPU or even TPU is being dynamically allocated to cope with fluctuating workload as well as system context, scalability is a major challenge⁵. And these systems must scale horizontally, adding or removing processing units based on the workload.

Such elastic scaling solutions as Kubernetes and Apache Mesos have gained a lot of attention by offering a kind of dynamism, especially in the use of heterogeneous facilities to scale up or down as needed. They adapt the resources used by the system where the workload is required according to the processing requirements needed for the workload. However, achieving efficient scaling between different approximate hardware types poses definite challenges. Workload distribution across GPUs, TPUs and CPUs is a complex process: the system has to take into account not only the number of processing

units but also their performance, energy consumption and other factors⁶.

Further complicating the adaptation is the fact that we are operating in dynamic environments, with additional external factors such as varying network conditions, node availability and unpredictable user workloads. This research has resulted in the development of adaptive resource allocation algorithms capable of dynamically adjusting resource distribution for both optimized performance and lowered cost (**Figure 1**).

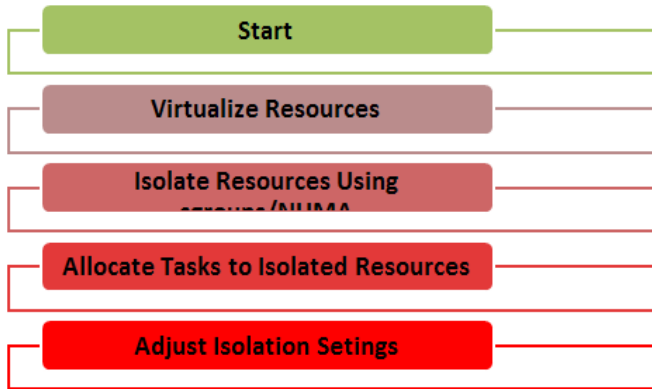


Figure 1: Resource Isolation and Virtualization.

Using resource isolation and virtualization, this diagram depicts efficient task execution without conflicts between workloads of different hardware types.

In the following, we describe an informative table summarizing the key resource management strategies that enable balancing GPU, TPU and CPU workloads in heterogeneous distributed systems⁷. The table illustrates some techniques, their use and the challenges that these techniques address.

The current table clearly explains in detail the resource management strategies in heterogeneous distributed systems along with the challenge of workload balancing of GPUs, TPUs and CPUs, fault tolerance, scalability and load balancing. Each strategy includes techniques that optimize performance and address system complexities. The essential property of robust scaling strategies and dynamic adaptation techniques can be used to satisfy the requirements of modern computing workloads for heterogeneous distributed systems, which provide the necessary flexibility and efficiency in the manner of using resources in the many processing units available.

Resource Management Strategy	Description	Purpose	Techniques/Methods	Challenges Addressed
Resource Scheduling	Involves dynamically allocating tasks to appropriate hardware based on workload requirements (CPU, GPU or TPU).	Maximize resource utilization, minimize latency and improve throughput.	Reinforcement Learning (RL) Multi-Objective Optimization (MOO) Kubernetes integration	Improper load balancing Performance degradation Inefficient resource usage
Workload Partitioning & Distribution	Efficient decomposition and distribution of tasks across CPUs, GPUs and TPUs to exploit hardware capabilities.	Ensure optimal usage of hardware accelerators (GPUs, TPUs) for specific workloads.	Task Cloning Hybrid Parallelism Data and Model Parallelism	Uneven workload distribution Underutilization or overloading of devices
Energy-Aware Resource Management	Power-efficient scheduling of tasks across heterogeneous hardware to minimize energy consumption while maintaining performance.	Reduce energy consumption without compromising performance.	Dynamic Voltage and Frequency Scaling (DVFS) Energy-Aware Scheduling Predictive Energy Modeling	Power inefficiency Balancing performance and energy consumption
Fault Tolerance & Reliability	Mechanisms to handle hardware failures or faults in CPUs, GPUs or TPUs, ensuring the system's reliability and performance in the event of failures.	Enhance system robustness and availability.	Redundancy Checkpointing Distributed Consensus Protocols (Paxos, Raft)	System downtime Inconsistent state after failure Fault recovery complexity
Scalability & Adaptation	The ability to scale resources horizontally and adapt to varying workloads in a dynamic environment.	Ensure flexible and responsive resource allocation as workloads change.	Elastic Scaling (e.g., Kubernetes) Predictive Resource Allocation (using ML models)	Scalability across different hardware types Handling resource demand fluctuations
Task Prioritization & QoS	Prioritizing tasks based on workload urgency and quality-of-service (QoS) requirements in heterogeneous environments.	Ensure critical tasks are executed first, maintaining performance levels for high-priority workloads.	Task Prioritization Algorithms Quality-of-Service (QoS) Policies Resource Allocation Scheduling	Resource starvation Managing latency and deadlines for time-sensitive tasks
Load Balancing	Distributes workloads across available processing units to avoid overloading any single unit, ensuring smooth operation and maximizing throughput.	Prevent bottlenecks and resource overloading.	Dynamic Load Balancing Algorithms Distributed Load Balancing Resource Migration	Uneven resource distribution Bottlenecks in processing units
Inter-device Communication	Efficient data exchange between heterogeneous hardware (CPUs, GPUs, TPUs) to optimize performance and reduce communication latency.	Improve inter-device communication efficiency and reduce overhead.	High-Performance Interconnects (e.g., NVLink, InfiniBand) Distributed Memory Models	Data transfer bottlenecks High communication latency

Resource Isolation & Virtualization	Ensures that workloads running on GPUs, TPUs and CPUs are isolated and don't interfere with each other while virtualizing hardware resources for efficient usage.	Prevent interference between workloads and ensure efficient utilization of each resource.	Virtualization Technologies (e.g., Docker, VMs) Resource Isolation Techniques (e.g., NUMA, groups)	Resource contention Virtualization overhead Inefficient isolation
Predictive Analytics for Resource Allocation	Predicts future resource requirements based on historical data, enabling proactive resource allocation to improve efficiency.	Optimize resource allocation and anticipate needs before they arise.	Machine Learning Models (e.g., neural networks, regression models) Historical Data Analysis	Unpredictable workloads Lack of accurate prediction Dynamic environment

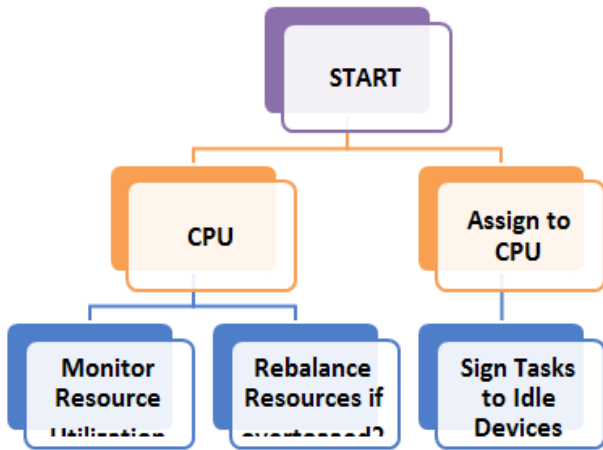


Figure 2: Resource Scheduling Flowchart.

In heterogeneous environments, this flowchart depicts a sequence of resource scheduling decisions as the decision making for assigning tasks to CPUs, GPUs or TPUs according to the characteristics of the workload.

This diagram shows how a workload on heterogeneous devices (CPUs, GPUs, TPUs) is partitioned and distributed. Task cloning and hybrid parallelism form the decision process here.

This flowchart illustrates the decision process for managing power and performance trade-offs⁸. It considers the energy efficient task scheduling techniques such as Dynamic Voltage and Frequency Scaling (DVFS).

This first diagram illustrates the fault tolerance integrated into the heterogeneous system. It deals with redundancy, checking out and system recovery protocols.

1.6. Applications of resource management in heterogeneous distributed systems

Balancing workloads across GPUs, TPUs and CPUs is an important problem that resource management strategies are increasingly important for in several high-performance computing domains. In general, these systems are used to govern machine learning, scientific computing, cloud computing, big data analytics and image processing.

In scientific computing, tasks such as simulations of molecular dynamics or weather modelling require significant computational power, where GPUs and TPUs can be employed to handle the parallelizable tasks, leaving CPUs to manage orchestration, I/O operations and system monitoring. Similarly, cloud computing services leverage heterogeneous environments to dynamically allocate resources based on user demand. This results in scalable and efficient infrastructure where workloads are balanced between CPUs for general-purpose processing and specialized processors (GPUs/TPUs) for high-performance tasks.

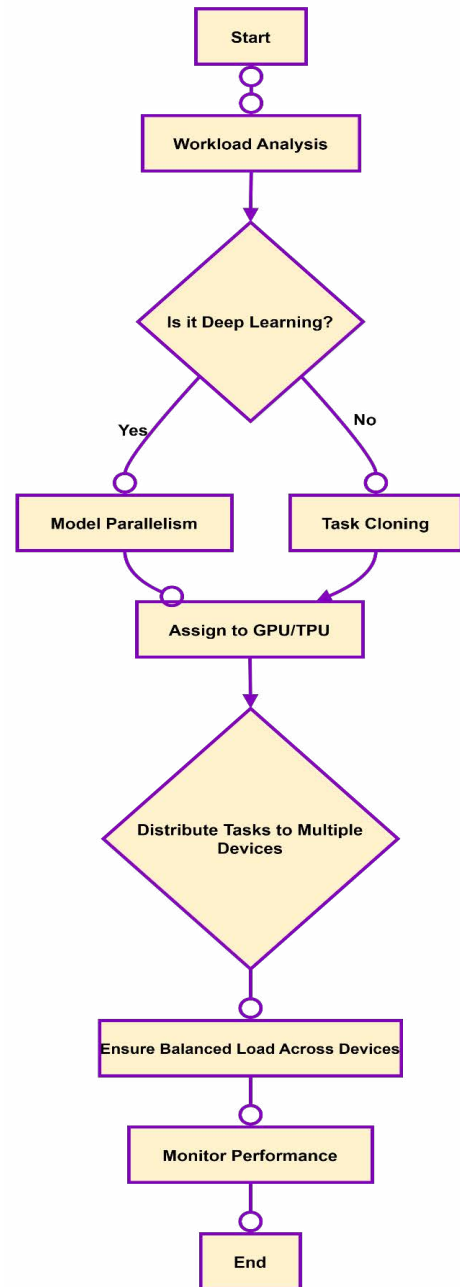


Figure 3: Workload Partitioning and Distribution.




Figure 4: Energy-Aware Scheduling Strategy.

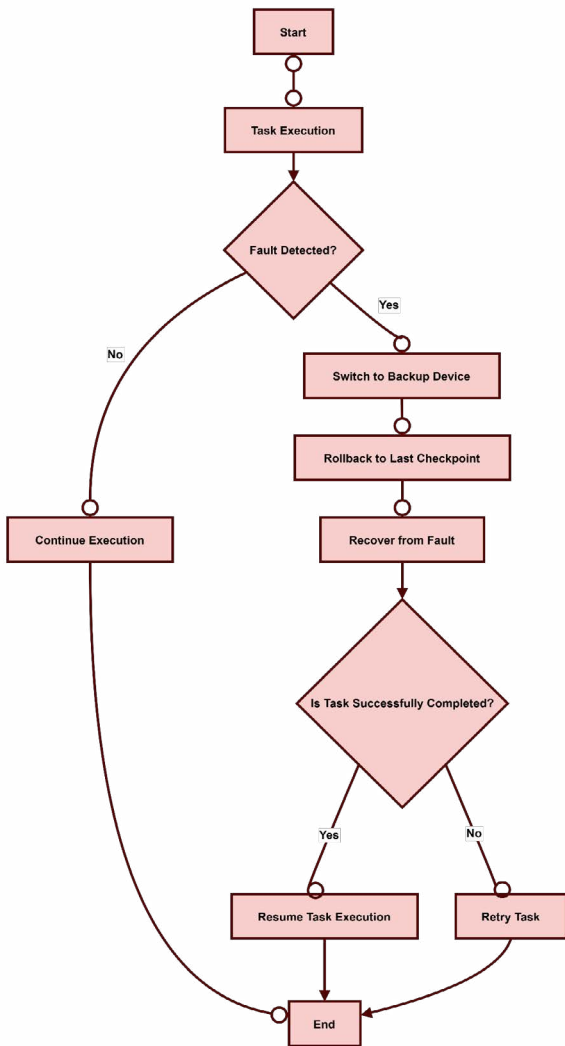


Figure 5: Fault Tolerance and Recovery Mechanism.

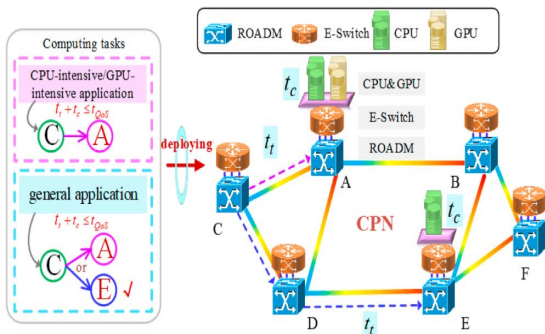


Figure 6: Applications in CPN with heterogeneous resources.

Big data analytics also benefits from resource management in these environments [10]. GPUs and TPUs accelerate data processing, enabling rapid analysis of large datasets, while CPUs manage the overall infrastructure and orchestration of data flow. In image processing and computer vision, GPUs and TPUs provide the computational capacity for handling large volumes of image data, making them crucial in industries like healthcare (e.g., for medical image analysis), autonomous driving and security surveillance systems.

1.7. Challenges in resource management in heterogeneous distributed systems

The resources in heterogeneous distributed systems that span GPUs, TPUs and CPUs are managed to present several challenges. The hardware itself is already heterogenous, which is one major

challenge. Some devices like CPUs and GPUs and TPUs have such different architectures and performance characteristics. GPUs and TPUs are made to perform well at parallel processing such as deep learning and matrix computations while CPUs are suited for attending sequential processing as well as general purpose computation.

The problem of loading balancing⁸ is another challenge. In multiple types of processor systems, some devices can be overloaded while other devices can remain underutilized. And it’s hard to get the optimal load distribution due to the fact that every processor type has different memory and computational requirements and that here can lead to bottlenecks. Real time monitoring of system performance and redistributing workload to avoid resource starvation or overuse is required to have dynamic load balancing algorithms.

Data movement and communication latency also pose significant challenges. In systems that schedule tasks across devices, GPUs and TPUs often need to talk to CPUs or each other. High data transfer latency between these devices¹² can severely degrade the overall performance of the system. Reducing these delays requires efficient protocols for communication and data transfer between all the nodes in the network and between them and the VMs.

Energy consumption and efficiency remain long-standing problems in heterogeneous distributed systems. The process of trying to balance energy spending and computational efficiency depends on the power consumption profile of the processors, which have many different architectures and it is hard to maintain a balance. Since data centers are becoming increasingly expensive, it is important to consider how to better manage resources and reduce energy consumption, using intelligent resource management strategies that take into account the GPUs, TPUs and CPUs that are used there.

1.8. Advantages of resource management in heterogeneous distributed systems

Heterogeneous distributed systems offer several notable advantages both in the strategic management of resources and by optimizing system performance, operational costs and scalability.

Performance optimization is one of the main benefits. These systems balance workloads between CPUs, GPUs and TPUs dynamically in order to place each task on the resource which will most effectively execute it, dramatically improving overall performance⁹. As a specific example, tasks like heavy parallel processing that rely on GPUs or TPUs for training data, for instance, deep learning, can be offloaded to those GPUs or TPUs, while tasks that are CPU heavy, such as data processing and control flow management, can be dedicated to the CPUs. Throughput in computationally demanding applications increases and it takes fewer resources to run those applications due to this fine-tuned allocation of resources.

An additional advantage is its scalability. In general, efficient scaling happens in heterogeneous distributed systems if our problem already has enough devices (more GPUs or TPUs) to accommodate growing workload needs. These systems easily accommodate small-size or large-scale workloads and can flexibly allocate resources based on real-time demand.

Heterogeneous systems significantly benefit from the cost-

effectiveness of the systems. These systems provide better resource utilisation, which means that companies can work at high performance without having to lay out large sums to invest in high-end hardware. By balancing the load across various processors, companies can make use of the existing resources to maximise their use without underspending expensive hardware like GPUs and TPUs while not overloading less expensive ones (CPUs).

Intelligent resource management improved energy efficiency as well. Systems can reduce the overall energy usage by large-scale computations via efficient allocation of tasks according to the power consumption profiles. Improving fault tolerance additionally guarantees that if one resource fails the others will step in without a steep performance penalty¹¹. As a result, it results in more resilient systems that can run satisfactorily even in the prese

2. Conclusion

In conclusion, this thesis concludes that efficient resource management in heterogeneous distributed systems, among GPUs, TPUs and CPUs, is necessary to maximize performance, scalability and energy efficiency. Deployment of computationally intensive applications is possible only by overcoming challenges such as hardware heterogeneity, load balancing and latency. The research also explores trade-offs between performance and energy consumption. However, in some cases, due to the long-running tasks, execution on an energy-efficient device that is slightly slower will provide much power savings, making this a workable design. Another energy-saving strategy is idle time management, where during periods of no activity, processors are powered down to save energy. The system is sustainable and yet able to meet performance targets through efficient power management.

3. Appendix: Pseudocode

```
# Energy-Aware Resource Scheduling Pseudocode
# Define resources (CPUs, GPUs, TPUs)
resources = {
    "CPU": {"type": "CPU", "power": 50, "performance": 2, "available": True}, # Power in watts, performance units
    "GPU": {"type": "GPU", "power": 200, "performance": 15, "available": True},
    "TPU": {"type": "TPU", "power": 250, "performance": 25, "available": True}
}
# Define workloads
workloads = [
    {"task_id": 1, "required_compute": 30, "task_type": "DeepLearning", "deadline": 100},
    {"task_id": 2, "required_compute": 10, "task_type": "DataProcessing", "deadline": 200},
    {"task_id": 3, "required_compute": 15, "task_type": "GraphicsRendering", "deadline": 150}
]
# Function to calculate energy efficiency of each resource
def calculate_energy_efficiency(resource):
    return resource["performance"] / resource["power"]

# Function to find the most suitable resource for a task
def find_best_resource_for_task(task):
    best_resource = None
    max_efficiency = -1

    # Check each available resource for suitability
    for resource_name, resource in resources.items():
        if resource["available"] and task["required_compute"] <= resource["performance"]:
            efficiency = calculate_energy_efficiency(resource)
            if efficiency > max_efficiency:
                best_resource = resource
                max_efficiency = efficiency
```

```

return best_resource

# Function to schedule tasks on the most suitable resources
def schedule_tasks(workloads):
    for task in workloads:
        # Find the best resource for each task
        best_resource = find_best_resource_for_task(task)

        if best_resource:
            # Assign task to the resource
            print(f"Task {task['task_id']} assigned to {best_resource['type']} with efficiency {calculate_energy_efficiency(best_resource):.2f}")
            best_resource["available"] = False # Mark resource as occupied
        else:
            print(f"Task {task['task_id']} could not be assigned due to insufficient resources.")
# Function to reset resource availability for the next scheduling round

def reset_resources():
    for resource in resources.values():
        resource["available"] = True
# Main scheduling loop
def main():
    # First scheduling round
    print("Scheduling round 1:")
    schedule_tasks(workloads)

    # Reset resources for the next round
    reset_resources()

    # Second scheduling round (different task set or changed conditions)
    print("\nScheduling round 2:")
    new_workloads = [
        {"task_id": 4, "required_compute": 40, "task_type": "DeepLearning", "deadline": 120},
        {"task_id": 5, "required_compute": 20, "task_type": "DataProcessing", "deadline": 180}
    ]
    schedule_tasks(new_workloads)
# Execute the main scheduling algorithm
if __name__ == "__main__":
    main()

```

4. References

1. Aggarwal Meenakshi, Khullar Vikas, Goyal Nitin, Rastogi Rashi, Singh Aman and Yelamos Vanessa, Albahar Marwan. Privacy preserved collaborative transfer learning model with heterogeneous distributed data for brain tumor classification. International Journal of Imaging Systems and Technology, 2023.
2. Thouheed Ahmed, Syed Kumar, Vinoth V, Mahesh TR, Prasad L, Velmurugan AK, Muthukumaran V, Niveditha VR. FedOPT: federated learning-based heterogeneous resource recommendation and optimization for edge computing. Soft Computing, 2024.
3. Anoushee Milad, Fartash Mehdi, Akbari Torkestani Dr. Javad. An intelligent resource management method in SDN based fog computing using reinforcement learning. Computing, 2023;106.
4. Bader J, Lehmann F, Thamsen L, Leser U and Kao O. "Lotar: Locally predicting workflow task runtimes for resource management on heterogeneous infrastructures," Future Generation Computer Systems, 2024;150: 171-185.
5. Buyya Rajkumar, Ilager Shashikant, Arroba Patricia. Energy-Efficiency and Sustainability in New Generation Cloud Computing: A Vision and Directions for Integrated Management of Data Centre Resources and Workloads, 2023.

6. <https://link.springer.com/article/10.1007/s11042-023-16399-2>.
7. <https://ieeexplore.ieee.org/document/10415079>.
8. Zahra A and N Mansouri. "A comprehensive survey on scheduling algorithms using fuzzy systems in distributed environments," *Artificial intelligence review*, 2024;57.
9. <https://arxiv.org/abs/2306.04207>.
10. <https://www.sciencedirect.com/science/article/pii/S2307187723002924?via%3Dihub>.
11. <https://arxiv.org/abs/2305.01974>.
12. <https://ieeexplore.ieee.org/document/10673918>.