

Rack Position Optimization in Large-Scale Heterogeneous Data Centers

Chang-Lin Chen¹, Jiayu Chen¹, Tian Lan², Zhaoxia Zhao³, Hongbo Dong³, and Vaneet Aggarwal¹

¹Purdue University

²George Washington University

³Meta Platforms, Inc.

Abstract

As rapidly growing AI computational demands accelerate the need for new hardware installation and maintenance, this work explores optimal data center resource management by balancing operational efficiency with fault tolerance through strategic rack positioning considering diverse resources and locations. Traditional mixed-integer programming (MIP) approaches often struggle with scalability, while heuristic methods may result in significant sub-optimality. To address these issues, this paper presents a novel two-tier optimization framework using a high-level deep reinforcement learning (DRL) model to guide a low-level gradient-based heuristic for local search. The high-level DRL agent employs Leader Reward for optimal rack type ordering, and the low-level heuristic efficiently maps racks to positions, minimizing movement counts and ensuring fault-tolerant resource distribution. This approach allows scalability to over 100,000 positions and 100 rack types. Our method outperformed the gradient-based heuristic by 7% on average and the MIP solver by over 30% in objective value. It achieved a 100% success rate versus MIP's 97.5% (within a 20-minute limit), completing in just 2 minutes compared to MIP's 1630 minutes (i.e., almost 4 orders of magnitude improvement). Unlike the MIP solver, which showed performance variability under time constraints and high penalties, our algorithm consistently delivered stable, efficient results—an essential feature for large-scale data center management.

1 Introduction

In large-scale data centers, efficient rack movement and resource allocation are important for balancing multiple operational objectives with respect to resource-efficiency, fault tolerance, and sustainability goals. Rising data center investments by major tech companies (Dynamics 2024; Reuters 2024) and increasing demands for data center efficiency highlight the importance of optimized rack positions, which directly impact movement counts, operational efficiency, and overall system resilience (Abbas et al. 2021). As data centers grow in complexity, continuous hardware updates and resource reallocation are needed to meet evolving service requirements. Optimized rack placement reduces movement counts, simplifies load balancing, and enhances fault tolerance—an increasingly important factor as data centers

host diverse, high-performance hardware with varied resource needs. Scalable and adaptable optimization solutions are thus important to sustain service quality, operational resilience, and environmental efficiency in modern data center environments, which is the focus of this paper.

Previous works have demonstrated the substantial impact of rack position optimization on data center performance. Seminal research by (Barroso, Hölzle, and Ranganathan 2019) established how rack layout designs significantly reduce operational energy usage, while (Micheliogiannakis et al. 2022) proved through production HPC experiments that rack-level resource disaggregation markedly improves utilization. Thermal management advances were made by (Zhang et al. 2023), demonstrating how optimized rack layouts enhance heat dissipation and energy efficiency. For heterogeneous environments, (Azimi, Zhan, and Reda 2014) developed effective rack-position mapping techniques for optimal cooling performance. Resource management innovations include (Zhu et al. 2020)'s microsecond-scale scheduler and (Jo et al. 2020)'s performance-boosting caching layer for rack-scale computing, collectively establishing rack positioning as important for both energy and resource optimization.

Composable data centers provide a promising solution by disaggregating resources like CPU, memory, and storage, enabling dynamic pooling and reallocation based on workload requirements (Li et al. 2017). This architectural approach stands in stark contrast to traditional systems that synchronize hardware replacement, which often leads to inefficiencies and delayed adoption of new technologies. The flexibility offered by composable architectures significantly enhances resource utilization and operational efficiency, making them particularly valuable for fast-evolving, heterogeneous AI environments. This architectural flexibility also removes many traditional constraints on rack positioning, enabling more opportunities for optimizing physical infrastructure layout. Strategic rack relocation simplifies scheduling, balances load distribution, and reduces maintenance overhead, especially as heterogeneity increases in data centers (Barroso, Hölzle, and Ranganathan 2019; Regaieg et al. 2021). Furthermore, the ongoing need to install new racks to meet growing demand underscores the importance of optimizing rack-position mapping at the data center level to maintain operational efficiency. Regular maintenance of

building infrastructure, power devices, and other facilities further necessitates the ability to efficiently move and position racks, ensuring minimal disruption during these routine activities.

This paper considers a region of rack-level disaggregated data centers consisting of heterogeneous racks and positions. To the best of our knowledge, this is the first work to optimize rack positioning for fault tolerance, ensuring resources from a subset of rack types are well-distributed across various scopes. Such fault tolerance empowers the hardware layer to provide reliability and availability guarantees to software services, important for optimal resource usage in large organizations. Consequently, we aim to identify a rack-position mapping that optimizes a weighted combination of rack movement count and fault tolerance metrics, while adhering to multi-layer resource constraints, including energy consumption, demand fulfillment, and placement limits. This novel problem is formulated as an integer non-linear optimization problem (INLP), for which computationally efficient approaches are lacking. To solve this problem, this paper proposes a two-tier optimization approach, with the Leader Reward approach (Wang et al. 2024) guiding the higher-level ordering of rack-type subproblems, and a gradient-based heuristic algorithm addressing each subproblem at the lower level, enabling efficient decomposition and targeted optimization for each rack type. Leader Reward, designed for standard combinatorial problems like the traveling salesman problem (TSP) and capacitated vehicle routing problem (CVRP), requires substantial memory due to its transformer-based architecture, making it impractical for direct application to our large-scale problem. Further, our problem does not have similar network/graph structure as in TSP/CVRP limiting direct usage of the approach. Therefore, we adapt it to our unique problem structure by incorporating an additional encoder, problem decomposition, and a tailored heuristic. The approach is shown to be scalable and outperform the MIP solvers as well as the greedy heuristic alone.

Summary of Contributions: The main contributions in this work are summarized as follows:

1. We introduce a rack movement optimization framework tailored for composable, heterogeneous data centers, which minimizes combined metrics associated with rack movement and fault tolerance. The problem is motivated by practical needs and can significantly advance datacenter intelligent management.
2. We formulate the problem as a computationally challenging INLP and address scalability issues by integrating a deep reinforcement learning (DRL) algorithm with a gradient-based heuristic in a tiered, collaborative architecture. The DRL algorithm employs Leader Reward, a technique well-suited for combinatorial tasks, to optimally order rack types, while the heuristic efficiently assigns specific positions. Experimental results demonstrate the heuristic’s competitiveness and speed.
3. Simulation results demonstrate the proposed framework’s superior performance, achieving a 7% improvement over a heuristic-only approach. Further, the proposed approach outperforms the MIP solver in terms of both the execution time

and objective value.

2 Related Works

Data center optimization has evolved through diverse methods, including ILP, MILP, heuristics, and more recent advancements in DRL, each targeting various challenges from power management to infrastructure optimization. However, rack layout optimization in disaggregated cloud systems remains largely unexplored.

Data center optimization methods Research on data center optimization spans several approaches. For power optimization, (Pahlavan, Momtazpour, and Goudarzi 2014) uses Integer Linear Programming (ILP) with heuristics to minimize active chassis while maintaining performance, while (Ajibola, El-Gorashi, and Elmirghani 2021) applies MILP and greedy algorithms for workload placement in composable architectures. For VM management, (Stefanello et al. 2019) combines MILP with genetic algorithms across geographically distributed data centers, and (Xia et al. 2017) uses hierarchical decomposition with MIP to address disk anti-colocation constraints. For infrastructure management, the Radar framework (Guo, Zukerman, and Wang 2023) tackles cooling challenges using MILP and heuristics, while (Xue and Zhu 2021) optimizes flow table installation in PDP switches using MINLP and k-median algorithms. Recent DRL applications include fault tolerance by logical cluster formation (Chen et al. 2025b), and network-aware resource allocation using GNNs (Shabka and Zervas 2023; Liu et al. 2022). However, optimizing rack layouts for new rack installations and maintenance, while accounting for broader management requirements such as fault tolerance, cooling, and energy efficiency, remains an unexplored area.

DRL methods for combinatorial problems (Kool, van Hoof, and Welling 2019) introduced attention-based architectures for neural combinatorial optimization (NCO), specifically tackling routing problems like TSP and VRP, which eliminated the need for handcrafted heuristics and established attention mechanisms as foundational in NCO. (Kwon et al. 2020) expanded on this with POMO, a multi-optima framework that optimizes across multiple optimal solutions within each instance, significantly enhancing convergence speed and solution diversity. (Grinsztajn, Furelos-Blanco, and Barrett 2023) then developed Poppy, a population-based reinforcement learning approach where a set of agents explores diverse solution paths, promoting thorough exploration and mitigating premature convergence. (Chalumeau et al. 2023) introduced Compass, which employs latent space search to dynamically adapt policies in response to problem structure variability, enabling the model to handle increasingly complex combinatorial tasks. Finally, (Wang et al. 2024) extended POMO by introducing a leader reward mechanism that prioritizes solutions near optimal regions, enhancing stability and precision by emphasizing high-quality solutions. Together, these methods illustrate substantial advancements in NCO through attention-driven models, policy diversity, adaptive exploration, and enhanced reward structures tailored to complex optimization problems.

Directly applying existing approaches is infeasible for our problem due to its large scale and unique structure, which differs from standard combinatorial problems like TSP and CVRP. To address this, we integrate the Leader Reward approach (Wang et al. 2024) with a gradient-based heuristic, selecting it over Compass and Poppy for its superior performance and memory efficiency (Chalumeau et al. 2023)—important for large-scale applications. Further, our algorithm can be directly used for online optimization for dynamic system environments since it captures the dynamic movement count.

3 System Model

This paper explores the strategic placement and adjustment of racks within data centers to address user needs and system requirements, and ensure fault-tolerant resource allocation. Users submit requests for a range of resources, such as CPUs, memory, disk storage, GPUs, and more. Consequently, the placement and position adjustment of racks is meticulously determined based on the evolving demands for these diverse resource types. Each rack is categorized into a specific type according to the combination of resources it offers. For example, racks equipped with GPUs and memory form one category, whereas those outfitted with CPUs and disk storage constitute another.

Within the complex ecosystem of data centers, designated spots known as positions define where each rack can be located, with the stipulation that each position can accommodate only one rack. These positions are aggregated into sets known as scopes, which are physical clusters such as regions, individual data centers, suites, and main switchboards (MSB). Some of the scopes are organized in a hierarchical fashion, where larger entities encompass several smaller ones, such as a region containing multiple data centers. The primary goal is to minimize system metrics, which include rack movement counts and penalties for resource imbalances across scopes, while simultaneously responding adaptively to user demands. This proactive rack repositioning is designed to enhance the system’s overall responsiveness and operational efficiency.

Let \mathbf{K} be the set of rack types. In this model, we assume that any two racks of the same type are fungible, so we do not work with individual rack. Let \mathcal{G} be a collection of subsets of \mathbf{K} , which we use to define metrics and constraints. Each $G \in \mathcal{G}$ may represent racks used for similar purposes, such as storage, computation, AI training, or racks used by specific product teams. Let \mathbf{P} be the set of all possible positions for hosting racks (expected to be on the order of 10^5), and let \mathcal{S} be a collection of subsets of \mathbf{P} . Each $s \in \mathcal{S}$ represents a position scope, such as a data center, MSB, or suite. In later formulations, we use S to represent a Boolean matrix where $S_{ij} = 1$ if position i is included in scope j . Let R be a resource matrix in which each row represents a rack type (indexed as in \mathbf{K}), and each column represents a resource type (e.g., power/cooling consumption, CPU, storage) or a Boolean flag (e.g., whether it is an AI rack). Define $X = [x_{p,k}]_{p \in \mathbf{P}, k \in \mathbf{K}}$ as a Boolean matrix of size $|\mathbf{P}| \times |\mathbf{K}|$, where the element $x_{p,k}$ at row p and column k is a binary variable indicating whether rack type k is present

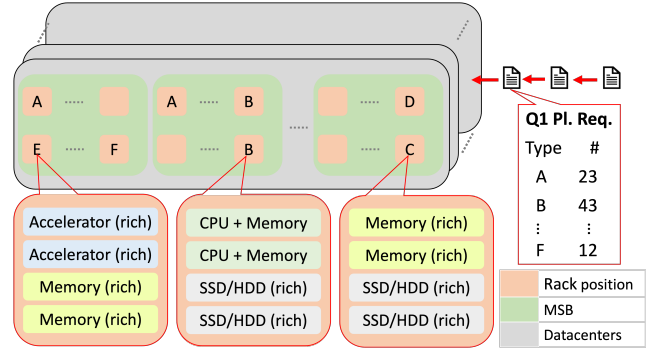


Figure 1: Simplified depiction of a rack movement system in heterogeneous data centers. The figure illustrates different rack types, labeled A, B, C, D, E, and F, each consisting of server nodes rich in specific resource types. The color-coded legend at the bottom represents different scopes within the datacenters, each with its resource limits. The figure also highlights placement requests, which indicate the number of racks required to meet resource usage demands. The rack movement system satisfies placement requests and service level objectives (SLOs) by dynamically moving racks or placing new racks within the data centers.

at position p . Position p is vacant if all entries in row p of X are zero, i.e., $x_{p,k} = 0$ for all k . The placement request for rack type k is denoted by d_k . To regulate resource overuse, a placement limit q is introduced to control the number of new racks placed in the system.

Position Mapping Constraint Assuming a position can accommodate at most one rack at a time, position-to-rack-type mapping, X , should satisfy the following constraint,

$$g_{1,p}(X) = \sum_{k \in \mathbf{K}} x_{p,k} \leq 1, \forall p \in \mathbf{P}. \quad (1)$$

Demand Fulfillment Moreover, to guarantee that the supply of rack types is enough to satisfy the placement request, X should satisfy the following constraint,

$$g_{2,k}(X) = \sum_{p \in \mathbf{P}} x_{p,k} - d_k \geq 0, \forall k \in \mathbf{K}. \quad (2)$$

Placement Limit The total number of racks newly placed should be less than the placement limit q .

$$g_3(X) = \sum_{k \in \mathbf{K}} \max \left(0, \sum_{p \in \mathbf{P}} x_{p,k} - \sum_{p \in \mathbf{P}} \bar{x}_{p,k} \right) - q \leq 0, \quad (3)$$

where $\bar{X} = [\bar{x}_{p,k}]_{p \in \mathbf{P}, k \in \mathbf{K}}$ represent the previous rack-position mapping. The number of newly placed racks are calculated by comparing the number of racks placed at current and previous time slots for all rack types and summing all the positive differences.

We aim to minimize rack movement and optimize resource distribution of all the resource types across various scopes.

Rack Movement Count When a rack moves to another position, it makes resources unavailable to the users and induces penalty of moving racks physically. The movement count of rack type k at position p is given as

$$o_{1,p,k}(X) = M_k \times \max(0, \bar{x}_{p,k} - x_{p,k}), \quad (4)$$

where M_k is weight of the movement count of rack type k .

Resource Spread Metrics For fault tolerance, resources from a subset of rack types should be well-distributed across a group of scopes to ensure resource availability and reliability in the event of a failure within any physical scope. Given the set of resource spread requirement r_i , G_i , \mathbf{c}_i , where r_i is a specific resource type, $G_i \subset \mathbf{K}$ is a subset of rack types, and \mathbf{c}_i is a union of mutually-exclusive scopes, the spread metrics is given as

$$o_{2,i}(X) = \tau(S[:, \mathbf{c}_i]^T \cdot X[\mathbf{c}_i, G_i] \cdot R[G_i, r_i]), \quad (5)$$

where $S[:, \mathbf{c}_i]^T \cdot X[\mathbf{c}_i, G_i] \cdot R[G_i, r_i]$ is the resource utilization of resource type r_i derived from rack types in G_i across the scopes in \mathbf{c}_i and τ is a multivariate function to measure the spread such as max, standard deviation, etc.

Scope Resource Consumption Limit To meet management requirements, such as limits for CPU allocation, energy consumption and bandwidth usage, we introduce scope-specific resource constraints for each resource type and scope. Let L represent the resource limit matrix of size $|\mathcal{S}| \times |\mathbf{R}|$, where each element $L[s, r]$ denotes the resource limit for resource type r in scope s . The resource limit constraint for each scope and resource type pair is then defined as:

$$o_{3,s,r}(X) = \zeta((L - S^T \cdot X \cdot R)[s, r]), \forall s \in \mathcal{S}, r \in \mathbf{R}. \quad (6)$$

where $\zeta(\cdot)$ is a function that returns a positive value for positive inputs and zero for non-positive inputs.

The following section shows the optimization problem considered and how it is optimized.

4 Problem Formulation

This section presents the problem considered in this work. The utility function is formulated considering the rack movement count and resource spread across scopes as,

$$\mathcal{U} = \sum_{k \in \mathbf{K}} \sum_{p \in \mathbf{P}} o_{1,p,k} + \beta_1 \sum_{i=1}^{|\mathcal{I}|} o_{2,i} + \beta_2 \sum_{s \in \mathcal{S}, r \in \mathbf{R}} o_{3,s,r}, \quad (7)$$

where β_1 and β_2 are the weights for the resource spread and resource limit penalties, respectively.

In this work, we aim to obtain high-quality solutions while considering the constraints of position mapping, demand fulfillment, and placement limit. We define our optimization problem and present it as:

$$\min_{x_{p,k}, \forall p \in \mathbf{P}, k \in \mathbf{K}} \mathcal{U}(X) \quad (8)$$

$$s.t. \quad g_{1,p}(X) \leq 1, \forall p \in \mathbf{P}, \quad (9)$$

$$g_{2,k}(X) \geq 0, \forall k \in \mathbf{K}, \quad (10)$$

$$g_3(X) \leq 0, \quad (11)$$

The above problem is a integer programming problem with the number of variables being $|\mathbf{P}||\mathbf{K}|$ and the number of constraints being $|\mathbf{P}| + |\mathbf{K}| + 1$. We note that the number of variables can go beyond millions for a region of data centers with over 100k positions and 100 rack types, limiting using commercial solvers directly. To obtain optimal results efficiently, we design a scalable heuristic approach integrated with a DRL-based combinatorial solver. The following section details each component of our proposed algorithm.

5 Proposed Algorithm

In this section, we present our proposed DRL-based approach to solving the rack movement optimization problem. Given the large scale of datacenters and the interdependencies between rack types as specified in functions (3), (5), and (6), we decompose the original optimization problem into subproblems by rack type. Each subproblem focuses on rack positioning for a single rack type, thereby reducing the number of variables and simplifying the overall complexity. The problem formulated in (8) can become a integer nonlinear programming problem (INLP) depending on the choice of functions $\tau(\cdot)$ in (5) and $\zeta(\cdot)$ in (6). As most existing solvers cannot handle such problems efficiently, we develop a gradient-based heuristic approach for the subproblems that can address any differentiable variant of the problem. Our gradient-based heuristic calculates the gradient of a composite function combining the objective function and the penalty for constraint violations. The algorithm iteratively flips either empty positions or positions occupied by the target rack type until the number of racks matches the requested quantity. The rack-position mappings are updated sequentially by applying the gradient-based heuristic to each rack type, one at a time, from the first to the last. The ordering of rack type optimization significantly impacts the solution quality due to varying characteristics of each rack type, including their resource compositions, current system quantities, and demand levels. For example, optimizing rack types with lower demand first provides greater flexibility for subsequent rack types to achieve better objective values while satisfying constraints. Therefore, we employ a DRL-based combinatorial problem solver to optimize the sequence of suboptimizations. The complete algorithm comprises a gradient-based heuristic for rack-position mapping and a DRL-based solver for sequencing suboptimizations. We detail the subproblems in Section 5.1, followed by the gradient-based heuristic and DRL-based solver in Sections 5.2 and 5.3, respectively.

5.1 Suboptimization

The original optimization problem in (8) is decomposed into subproblems based on rack types:

$$\min_{x_{p,k}, \forall p \in \mathbf{P}} \mathcal{U}(X^{\bar{\mathbf{K}}}) \quad (12)$$

$$s.t. \quad g_{1,p}(X^{\bar{\mathbf{K}}}) \leq 1, \forall p \in \mathbf{P}, \quad (13)$$

$$g_{2,k}(X^{\bar{\mathbf{K}}}) \geq 0, \quad (14)$$

$$g_3(X^{\bar{\mathbf{K}}}) \leq 0, \quad (15)$$

where $k \notin \bar{\mathbf{K}}$ denotes the rack type for the current subproblem, $X^{\bar{\mathbf{K}}}$ represents the rack-position mapping matrix with

Algorithm 1: Proposed heuristic algorithm

Input: $\{X, S, R, L, d_k, q, \{(c_i, G_i, r_i)\}, \bar{\mathbf{K}}, \Gamma\}$
Output: $\{x_{p,k}, \forall p \in \mathbf{P}\}$

```

1 for  $k = 1$  to  $|\bar{\mathbf{K}}|$  do
2    $allocations \leftarrow \sum_{p \in \mathbf{P}} x_{p,k} - d_k$ 
3    $iter \leftarrow abs(allocations)$ 
4   for  $i = 1$  to  $iter$  do
5     if  $allocations < 0$  then
6        $\Phi \leftarrow \{p' | \sum_{k \in \bar{\mathbf{K}}} x_{p,k} == 0\}$ 
7        $x_{p',k} \leftarrow 1$ , where
8          $p' = argmax \nabla_{X^{\bar{\mathbf{K}}}} f(X^{\bar{\mathbf{K}}}) \cdot \sum_{p \in \Phi} e_p$ 
9         and  $e_p$  is an elementary vector
10      else
11         $\Phi \leftarrow \{p' | \sum_{k \in \bar{\mathbf{K}}} x_{p,k} == 1\}$ 
12         $x_{p',k} \leftarrow 0$ , where
13           $p' = argmax - \nabla_{X^{\bar{\mathbf{K}}}} f(X^{\bar{\mathbf{K}}}) \cdot \sum_{p \in \Phi} e_p$ 
14      end
15    end
16  end
17  for  $i = 1$  to  $\Gamma$  do
18     $\Phi \leftarrow \{p' | \sum_{k \in \bar{\mathbf{K}}} x_{p,k} == 1\}$ 
19    if  $\Phi \neq \emptyset$  then
20       $x_{p',k} \leftarrow 0$ , where
21         $p' = argmax - \nabla_{X^{\bar{\mathbf{K}}}} f(X^{\bar{\mathbf{K}}}) \cdot \sum_{p \in \Phi} e_p$ 
22       $x_{p',k} \leftarrow 1$ , where
23         $p' = argmax \nabla_{X^{\bar{\mathbf{K}}}} f(X^{\bar{\mathbf{K}}}) \cdot \sum_{p \in \Phi^c} e_p$ 
24    end
25  end
26 end

```

rack types in $\bar{\mathbf{K}}$ determined by previous subproblem solutions, and constraints $g_{1,p}$ and g_3 follow the same structure as in (8) but operate on $X^{\bar{\mathbf{K}}}$.

5.2 Gradient Based Heuristic Algorithm

We propose a gradient-based heuristic algorithm that sequentially solves the subproblem in (12) for each rack type to obtain a solution to the original optimization problem in (8). The algorithm iteratively updates rack positions by following the gradient of an augmented objective function:

$$f(X^{\bar{\mathbf{K}}}; k) = \mathcal{U}(X^{\bar{\mathbf{K}}}) + \gamma g_3(X^{\bar{\mathbf{K}}}), \quad (16)$$

where k denotes the rack type being optimized (corresponding to the variable column in $X^{\bar{\mathbf{K}}}$), g_3 is the inequality constraint from (12), and $\gamma > 0$ is the penalty coefficient for constraint violations. For each rack type, the algorithm enforces constraints $g_{1,p,k}$ and $g_{2,k}$ by: (1) using a masking mechanism to exclude non-applicable positions (positions already assigned to other rack types when additional positions are needed, or vacant positions when some need to be cleared), and (2) ensuring the required number of racks d_k for type k . Finally, Γ rounds of adjustments are applied to further minimize the objectives. The complete algorithm is presented in Algorithm 1.

5.3 Proposed Two-Tier Algorithm

In this section, we introduce our proposed DRL-based approach for the rack movement optimization problem. As out-

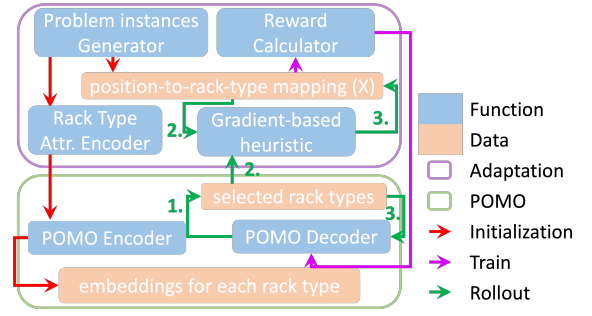


Figure 2: Our proposed algorithm integrates Leader Reward and adaptation components for rack positioning optimization. It comprises two main modules: an adaptation module (purple box) containing the problem generator and gradient-based heuristic; and a POMO module (green box) featuring an encoder-decoder architecture. The workflow encompasses initialization (red arrows), training (purple arrows), and rollout (green arrows) phases, with data blocks (peach boxes) showing the information passed between functions.

lined in Section 4, the problem’s complexity scales with the number of positions, rack types, scopes, and resource types, making real-time optimization challenging. To address this, we decompose the original problem into subproblems by rack type in Section 5.1 and employ a gradient-based heuristic to sequentially solve each subproblem in Section 5.2. The order of these suboptimizations is important to achieving optimal system performance.

To determine the optimal sequence for suboptimizations, we leverage Leader Reward (Wang et al. 2024), a method proven effective for NP-hard combinatorial problems such as the TSP and CVRP. We further adapt Leader Reward to our context by encoding each rack type’s attributes—demand, current quantity, resource types, and rack groups—into two float values. These values are then converted into embeddings by the encoder, allowing the decoder to learn the optimal sequence for rack type suboptimizations.

During training, the decoder generates permutations of rack types, representing different suboptimization sequences. The new rack-position mapping is then obtained by the greedy heuristic in Algorithm 1 given the permutations (line 4 in Algorithm 2). Each permutation’s performance is evaluated using cumulative rewards, which reflect objective in (16) with the new rack-position mapping as input. By leveraging multiple trajectories and a shared reward baseline, the proposed framework effectively explores the solution space, learning an optimal policy for suboptimization sequencing.

In summary, our algorithm combines a gradient-based heuristic for rack-position mapping with a DRL-based solver to determine the optimal suboptimization order, as outlined in Algorithm 2.

6 Simulation

In this section, we evaluate the performance of our proposed algorithm. We begin with the system setup in Section 6.1,

Parameter	\mathbf{P}	\mathbf{K}	\mathbf{R}	\mathcal{S}_1	\mathcal{S}_2	\mathcal{S}_3	$[d_{min}, d_{max}]$	$[q_{min}, q_{max}]$	$[L_1^{max}, L_1^{min}]$	$[L_2^{min}, L_2^{max}]$	$[L_3^{min}, L_3^{max}]$
Value	1000	10	10	2	10	50	[20, 60]	[800, 1000]	[3,6]	[30,50]	[90, 110]

Table 1: Parameter Setup

Rack Type	0	1	2	3	4	5	6	7	8	9	ϕ
Probability	0.033	0.014	0.030	0.010	0.009	0.102	0.029	0.204	0.018	0.051	0.5

Table 2: Probability Distribution of Rack Types for Each Position

Algorithm 2: Adapted Leader Reward Training

Input: Number of training epochs T , batch size B , problem set P , number of rack types \mathbf{K}

Output: Updated policy network parameter θ

```

1 Initialize policy network parameter  $\theta$ ;
2 for epoch = 1 to  $T$  do
3    $\delta_i \leftarrow \text{GENERATEPROBLEMS}(P), \forall i \in \{1, \dots, B\}$ 
   where
    $\delta_i = (\bar{X}, d_1, \dots, d_{\mathbf{K}}, q, (a_1, b_1), \dots, (a_{\mathbf{K}}, b_{\mathbf{K}}))_i$  with
    $\bar{X}$ : initial rack-position mapping and  $(a_k, b_k)$ : Inputs
   for the encoder derived from rack type attributes for
   each  $k \in \mathbf{K}$ ;
4    $\tau_i^j, X_i^j \leftarrow \text{ROLLOUT+GRADIENTHEURISTIC}(j, \delta_i, \pi_\theta), \forall i \in \{1, \dots, B\}, \forall j \in \{1, \dots, \mathbf{K}\}$ ;
5    $b_i \leftarrow \frac{1}{N} \sum_{j=1}^N \text{Rew}(X_i^j, \delta_i), \forall i \in \{1, \dots, B\}$ ;
6    $\nabla_\theta J(\theta) \leftarrow \frac{1}{BN} \sum_{i=1}^B \sum_{j=1}^N (\text{Rew}(X_i^j, \delta_i) - b_i) \nabla_\theta \log p_\theta(\tau_i^j)$ ;
7    $\theta \leftarrow \theta + \alpha \nabla_\theta J(\theta)$ ;
8 end

```

clarify the baselines in Section 6.2, and show the performance comparison in Section 6.3.

6.1 System Setup

We simulate the hierarchical data centers by three sets of position scopes, \mathcal{S}_1 , \mathcal{S}_2 , and \mathcal{S}_3 , representing data centers, suites, and MSBs. Each scope is a subset of \mathbf{P} and any scope in \mathcal{S}_i is contained in one of the scopes in \mathcal{S}_j for $(i, j) \in \{(3, 2), (2, 1)\}$. Scopes within each set are equal in their size and exclusive to each other. The union of scopes in each set is \mathbf{P} and the union of \mathcal{S}_1 , \mathcal{S}_2 , and \mathcal{S}_3 is \mathcal{S} in (6). The resource limit matrices for sets \mathcal{S}_1 , \mathcal{S}_2 , and \mathcal{S}_3 are denoted as L_1 , L_2 , and L_3 , respectively. Concatenation of L_1 , L_2 , and L_3 along the row dimensions is equal to L in (6). Table 1 delineates the key parameters considered in this work. The counts of positions, rack types, resource types are symbolized as $|\mathbf{P}|$, $|\mathbf{K}|$, $|\mathbf{R}|$, $|\mathcal{S}|$, d , and q , respectively. Each scope in \mathcal{S}_1 , \mathcal{S}_2 , \mathcal{S}_3 has 500, 100, 20 positions, respectively. The demands for each rack types are uniformly sampled in the range $[d_{min}, d_{max}]$ and placement limit is sampled uniformly from the range $[q_{min}, q_{max}]$. Finally, we use the standard deviation and softplus functions for $\tau(\cdot)$ and $\zeta(\cdot)$, respectively, in our implementation to measure resource spread and penalize resource limit violations, as formulated in (5) and (6), making the problem in (8) a INLP.

This datacenter setup is designed by domain experts among the authors, approximating a real-life scenario.

The previous rack-position mapping is generated using the probability distribution in Table 2, which has been reviewed and approved by a domain expert among the authors to reflect a realistic scenario. This distribution shows that the no-rack category (ϕ) has the highest probability at 0.5, meaning that 50% of positions are unassigned. Among rack types, rack type 7 has the highest probability at 0.204, making it the most likely to be assigned when a rack is present, followed by rack types 5 and 9, with probabilities of 0.102 and 0.051, respectively. The remaining rack types have lower probabilities, ranging from 0.009 to 0.029. The matrix R representing the rack-type-to-resource-type mapping is shown in Table 3.

Our proposed algorithm is trained using a single CPU core and a single GPU core on a system with two Milan CPUs @ 2.45GHz and four NVIDIA A100 GPUs.

Resource Type	0	1	2	3	4	5	6	7	8	9
Rack Type 0	0	0	0	0	0	1	1	0	1	0
Rack Type 1	0	1	0	0	1	1	0	0	0	0
Rack Type 2	0	0	0	0	1	0	1	0	1	0
Rack Type 3	0	1	1	0	0	0	1	0	0	1
Rack Type 4	0	1	0	1	0	0	0	0	1	0
Rack Type 5	0	0	1	1	1	0	1	0	1	0
Rack Type 6	1	0	1	1	0	0	0	0	0	0
Rack Type 7	1	1	0	1	0	0	1	0	0	0
Rack Type 8	0	0	0	0	1	0	0	0	1	0
Rack Type 9	1	1	0	0	0	0	0	0	0	1

Table 3: Summary of the Rack Type Resources

6.2 Baseline Algorithms

This section describes the baseline algorithms. We include the MIP solver, which was widely used for operation optimizations, and the gradient-based heuristic proposed in this paper. **MIP solver:** We implemented the solver using Gurobi v11.0.3 (Gurobi Optimization, LLC 2024). The solver’s time limit configuration and performance analysis are presented in the Section 6.3. Gurobi was run using 18 threads on an Intel(R) Core(TM) i9-9980XE CPU @ 3.00GHz. When Gurobi cannot find a solution within the specified time limit, we retain the previous rack-type-to-position mapping as the solution. **Gradient-based heuristics:** This is implemented by PyTorch (Paszke et al. 2019) following Algorithm 1 to allocate positions for one rack type at a time. The order of rack is predefined.

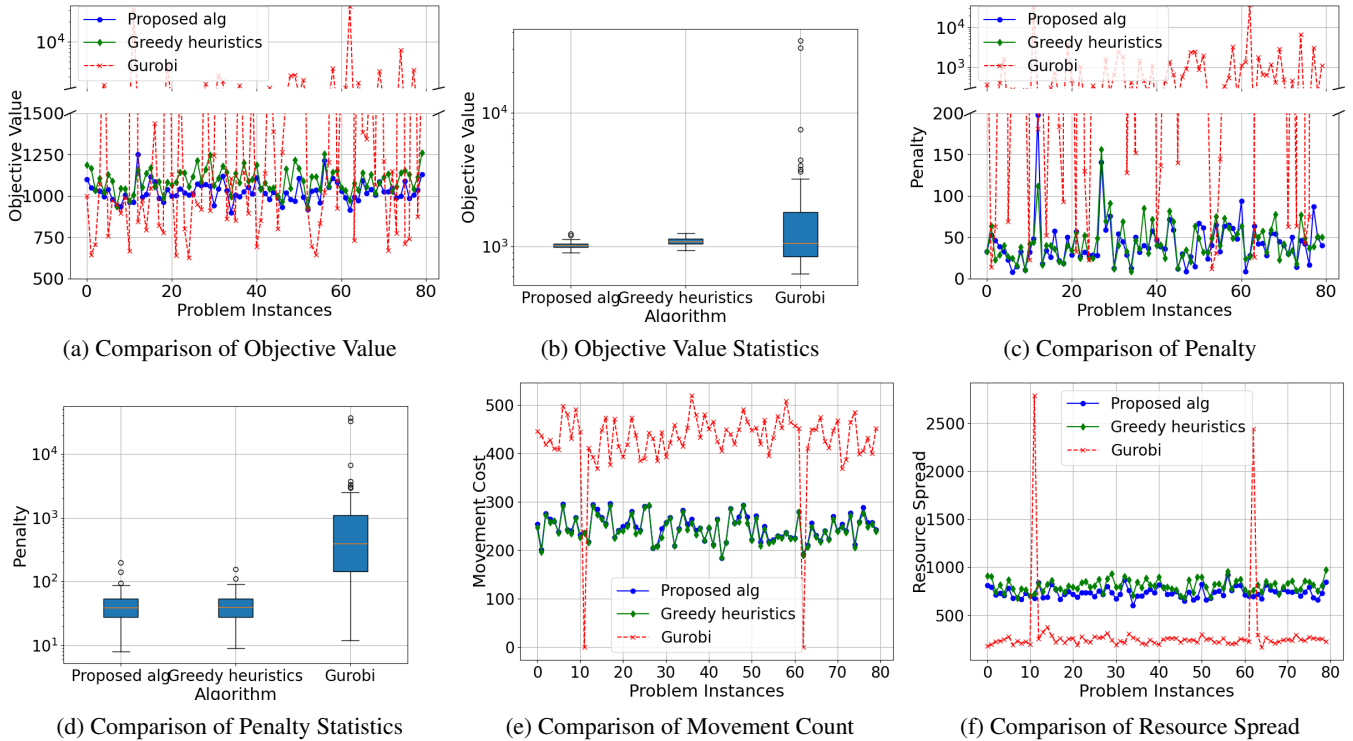


Figure 3: Comparison of the Proposed Algorithm, gradient Heuristics, and Gurobi across different metrics.

Algorithm	MIP	MIP	MIP	MIP	Proposed
Time limit (min)	1	5	10	20	N/A
Solving percentage (%)	0	6.25	81.25	97.5	100
Running time (min)	110	630	830	1630	2
Training time (min)	N/A	N/A	N/A	N/A	600
Performance \mathcal{U}	38816	1600	1563	1475	1025

Table 4: Summary of the algorithm capability

6.3 Main Results

We compare the performance of our proposed algorithms with the baseline algorithms in this section. In Figure 3, we compare the baselines and our algorithm on the overall performance, resulting in movement count, resource limit violation penalty, and resource spread metric. We compare the execution time of MIP and our proposed algorithm in Table 4. Finally, we discuss the decision made by Leader Reward given attributes of the rack types in Figure 4.

Figure 3 compares the performance of three algorithms — our proposed algorithm, gradient-based heuristics, and MIP solver — across 80 problem instances. The comparison is based on four key metrics: movement count, objective value, resource limit penalty, and resource spread metric, where the objective value is the summation of the other three metrics. In Figure 3a, the proposed algorithm and gradient-based heuristics show similar performance, with objective values ranging between 1000 and 1400. Moreover, the proposed algorithm outperforms the gradient-based heuristics for nearly all of the 80 problem instances and the average objective values of the proposed algorithm outperforms the

gradient-based heuristics by 7%. Comparing with that the results for the other two algorithms remain consistent and relatively smooth across all problem instances, Gurobi presents a more erratic behavior. Although it maintains reasonable objective values for most instances, there are extreme spikes where the objective value exceeds 10^4 , where the solver reaches the time limit of 20 minutes and cannot get a solution. Among the 80 problem instances, the proposed algorithm outperforms the MIP solver for 46 instances and the average objective value of the proposed algorithm outperforms the one for the MIP solver by more than 30%. Figure 3e, 3c, and 3f show that the MIP solver succeeds in minimizing resource spread metrics, but incurs high movement count and resource limit penalty. At instances 12 and 63, where the MIP solver cannot get a solution in time, the previous rack-type-to-position mapping is used as the action and thus incurs zero movement count and high penalty and resource spread. In conclusion, the proposed algorithm and gradient-based heuristic demonstrate minimal variability, highlighting their robustness in addressing problem instances—an essential feature for large-scale data center management. In contrast, the MIP solver demonstrates significant instability, with sharp spikes in all metrics for certain instances. This suggests that while the MIP solver may occasionally achieve competitive results, it struggles to maintain consistent performance, particularly in terms of balancing the components of the objective function.

Table 4 presents a comparative analysis of the MIP solver with various time limits (1, 5, 10, and 20 minutes) and the proposed algorithm, evaluated across multiple metrics:




Figure 4: Top-three reward rankings for Rack Type 1, Rack Type 7, and Rack Type 9 across 80×10 problem instances. Each figure shows the number of instances where the respective rack type, placed in a specific position within the order, achieved a top-three reward ranking (rank_1, rank_2, or rank_3). The x-axis represents the order position of the rack type, while the y-axis shows the number of instances. Rack Type 1 performs best when placed in position 3, Rack Type 7 performs optimally in position 6, and Rack Type 9 achieves the highest rewards when placed in position 8, indicating preferred order positions for each rack type to maximize performance.

solving percentage, running time, training time, and objective value. The proposed algorithm demonstrates superior performance in all metrics. It successfully solves 100% of the problem instances, outperforming all MIP configurations, even those with extended time limits. In contrast, MIP with shorter time limits (1 and 5 minutes) struggles, solving 0% and 6.25% of the instances, respectively. MIP shows improved performance with longer time limits—solving 81.25% and 97.5% of instances at 10 and 20 minutes, respectively—though it still falls short of the proposed method. Regarding running time, the proposed approach is remarkably efficient, requiring just 2 minutes to complete all problem instances. By comparison, MIP algorithms require significantly more time, with running times ranging from 110 minutes for the 1-minute time limit to 1630 minutes for the 20-minute time limit. This stark contrast in computational efficiency highlights the practicality of the proposed method. While training time is not applicable to MIP, the proposed algorithm requires 600 minutes of training time. However, this one-time training time is offset by its exceptional runtime efficiency. In terms of objective value, the proposed algorithm achieves the best performance with a score of 1025, outperforming all MIP variants. MIP’s performance improves as the time limit increases, with the best MIP configuration (20 minutes) yielding a performance of 1475. However, even this best-case MIP performance lags behind the proposed algorithm.

Fig. 4 illustrate the performance of different rack type orders for three specific rack types - Rack Type 1, Rack Type 7, and Rack Type 9 - evaluated across 80×10 problem instances using a Leader Reward agent. As shown in Table 3, Rack Type 1, Rack Type 7, and Rack Type 9 are selected for the plots because they are either very unlikely or likely to be found in the data center positions previously. Each figure shows the number of instances where the respective rack type, placed in a specific order position, achieved a top-three reward ranking (rank_1, rank_2, or rank_3). The objective is to identify the optimal position for each rack Type within the order that maximizes its reward. For Rack Type 1, the opti-

mal position is found to be position 3, where it achieves the highest number of top-three rewards across the problem instances. Positions beyond 4, however, show a sharp decline in performance, indicating that Rack Type 1 is most effective when placed early in the sequence, particularly around the third position. The performance of Rack Type 7 is clearly optimized when it is placed in position 6. Positions 4 and 5 also perform well, although to a lesser degree, suggesting that Rack Type 7 benefits from being placed toward the latter half of the order. For Rack Type 9, the ideal position is position 8, which consistently delivers the highest number of top-three rewards across the instances. Positions 7 and 9 also contribute to a significant number of top-three results, particularly in rank_2 and rank_3, indicating that Rack Type 9 performs best when placed at or near the end of the sequence.

7 Conclusions

This work presents a scalable, DRL-based solution for optimizing rack movement in large, heterogeneous data centers. Our two-tier approach, combining Leader Reward with a greedy heuristic, outperforms traditional MIP solvers and heuristics by reducing movement counts, ensuring fault tolerance, and improving execution times. Simulations confirm that our method achieves efficient, consistent results across varying demands, highlighting its practical application in dynamic data center environments.

Acknowledgments

C. Chen, J. Chen, T. Lan, and V. Aggarwal acknowledge research award from Meta Platforms, Inc., and Office of Naval Research under grant N00014-23-1-2532. The paper with supplementary is available at (Chen et al. 2025a).

References

Abbas, A.; Huzayyin, A.; Mouneer, T.; and Nada, S. 2021. Thermal management and performance enhancement of data centers architectures using aligned/staggered in-row cooling arrangements. *Case Studies in Thermal Engineering*, 24: 100884.

- Ajibola, O. O.; El-Gorashi, T. E. H.; and Elmirghani, J. M. H. 2021. Energy Efficient Placement of Workloads in Composable Data Center Networks. *Journal of Lightwave Technology*, 39(10): 3037–3063.
- Azimi, R.; Zhan, X.; and Reda, S. 2014. Thermal-aware layout planning for heterogeneous datacenters. In *Proceedings of the 2014 International Symposium on Low Power Electronics and Design*, 245–250.
- Barroso, L. A.; Hölzle, U.; and Ranganathan, P. 2019. *The datacenter as a computer: Designing warehouse-scale machines*. Springer Nature.
- Chalumeau, F.; Surana, S.; Bonnet, C.; Grinsztajn, N.; Pretorius, A.; Laterre, A.; and Barrett, T. 2023. Combinatorial optimization with policy adaptation using latent space search. *Advances in Neural Information Processing Systems*, 36: 7947–7959.
- Chen, C.-L.; Chen, J.; Lan, T.; Zhao, Z.; Dong, H.; and Aggarwal, V. 2025a. Rack Position Optimization in Large-Scale Heterogeneous Data Centers. arXiv:2504.00277.
- Chen, C.-L.; Zhou, H.; Chen, J.; Pedramfar, M.; Lan, T.; Zhu, Z.; Zhou, C.; Mauri Ruiz, P.; Kumar, N.; Dong, H.; and Aggarwal, V. 2025b. Learning-Based Two-Tiered Online Optimization of Region-Wide Datacenter Resource Allocation. *IEEE Transactions on Network and Service Management*, 22(1): 572–581.
- Dynamics, D. C. 2024. Meta to “accelerate infrastructure investments” for AI, raises capex to \$35bn-\$40bn, <https://www.datacenterdynamics.com/en/news/meta-to-accelerate-infrastructure-investments-for-ai-raises-capex-to-35bn-40bn>.
- Grinsztajn, N.; Furelos-Blanco, D.; and Barrett, T. D. 2023. Population-Based Reinforcement Learning for Combinatorial Optimization Problems, <https://openreview.net/forum?id=JRFSLFyYAIL>.
- Guo, C.; Zukerman, M.; and Wang, T. 2023. Radar: Reliable Resource Scheduling for Composable/Disaggregated Data Centers. *IEEE Transactions on Industrial Informatics*, 19(8): 8551–8563.
- Gurobi Optimization, LLC. 2024. Gurobi Optimizer Reference Manual.
- Jo, C.; Kim, H.; Geng, H.; and Egger, B. 2020. RackMem: A Tailored Caching Layer for Rack Scale Computing. In *Proceedings of the ACM International Conference on Parallel Architectures and Compilation Techniques*, 467–480.
- Kool, W.; van Hoof, H.; and Welling, M. 2019. Attention, Learn to Solve Routing Problems! In *International Conference on Learning Representations*.
- Kwon, Y.-D.; Choo, J.; Kim, B.; Yoon, I.; Gwon, Y.; and Min, S. 2020. Pomo: Policy optimization with multiple optima for reinforcement learning. *Advances in Neural Information Processing Systems*, 33: 21188–21198.
- Li, C.-S.; Franke, H.; Parris, C.; Abali, B.; Kesavan, M.; and Chang, V. 2017. Composable architecture for rack scale big data computing. *Future Generation Computer Systems*, 67: 180–193.
- Liu, W.-X.; Lu, J.; Cai, J.; Zhu, Y.; Ling, S.; and Chen, Q. 2022. DRL-PLink: Deep Reinforcement Learning With Private Link Approach for Mix-Flow Scheduling in Software-Defined Data-Center Networks. *IEEE Transactions on Network and Service Management*, 19(2): 1049–1064.
- Michelogiannakis, G.; Klenk, B.; Cook, B.; Teh, M. Y.; Glick, M.; Dennison, L.; Bergman, K.; and Shalf, J. 2022. A Case For Intra-rack Resource Disaggregation in HPC. *ACM Trans. Archit. Code Optim.*, 19(2).
- Pahlavan, A.; Momtazpour, M.; and Goudarzi, M. 2014. Power reduction in HPC data centers: a joint server placement and chassis consolidation approach. *The Journal of Supercomputing*, 70(2): 845–879.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Köpf, A.; Yang, E.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J.; and Chintala, S. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. arXiv:1912.01703.
- Regaieg, R.; Koubàa, M.; Ales, Z.; and Aguilu, T. 2021. Multi-objective optimization for VM placement in homogeneous and heterogeneous cloud service provider data centers. *Computing*, 103(6): 1255–1279.
- Reuters. 2024. Microsoft, OpenAI plan \$100 billion data-center project.
- Shabka, Z.; and Zervas, G. 2023. Network-aware compute and memory allocation in optically composable data centers with deep reinforcement learning and graph neural networks. *Journal of Optical Communications and Networking*, 15(2): 133–143.
- Stefanello, F.; Aggarwal, V.; Buriol, L. S.; and Resende, M. G. 2019. Hybrid algorithms for placement of virtual machines across geo-separated data centers. *Journal of Combinatorial Optimization*, 38: 748–793.
- Wang, C.; Cheng, P.; Li, J.; and Sun, W. 2024. Leader Reward for POMO-Based Neural Combinatorial Optimization. arXiv:2405.13947.
- Xia, Y.; Tsugawa, M.; Fortes, J. A. B.; and Chen, S. 2017. Large-Scale VM Placement with Disk Anti-Colocation Constraints Using Hierarchical Decomposition and Mixed Integer Programming. *IEEE Transactions on Parallel and Distributed Systems*, 28(5): 1361–1374.
- Xue, Y.; and Zhu, Z. 2021. Hybrid Flow Table Installation: Optimizing Remote Placements of Flow Tables on Servers to Enhance PDP Switches for In-Network Computing. *IEEE Transactions on Network and Service Management*, 18(1): 429–440.
- Zhang, Y.; Shan, K.; Li, X.; Li, H.; and Wang, S. 2023. Research and Technologies for next-generation high-temperature data centers – State-of-the-arts and future perspectives. *Renewable and Sustainable Energy Reviews*, 171: 112991.
- Zhu, H.; Kaffes, K.; Chen, Z.; Liu, Z.; Kozyrakis, C.; Stolica, I.; and Jin, X. 2020. RackSched: A Microsecond-Scale Scheduler for Rack-Scale Computers. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*, 1225–1240.