

ShuntServe: Cost-Efficient LLM Serving on Heterogeneous Spot GPU Clusters

Seungwoo Jeong^a, Moohyun Song^b, Juhyun Park^a, Kyungyong Lee^{a,*}

^aHanyang University, Department of Data Science, Seoul, 04763, Republic of Korea

^bHanyang University, Department of Artificial Intelligence, Seoul, 04763, Republic of Korea

Abstract

As large language model (LLM) services become widely adopted, the cost of GPU resources for serving these models in cloud environments has emerged as a critical concern. Spot instances offer up to 90% cost savings over on-demand instances, but their frequent interruptions and limited availability pose significant challenges for continuous LLM serving. GPU spot instances, in particular, exhibit lower and more volatile availability than CPU-based instances, making homogeneous clusters that depend on a single GPU type vulnerable to correlated failures. Heterogeneous clusters spanning multiple GPU types can address this by leveraging complementary availability patterns across diverse spot pools, yet existing LLM serving systems are designed for homogeneous environments and suffer from load imbalance when deployed on heterogeneous GPUs. This paper presents ShuntServe, a cost-efficient LLM serving system for heterogeneous spot GPU clusters. ShuntServe employs a profiling-free roofline model-based serving performance estimator and a dynamic programming-based model placement optimizer that jointly determines node configuration, parallelization strategy, and layer assignment to maximize throughput across heterogeneous GPUs. To enhance fault tolerance when using spot instances, ShuntServe combines output-preserving request migration with concurrent initialization via a shared tensor store, minimizing migration downtime by overlapping replacement node preparation with ongoing serving. Evaluation on Llama-3.1-70B and Qwen3-32B with a heterogeneous AWS cluster of L4, A10G, and L40S GPUs shows that ShuntServe achieves 1.42× and 1.35× higher throughput than state-of-the-art baselines and attains 31.9% and 31.2% cost efficiency improvements over on-demand instances for offline and online serving, respectively.

1. Introduction

Large language models (LLMs) become a fundamental component for a broad spectrum of applications, from conversational assistants and code generation to domain-specific reasoning [1]. While commercial LLM APIs are widely available, many organizations choose to deploy models independently due to prohibitive per-query costs at scale, data privacy and regulatory requirements, or the need for specialized fine-tuned models that public services cannot accommodate.

Running self-hosting LLM services, however, can become costly. LLM inference is GPU-intensive, and demand can fluctuate significantly, making on-premise clusters with large upfront capital expenditures difficult to justify. Cloud computing offers elastic GPU provisioning as an alternative, yet GPU instances are priced significantly higher than their CPU-only counterparts. For example, serving large models such as Llama-3-70B or Llama-3-405B on clusters of high-end GPUs can cost tens of thousands of dollars per month.

Spot instances can reduce these costs by up to 90% by leveraging providers' idle capacity [2, 3], but they may be reclaimed with only a short grace period, threatening service continuity. This risk is especially pronounced for GPU workloads, where empirical studies show that GPU spot instances exhibit significantly lower and more volatile availability than CPU-based

instances [4, 5, 6, 7]. Since spot instances are provisioned from the provider's surplus capacity for each specific instance type, their availability is inherently tied to the demand fluctuations of that type. Thus, relying on a single GPU type for a homogeneous cluster is inherently fragile, as correlated shortages can simultaneously disable all instances of that type. Heterogeneous clusters spanning multiple GPU types can instead exploit complementary availability patterns across diverse spot pools, improving both resilience and cost efficiency.

However, existing serving systems are mostly designed for homogeneous GPU environments, assuming identical model partitions and uniform workload distribution. In heterogeneous clusters, GPUs with disparate compute capabilities and memory bandwidths receive equal workloads, causing slower devices to bottleneck the pipeline. Moreover, the distinct characteristics of LLM inference, compute-bound prefill and memory-bound decode [8, 9, 10, 11], create decoupled bottlenecks across heterogeneous stages. This necessitates a serving system that jointly optimizes model placement, parallelization strategy, and layer assignment based on each device's unique capabilities.

To this end, we propose ShuntServe, a cost-efficient LLM serving system designed for heterogeneous spot GPU clusters. ShuntServe combines a profiling-free roofline model-based serving performance estimator with a dynamic programming algorithm to determine throughput-maximizing model placements that account for the distinct compute and memory characteristics of different GPU types. To handle spot GPU node interruptions, ShuntServe employs output-preserving request migration, which

*Corresponding author

Email address: kyungyong@hanyang.ac.kr (Kyungyong Lee)

recovers in-progress requests by recomputing their generated outputs on replacement instances rather than transferring state from the interrupted nodes. ShuntServe further minimizes migration downtime through concurrent initialization, which decouples the inference engine lifecycle from model weights loading via a shared tensor store. This allows the replacement node’s model loading and engine initialization to proceed in parallel with ongoing serving on the existing pipeline, significantly reducing the time required to resume full pipeline operation after an interruption.

We evaluate ShuntServe using Llama-3.1-70B and Qwen3-32B on a heterogeneous AWS GPU cluster comprising L4, A10G, and L40S instances with realistic workload traces. Through optimized heterogeneous model placement, ShuntServe achieves up to $1.31\times$ higher throughput than the state-of-the-art heterogeneous serving system and up to $1.43\times$ higher throughput than homogeneous baselines, without sacrificing request latency. By leveraging spot instances with its fault tolerance mechanisms, ShuntServe attains $1.21\times$ and $1.29\times$ cost efficiency improvements over on-demand instances for offline and online serving, respectively.

The key contributions of this paper are as follows:

- **Profiling-free serving performance estimator.** An analytical roofline model-based approach that estimates serving throughput across diverse parallelization configurations and heterogeneous GPU types without comprehensive profiling.
- **Dynamic programming-based model placement optimizer.** A dynamic programming algorithm augmented with beam search that jointly optimizes node configuration, parallelization strategy, and layer assignment to maximize throughput in heterogeneous clusters.
- **Fault tolerance for spot interruptions.** recomputation-based output-preserving request migration combined with concurrent initialization via a shared tensor store, minimizing migration downtime and maximizing spot utilization.

2. LLM Serving with Heterogeneous GPUs

The scale of modern Large Language Models has reached tens to hundreds of billions of parameters, rendering single-GPU serving infeasible [12, 13]. This challenge is typically addressed using two primary parallelization strategies of *Tensor Parallelism (TP)* and *Pipeline Parallelism (PP)*. TP partitions operators within each layer to run in parallel, which reduces request latency but is highly network-sensitive due to collective communications (e.g., AllReduce) [14]. In contrast, PP divides the model into stages, groups of layers, that are assigned to different devices [15]. While this does not reduce single-request latency, it increases overall throughput by processing micro-batches in parallel and is more broadly applicable under constrained network environments, as communication only occurs at stage boundaries.

Serving large-scale LLMs requires many GPUs that incurs significant operational costs. This economic constraint renders performance-only optimization insufficient, and new approaches




Figure 1: Time-series Graph of Available GPU Instance Counts by Type in AWS (us-west-2).

are necessary to co-optimize for high throughput and cost-effectiveness.

2.1. Cost Savings and Vulnerability of Spot Instances

Major cloud providers offer spot instances at discounts of up to 90% by utilizing idle compute resources [2, 3]. However, these instances are subject to interruption that a provider reclaims the instance after a short grace period when resources become scarce.

Because of the volatility, spot instances are suitable for stateless workloads, but using them for distributed LLM serving is challenging. Unlike traditional single-node inference that could often finish serving tasks within the grace period [16, 17], LLM serving is quasi-stateful that requires maintaining a persistent KV cache for long-running, autoregressive requests. In a distributed setting, this state creates a strong inter-node dependency.

Consequently, this stateful and distributed nature renders a simple finish and migrate strategy on a small-scale serving environment impractical. The interruption of a single node triggers a cascading failure that stalls the entire pipeline. This inherent fragility makes leveraging spot instances exceptionally difficult, mandating an effective fault tolerance mechanism.

2.2. Spot GPU Scarcity and Necessity of Heterogeneity

The spot availability of GPU instances is markedly lower than that of CPU-based instances and varies sharply across GPU types, with high-end accelerators particularly scarce [4, 5, 6, 7]. Figure 1 illustrates the hourly spot availability of GPU instance types in the AWS us-west-2 region from March 13 to 18, 2026 (UTC) where the y-axis represents the number of GPUs and the x-axis shows the timeline, high-end GPU instances like the p5.48xlarge (H100) are frequently unavailable that is consistent




Figure 2: Asymmetric model partitioning methods to mitigate performance bottlenecks in heterogeneous GPU clusters.

Table 1: Comparison of hardware specifications for various GPUs. FLOPS values are based on BF16 non-sparse operations.

GPU	Mem. (GB)	FLOPS (TFLOPS)	Mem. BW. (GB/s)
L4	24	121	300
A10G	24	70	600
L40S	48	362	864
A100	40	312	1,555
H100	80	989	3,350
B200	180	4,500	7,700

with prior reports on GPU scarcity [18]. Our 6-day monitoring revealed that H100 GPUs adequate for serving Llama-3-405B were available only 28.64% of the time, and even B200 GPUs were never successfully provisioned throughout the entire observation period. In contrast, Mid-tier instances (e.g., g6, g6e) exhibit different and often more stable availability patterns.

This high volatility and scarcity pose a significant operational challenge for systems relying on homogeneous clusters. Attempting to build a cluster from a single, high-end GPU type would nullify potential cost savings, as those instances are rarely available.

Heterogeneous GPU clusters are an effective solution to this limitation. Different instance types often exhibit complementary availability patterns; for example, Figure 1 shows g5/g6/g6e instances are available when p4d/p5/p6 instances are not. This approach enables cost savings by maximizing spot instance utilization from diverse pools, minimizing reliance on expensive on-demand instances.

2.3. Asymmetric Model Partitioning

Deploying distributed LLM serving on a heterogeneous GPU cluster introduces a new challenge for model partitioning. This arises from the disparities in compute capability (FLOPS) and memory bandwidth (Mem. BW.) across different GPU types, as shown in Table 1.

Uneven Layer Partitioning in Pipeline Parallelism. In homogeneous environments, a balanced partitioning scheme that assigns an equal number of layers per GPU is standard. However, applying this naive approach to a heterogeneous setting causes the least powerful GPU to become a performance bottleneck, severely limiting pipeline throughput. Figure 2a illustrates that the high-performance GPU (Stage 0) is forced to be idle while




Figure 3: Decoupled bottlenecks in the prefill and decode phase on a pipeline with heterogeneous capability GPUs.

waiting for the low-performance GPU (Stage 1). To resolve this, layers must be partitioned asymmetrically, proportional to GPU’s processing capability, to balance the execution time of each stage, reducing idle time and improving pipeline throughput (Figure 2b).

Per-Stage Tensor Parallelism. While uneven layer partitioning balances stages to improve throughput, the end-to-end request latency, which is the sum of all stage times, may remain high due to the inclusion of low-performance GPUs. A complementary approach is to configure multiple low-performance GPUs as a single logical stage using TP, thereby increasing the processing speed of a stage. For instance, as shown in Figure 2c, configuring four low-performance L4 GPUs with a TP degree of 4 achieves comparable performance to that of a single high-performance A100 GPU in Stage 0. This allows stages composed of weaker GPUs to balance their execution time against stages with high-performance GPUs, reducing pipeline bottlenecks and increasing overall throughput [19, 18].

The Optimization Challenge. Although uneven layer partitioning and per-stage tensor parallelism are effective partitioning strategies, determining the optimal model placement remains a complex challenge. LLM inference exhibits distinct performance characteristics where the prefill phase is typically compute-bound, and the decode phase is memory-bound [8, 9, 10, 11]. In a heterogeneous cluster, this behavior leads to decoupled bottlenecks, where the bottleneck stage differs between phases. As shown in Figure 3, in a four-stage pipeline of heterogeneous GPUs, Stage 0 is the bottleneck during prefill, while Stage 1 becomes the bottleneck during decode. Maximizing throughput requires jointly optimizing a vast set of factors, such as GPU types, pipeline depth, TP degree, and layer allocation. The re-




Figure 4: ShuntServe System Overview. The global server coordinates serving performance estimation, model placement optimization, and spot instance lifecycle management. Each GPU across pipelines runs an inference engine and a shared tensor store. Requests are dispatched to heterogeneous GPU pipelines via weighted round-robin.

sulting search space can become exponentially large, rendering a naive exhaustive search intractable.

3. ShuntServe Design Overview

Implementing cost-effective LLM serving on heterogeneous spot GPU clusters requires addressing two primary challenges. First, a spot instance interruption triggers a system-wide pipeline stall, resulting in significant downtime and loss of progress for ongoing requests. Second, determining the optimal model placement in a heterogeneous cluster is complicated by a vast search space of possible configurations, rendering profiling and exhaustive search methods intractable.

ShuntServe addresses these challenges through *output-preserving request migration* and *concurrent initialization* techniques (Section 4) that substantially reduce downtime, and a *partitioned model placement optimizer* (Section 5) that efficiently determines optimal configurations.

Figure 4 shows ShuntServe’s architecture. The global server hosts three key components: the serving performance estimator, model placement optimizer, and instance manager. The serving performance estimator predicts performance for candidate configurations, where the optimizer uses to determine the best model placement deployed through the instance manager. The instance manager monitors spot interruption events and availability fluctuations to handle GPU instance lifecycle. Each deployed pipeline consists of heterogeneous GPU instances, either spot or on-demand, with each node running an inference engine and a




Figure 5: KV cache transfer vs recomputation latency on Llama 3 family (3B, 8B, 70B).

shared tensor store. Incoming requests are scheduled to pipelines via weighted round-robin based on their throughputs. When spot instance interruption events occur, the instance manager invokes concurrent initialization and output-preserving request migration to handle the events.

4. Handling Spot Interruption

Effectively solving fault tolerance challenge raised from spot instance interruption requires addressing two specific sub-problems. First is how to preserve the state of ongoing inference requests. In autoregressive decoding, simply discarding the generated tokens and KV cache wastes all prior computation and results in the loss of the user’s generated output. Second is how to rapidly reconfigure the pipeline with new resources substituting interrupted ones. ShuntServe addresses these sub-problems through *output-preserving request migration* and *concurrent initialization* based on a shared tensor store.

4.1. Output-Preserving Request Migration

Rather than simply discarding the generated tokens and KV cache, two primary approaches exist to migrate requests while preserving their computation progress. The first is KV cache transfer-based migration, which transfers the KV cache from the interrupted instances to new replacement instances, allowing the receiving instances to resume inference requests without recomputation. The second is recomputation-based migration, which routes the interrupted requests’ input sequence and already-generated output sequence to a pipeline containing new instances to reconstruct the KV cache.

Figure 5 shows the latency overhead of the two approaches measured on AWS g6.xlarge (L4) and g6e.xlarge (L40S) instances across varying context lengths for the Llama 3 family with 3B, 8B and 70B model sizes. The x-axis represents the context length, and the y-axis represents latency on a log scale. For the 70B model, since all layers cannot fit on a single instance, we measured the latency of a subset of layers and normalized it to a per-layer basis. The results show that recomputation-based migration exhibits lower latency overhead in most cases, demonstrating the advantage of the recomputation-based approach.

On the other hand, KV cache transfer can be effective in certain environments (i.e., very large models, very long contexts, GPUs with limited compute capacity, and high-bandwidth networks). For example, in the case of Llama-3-70B (Figure 5c),




Figure 6: Output-preserving request migration. In b^j_t , j denotes the micro-batch index, and t indicates the iteration of the decoding step.

recomputation exhibited higher latency than transfer across all context lengths on L4. On L40S, recomputation showed lower latency up to 32K context length, but at 64K it incurred a recomputation latency of 7.59s compared to 6.92s for transfer, resulting in approximately 9.6% higher latency. However, in spot instance environments, KV cache transfer introduces challenges beyond latency, particularly in terms of fault tolerance. If an interruption occurs during KV cache transfer, the system must ultimately fall back to recomputation, incurring the overhead of both approaches. Moreover, spot interruption environments impose a strict constraint that the entire process, from new instance initialization to transfer completion, must be completed within the grace period, rendering the transfer-based approach impractical for spot environments. Therefore, ShuntServe adopts the recomputation-based approach, which is practical under spot instance constraints.

Figure 6 illustrates the process of ShuntServe’s recomputation-based output-preserving request migration.

4.2. Minimizing Downtime via Concurrent Initialization

Selectively replacing only the interrupted node while continuing inference on the unaffected nodes reduces resource waste and downtime. However, this partial pipeline replacement introduces technical challenges unsolved by existing LLM serving systems. When replacing only part of a pipeline, the replacement node must be added to the cluster and the distributed inference engine must be re-initialized with the unaffected nodes.

The core problem is that existing LLM serving systems, such as vLLM [20] and TensorRT-LLM [21], tightly couple the inference engine lifecycle with model weight management. Ideally, the new pipeline incorporating the replacement node would be initialized while the existing pipeline continues serving requests, requiring both old and new engine processes to temporarily co-exist on each unaffected node. However, because each engine process independently loads its own copy of the model weights and allocates KV cache space in GPU memory, this coexistence results in duplicate memory allocation that triggers an out-of-memory (OOM) error. Consequently, the new pipeline can only start after the old one is fully terminated, precluding any overlap between preparation and serving and incurring significant downtime.

Shared Tensor Store. To resolve inefficiencies arising from the coupled lifecycle, ShuntServe manages model weights in a separate *Shared Tensor Store* process, decoupling the lifecycle of the




Figure 7: Concurrent initialization with shared tensor store.

inference engine from that of the model. In the *Shared Tensor Store*, model weight tensors and KV cache space are shared with the inference engine process in a zero-copy manner through CUDA IPC handles, allowing the engine to directly reference the tensors during inference. This eliminates copy and communication overhead during inference, and prevents duplicate memory allocation upon re-initialization, as a new engine process can be attached to the existing store.

Figure 7 illustrates the concurrent initialization in ShuntServe. When the global server detects a spot interruption, it adds the new instance to the cluster and starts a *Shared Tensor Store* process on the new instance. It simultaneously loads the model partition from remote storage and initializes inference engines across the new instance and the unaffected instances. This preparation occurs in the background while the existing pipeline continues serving.

Once the initialization of the new pipeline is complete, the global server stops routing requests to the old pipeline, terminates its inference engine, and shunts the request path to the new pipeline. Since the new pipeline initialization overlaps with ongoing serving, the transition is achieved with near-zero downtime when the initialization completes within the grace period. Note that even when the initialization time exceeds the grace period, downtime is incurred only for the portion that extends beyond the grace period.

5. Model Placement for Heterogeneous GPUs

This section addresses the method for finding the optimal model layer placement within the vast search space of a heterogeneous GPU cluster with the parallelization strategies.

5.1. Profiling-free Serving Performance Estimator

To make an optimal model layer placement on heterogeneous GPU devices, it is important to estimate the performance on disparate devices. Prior studies have relied on thorough empirical profiling to predict the performance of LLM inference systems [22, 18, 23]. However, the factors influencing throughput, such as model parameter size, sequence length, batch size, and GPU type, create an extensive combinatorial space. This combinatorial explosion is further exacerbated when considering parallelization strategies of TP and PP, making the cost of exhaustive profiling intractable. Consequently, an approach

that minimizes or eliminates profiling is essential for efficiently finding an optimal, cost-effective LLM serving configuration.

5.1.1. Static Latency Estimation

To accurately calculate system throughput, one must first estimate inference latency. Following the roofline model [24, 25], which dictates that an operation’s latency is determined by the maximum of its compute latency and memory scan latency. Equation 1 formulates the latency of each operation based on its arithmetic intensity. The total end-to-end inference latency for a given device is then computed by summing the latencies of all operations it must perform.

$$\begin{aligned} L_{compute} &= \frac{FLOPs}{FLOPS} \\ L_{memory} &= \frac{MemScanCost \cdot E}{Mem. BW.} \\ L_{ops} &= \max(L_{compute}, L_{memory}) \end{aligned} \quad (1)$$

The operations are defined based on Llama 3 [1], a widely adopted open-source generative LLM. We formulate the three components in which most computations occur: (1) Attention, (2) Feed Forward Network, and (3) Output Embedding. These components are the de-facto standard of transformer-based non-MoE generative LLMs, and the formulations therefore generalize to other models that adopt the same structure. Table 2 details the FLOPs and memory scan cost required for each defined operation. This formulation-based static latency estimation derives latency for any pipeline configuration through fast calculation, applying flexibly to new configurations without requiring prior profiling. Table 3 summarizes the notation used throughout this paper.

To estimate the end-to-end inference latency in a distributed serving system, network latency should be considered alongside the compute latency. ShuntServe models the network latency for pipeline communication using the α - β model [26, 27], which estimates communication cost as the sum of a constant network latency α and the data transfer time proportional to message size over bandwidth β , and the communication topology for TP operations, e.g., AllReduce/AllGather, is modeled using a Ring-based [28] approach. These communication latency costs are factored into the calculation for each stage and phase.

In the Ring-based topology, a single *AllReduce* operation consists of one *ReduceScatter* and one *AllGather* operation, where each operation transmits data of size N/P for $P - 1$ iterations. This can be expressed as $(\alpha + \frac{N/P}{\beta})(P - 1)$. Here, α and β denote network latency and bandwidth, respectively, N represents the size of the original data to be shared, and P is the number of processes participating in the communication, which corresponds to D_{TP} in TP communication. These AllReduce operations occur twice per Transformer layer, and we ultimately calculate the PP and TP communication latencies as follows:

$$L_{PP} = \alpha_{inter} + \frac{BSHE}{\beta_{inter}} \quad (2)$$

$$L_{TP} = 4 \cdot (\alpha_{intra} + \frac{BSHE}{D_{TP} \cdot \beta_{intra}}) \cdot (D_{TP} - 1) \cdot l \quad (3)$$

α_{inter} and β_{inter} denote the inter-stage (e.g., Ethernet, InfiniBand) network latency and bandwidth, and α_{intra} and β_{intra} denote the intra-stage (e.g., PCIe, NVLink) network latency and bandwidth.

5.1.2. Throughput Modeling

While the latency of a single request is critical, the total number of requests the system can process per unit of time, throughput, is an equally important performance metric for a production LLM serving system. This is often measured in Requests per Second (RPS) and can be expressed as Equation 4.

$$RPS = \frac{B}{\sum L_{ops}(B)} \quad (4)$$

ShuntServe models the throughput of a heterogeneous pipeline by first using the static latency estimator to calculate the distinct prefill and decode latencies for each stage i . The overall throughput is then derived by computing the latency using Equation 5 and then combining it with Equation 4.

$$L(B) = \max_{0 \leq i < D_{PP}} (L_{prefill}^i(B)) + \max_{0 \leq i < D_{PP}} (L_{decode}^i(B)) \quad (5)$$

5.2. Partitioned Model Placement Optimizer

ShuntServe determines the optimal model placement using the throughput modeling described in the previous section. The most intuitive approach, an exhaustive search over all possible placements, is computationally intractable as the search space is exponentially large. To address this, ShuntServe employs a dynamic programming (DP) algorithm augmented with beam search that efficiently navigates this search space. In this section, we describe the details of the algorithm, the rationale for its batch size determination, and the optimization objective.

5.2.1. Model Placement with Dynamic Programming

A key capability of ShuntServe’s approach is that the operation-level latency estimation and throughput modeling from Section 5.1 enable evaluating partial model placements within DP subproblems, exploiting the optimal substructure property. Following prior studies [18, 19, 10] and our empirical experience, tensor parallelism is restricted to intra-node GPUs, as inter-node network latency and bandwidth is insufficient. Furthermore, in the spot environment, ShuntServe assigns each instance exclusively to a single pipeline, as sharing an instance across multiple pipelines might cause failures where the impact of interruption propagates to all pipelines containing it.

Algorithm 1 formalizes this optimization approach. The problem of placing l layers across $s + 1$ stages is decomposed into a subproblem of placing $l' < l$ layers across s stages and the assignment of the remaining $l - l'$ layers to a new stage with a

Table 2: FLOPs and memory scan requirements for generative LLM inference.

Operation	Phase	FLOPs	Memory Scan Cost
QKV	Prefill	$B(2S_{in}H^2 + 4S_{in}HH_{kv})/D_{TP}$	$BS_{in}H + (H^2 + 2HH_{kv})/D_{TP}$
Projection	Decode	$BS_{out}(2H^2 + 4HH_{kv})/D_{TP}$	$S_{out}(BH + (H^2 + 2HH_{kv})/D_{TP})$
Attention	Prefill	$4BS_{in}^2H/D_{TP}$	$(BS_{in}H + 2BS_{in}H_{kv})/D_{TP}$
	Decode	$\sum_{t=1}^{S_{out}} 4(s_{in} + t)H/D_{TP}$	$\sum_{t=1}^{S_{out}} (BH + 2B(S_{in} + t)H_{kv})/D_{TP}$
Output	Prefill	$2BS_{in}H^2/D_{TP}$	$(BS_{in}H + H^2)/D_{TP}$
Projection	Decode	$2BS_{out}H^2/D_{TP}$	$S_{out}(BH + H^2)/D_{TP}$
Up and Gate	Prefill	$4BS_{in}HH_{ffn}/D_{TP}$	$BS_{in}H + 2HH_{ffn}/D_{TP}$
Projection	Decode	$4BS_{out}HH_{ffn}/D_{TP}$	$S_{out}(BH + 2HH_{ffn}/D_{TP})$
Down	Prefill	$2BS_{in}HH_{ffn}/D_{TP}$	$(BS_{in}H_{ffn} + HH_{ffn})/D_{TP}$
Projection	Decode	$2BS_{out}HH_{ffn}/D_{TP}$	$S_{out}(BH_{ffn} + HH_{ffn})/D_{TP}$
Logits	Prefill	$2BS_{in}H\mathcal{V}/D_{TP}$	$BS_{in}H + H\mathcal{V}/D_{TP}$
Calculation	Decode	$2BS_{out}H\mathcal{V}/D_{TP}$	$S_{out}(BH + H\mathcal{V}/D_{TP})$

Table 3: Notations

Symbol	Description
L	Latency
B	Batch size
S_{in}	Input sequence length
S_{out}	Output sequence length
H	Hidden dimension
H_{kv}	Key/Value head dimension
H_{ffn}	Intermediate dimension in FFN
$\mathcal{V}$	Vocabulary size
D_{PP}	Pipeline parallelism degree
D_{TP}	Tensor parallelism degree
l	Number of transformer layers
E	Element size in bytes
t	Iteration of current step

specific instance type (lines 2–4). For each prior configuration in $DP[l'][s]$, the configuration is extended with each available instance type (lines 8–9). The maximum batch size is then computed for the resulting placement under memory constraints (line 10), and the throughput is estimated using the serving performance estimator (line 11). The DP table $DP[l][s+1]$ is updated if the new configuration achieves a higher score (line 12), and after all layers are placed, the optimal placement is extracted from $DP[N_L][*]$ (line 17).

Even with a DP approach, the search space for model placement can grow exponentially with the number of configurable stages and instance types. The number of stages can, in theory, be as high as the total number of layers, N_L , and each stage can be assigned one of $|I|$ instance types. When including the combinations for layer allocation, the time complexity of exhaustive DP implementation reaches $O(N_L^3 \cdot N_S! / \prod_{i \in I} \binom{N_S}{|i|})$, which becomes computationally infeasible. To manage this, ShuntServe incorporates beam search, retaining only the top- k partial placement candidates in each $DP[l][s]$ entry (line 7). This optimization prunes the search space, dramatically reducing the time complexity to $O(N_L^3 \cdot k \cdot |I|)$. ShuntServe employs this efficient optimization process iteratively, allowing it to greedily extract

Algorithm 1 Partitioned Model Placement Optimizer

Require: Beam size k (default 1), instance types I , total layers N_L

Ensure: Model Placement $\mathcal{M}$

- 1: Initialize $DP[l][s] \leftarrow \emptyset$ for all $l \in [0, N_L], s \in [0, N_L]$
- 2: **for** $l = 1$ **to** N_L **do**
- 3: **for** $l' = 0$ **to** $l - 1$ **do**
- 4: $l_{new} \leftarrow l - l'$
- 5: **for** $s = 0$ **to** l' **do**
- 6: $s_{new} \leftarrow s + 1$
- 7: $C \leftarrow$ Top- k stage combinations from $DP[l'][s]$
- 8: **for** $(c, i) \in C \times I$ **do**
- 9: $c_{new} \leftarrow c.add_stage(i, l_{new})$
- 10: Compute maximum batch size B_{max} for c_{new}
- 11: Evaluate throughput of c_{new}
- 12: Update $DP[l][s_{new}]$ with c_{new} if better
- 13: **end for**
- 14: **end for**
- 15: **end for**
- 16: **end for**
- 17: $\mathcal{M} \leftarrow \arg \max_{c \in DP[N_L][*]} objective(c)$

the desired number of pipeline configurations to populate the serving system.

5.2.2. Batch Size Determination

In LLM serving, increasing the batch size significantly improves throughput. This efficiency stems from a disparity between the inference phases. Latency in the compute-bound prefill phase scales linearly with batch size, whereas latency in the memory-bound decode phase exhibits minimal growth, as multiple requests in a batch share the same weight matrix and require only a single memory access, making batching highly efficient in the decode phase. Since the decode phase typically dominates the overall request latency [9, 20, 29, 30], this batching efficiency directly translates into throughput improvements. Thus, ShuntServe maximizes serving throughput by setting the largest possible batch size B that satisfies the memory constraint

of a given model placement, as in Equation 6.

$$B = \min_{0 \leq i < D_{PP}} (\lfloor \frac{M^i - M_{\text{weight}}^i - M_{\text{activation}}}{N_l^i \cdot (S_{\text{in}} + S_{\text{out}})} \rfloor) \quad (6)$$

In the equation, M^i , M_{weight}^i , and $M_{\text{activation}}$ denote the total memory capacity of stage i , the model weight size, and the maximum activation tensor size during a forward pass, respectively. Using the determined batch size B , the model placement optimizer evaluates the maximum throughput of the given placement.

5.2.3. Optimization Objective

Using throughput as the sole optimization objective is problematic. Adding instances increases the available memory, which in turn allows for larger batch sizes and higher throughput. This creates an incentive for the algorithm to indiscriminately add instances, regardless of cost-efficiency. To prevent this, ShuntServe defines its objective function as throughput per cost, thereby evaluating the cost-effectiveness of each potential instance addition. We formulate this as:

$$\max(\frac{\text{Throughput}}{\text{Cost}} \cdot (1 - \gamma \cdot \max(0, \frac{\text{latency}}{SLO} - 1))) \quad (7)$$

We primarily optimize for throughput per cost, but include a soft penalty term for latency to satisfy the desired SLO. The parameter γ controls the sensitivity to this latency penalty; setting $\gamma = \infty$ enforces it as a hard constraint.

6. Implementation

The ShuntServe inference engine was implemented by modifying vLLM [20] v0.8.1. The modifications comprise 1,077 lines of Python code. These modifications include changes to the communication methods to support asymmetric parallelism and code for fetching model weights from the Shared Tensor Store. The Shared Tensor Store itself, along with the global server and scheduler for managing pipelines, were written in 10,363 lines of Python code. The scheduler implements load balancing via a weighted round-robin policy based on the throughput of each pipeline.

The global server uses Ray [31] to manage physical instances, handling their participation in the cluster and subsequent eviction. Model loading is handled by downloading from a remote object storage server. To minimize loading time, model tensors are sharded into individual binary files in a custom raw binary format. In our observation, `torch.save()` retains the full original tensor size when saving sliced tensors, resulting in significant network overhead for downloading partitioned shards. The custom binary format eliminates this overhead by storing only the actual partition data. This design allows each distributed node to download only its required partition via multi-threading and load it directly onto the GPU.

7. Evaluation

This section comprehensively evaluates the key contributions of ShuntServe. The evaluation is structured into two main parts: an assessment of the model placement optimizer’s performance on heterogeneous GPU clusters (Section 7.1), and an evaluation of the fault tolerance mechanism’s effectiveness under spot interruptions (Section 7.2).

Model and Cluster Setup. ShuntServe is evaluated using the half-precision (BF16) Llama-3.1-70B-Instruct [1] and Qwen3-32B [32] models. The evaluation cluster consists of heterogeneous AWS GPU instances: three g6.12xlarge instances (each with 4x L4 24GB GPUs), two g5.12xlarge instances (each with 4x A10G 24GB GPUs), and four g6e.xlarge instances (each with 1x L40S 48GB GPU). This provides a total of 24 GPUs and 672GB of GPU memory. The evaluation experiments were conducted on AWS at a total cost of approximately \$3,000, excluding costs for development and debugging.

Request Scenario. We use the Azure Conversation Dataset [9], a production trace from Azure LLM inference services, which spans one hour with fluctuating request arrival patterns and input/output length distributions. We prune requests with input lengths exceeding 2048 tokens, as peak context lengths in the dataset cause OOM errors on homogeneous baselines whose memory of each pipeline is limited, rendering comparison infeasible. The resulting dataset has an average input length of 763 and an average output length of 232, and an average request rate of 4.67 req/s. As also noted in [18], this distribution is more challenging than those used in prior work [20] due to its longer request lengths.

Metrics. We evaluate serving performance under two workload types. In offline workload, where the system is fully saturated with requests, we measure throughput in terms of Requests per Second (RPS). In online workload, where requests arrive dynamically and the system processes them in real time, we measure request latency using Time to First Token (TTFT), which captures the duration from request dispatch through prefill to the generation of the first output token, Time per Output Token (TPOT), which measures the average generation time of each token during the decode phase, and end-to-end latency, which captures the total duration from request dispatch to completion. In the fault tolerance evaluation, we additionally compare the monetary cost for operating each system. We define cost efficiency as cost per throughput for offline workloads and as request latency $\times$ cost for online workloads.

7.1. Model Placement Evaluation

We evaluate the model placement optimizer of ShuntServe, which determines the optimal node assignment and layer partitioning for inference pipelines across a heterogeneous GPU cluster. We first assess the accuracy of the static latency estimation technique, a foundational component of the optimizer (Section 7.1.1). We then evaluate the performance of the model placement decisions produced by ShuntServe under offline (Section 7.1.2) and online (Section 7.1.3) workloads. Finally, we




Figure 8: Comparison of estimated and measured inference latency for Llama-3.1-70B-Instruct across parallelism configurations (PP $\times$ TP), GPU types (A10G, L4, L40S), and batch sizes (BS). Each row corresponds to a parallelism strategy, each column to a batch size, and the x-axis within each subplot represents the GPU type. Red crosses indicate out-of-memory (OOM) cases.

conduct a sensitivity analysis on the beam search width k with respect to optimizer execution time and the quality of the resulting placement (Section 7.1.4).

7.1.1. Serving Performance Estimation Accuracy

To validate the accuracy of ShuntServe’s roofline model-based performance estimator, we compare predicted inference latency against measured values on three GPU instance types (g5.48xlarge with 8 A10G, g6.48xlarge with 8 L4, and g6e.48xlarge with 8 L40S) across four parallelism strategies (PP1/TP8, PP2/TP4, PP4/TP2, PP8/TP1) and six batch sizes (1, 2, 4, 8, 16, 32). The input and output sequence lengths are set to 763 and 232, respectively, matching the average of the evaluation dataset. To isolate GPU execution performance from CPU scheduling overhead, we use the gptBench benchmark from TensorRT-LLM [21], which employs static batching with automatically optimized kernels. We adopt Mean Absolute Percentage Error (MAPE) as the accuracy metric across all feasible configurations, excluding cases where the configuration results in out-of-memory (OOM) errors.

Figure 8 presents the comparison across all GPU types, parallelism strategies, and batch sizes. The estimator achieves an average MAPE of 3.49% on L4, 6.68% on L40S, and 9.67% on A10G, yielding an overall MAPE of 6.63% across 54 feasible configurations. The minimum error is 0.03% (PP1/TP8, batch size 32, L40S) and the maximum is 23.72% (PP2/TP4, batch size 32, L40S). This variance arises because the static latency estimation assumes linear scaling of each operation, whereas actual GPU execution involves non-linearity that the model does not fully capture. We further discuss this limitation and potential refinements in the Discussion (Section 8).

Despite this variance, the achieved overall MAPE of 6.63% demonstrates sufficient accuracy as the foundational component




Figure 9: Offline throughput comparison under pipeline configurations determined by each baseline system’s model placement algorithm. (a) shows results for Llama-3.1-70B, and (b) shows results for Qwen3-32B.

of the model placement optimizer, while requiring only a one-time hardware calibration (FLOPS, memory bandwidth, and network bandwidth) per GPU type rather than per-configuration profiling.

7.1.2. Offline Serving Throughput across Placement Algorithms

We evaluate the model placement optimizer of ShuntServe by comparing against three baseline systems. The first is HEXGEN [19], a state-of-the-art system for heterogeneous GPU clusters that uses a genetic algorithm and dynamic programming for model placement. The second is AlpaServe [23], a state-of-the-art system for homogeneous GPU clusters that uses dynamic programming to equalize stage latencies. The third, vLLM [20], is a widely adopted serving framework that applies even layer partitioning across homogeneous GPU clusters.

To measure the throughput under saturated conditions while prohibitive experimental cost, we use the first 5 minutes of the trace for Llama-3.1-70B (1,309 requests) and the first 20 minutes

for Qwen3-32B (5,299 requests), as the smaller model requires a longer trace to maintain saturation. Figure 9 presents the results.

For Llama-3.1-70B (Figure 9a), ShuntServe achieves the highest throughput of 1.53 req/s, outperforming all baselines. In particular, ShuntServe improves throughput by 1.17 $\times$ over HexGen, the state-of-the-art heterogeneous baseline. This improvement is attributed to the DP-based optimizer of ShuntServe, which explores the layer distribution space using roofline-based throughput estimation that accounts for the distinct characteristics of prefill and decode phases. In contrast, HexGen explores layer distributions through local perturbation within its genetic algorithm, which can become trapped in local optima when GPU capabilities are not ordered along the pipeline. Although HexGen mitigates this by distributing layers proportionally to the memory capacity of each stage, memory capacity does not necessarily correlate with GPU compute performance, leaving the resulting placement suboptimal.

Nevertheless, in Llama-3.1-70B, HexGen and ShuntServe, both heterogeneous systems, outperform the homogeneous baselines (AlpaServe and vLLM). This advantage stems from the ability of heterogeneous systems to flexibly compose pipelines from diverse GPU types, allowing each pipeline to retain sufficient memory for KV cache. For example, ShuntServe and HexGen each construct two pipelines for the cluster. In contrast, AlpaServe and vLLM construct three pipelines by grouping each GPU type separately, leaving insufficient memory for KV cache in each pipeline.

For Qwen3-32B (Figure 9b), the results show a different trend. ShuntServe still achieves the highest throughput at 4.59 req/s, but the improvement over vLLM, which applies the intuitive even-partitioning placement in a homogeneous configuration, is only 1.05 $\times$, compared to 1.43 $\times$ for Llama-3.1-70B. This narrow gap arises from the evaluation cluster providing sufficient memory capacity for each pipeline even under homogeneous configurations for the smaller Qwen3-32B. In contrast, HexGen initializes pipeline groups based on communication topology, inherently producing heterogeneous configurations that persist through subsequent mutations. As a result, HexGen yields the lowest throughput among all baselines at 3.39 req/s. AlpaServe achieves 4.14 req/s, slightly below vLLM, as its SLO-oriented design selects additional replication. This increases the number of pipelines but reduces the batch size of each pipeline.

This contrast with the Llama-3.1-70B results indicates that when the cluster provides sufficient resources relative to the model size, an intuitive homogeneous configuration can be effective. ShuntServe’s optimizer reflects this adaptability. For Qwen3-32B, it constructs each pipeline homogeneously to minimize decoupled bottlenecks, and differs from vLLM only in selecting replication over pipeline parallelism for the g5.12xlarge instance group based on throughput per cost.

7.1.3. Analysis of Online Serving Latency

We next evaluate request latency under online workload to confirm that the throughput improvements do not sacrifice request latency. For Llama-3.1-70B, we replay the first 3 minutes of the trace over a 15-minute window at approximately 0.7 req/s, and for Qwen3-32B, the first 9 minutes over a 15-minute window




Figure 10: Online serving latency comparison under pipeline configurations determined by each baseline system’s model placement algorithm. For each system, the solid box represents Qwen3-32B, while the hatched box represents Llama-3.1-70B. The box spans the 25th to 75th percentile, the whiskers indicate the 10th and 90th percentile, and the yellow line denotes the median. The annotated multipliers denote the factor of increase in median latency when scaling from Qwen3-32B to Llama-3.1-70B for each system.

at approximately 2.4 req/s. These arrival rates remain within the serving capacity of all baselines, ensuring a fair latency comparison without distortion from queuing delays. Figure 10 presents the TTFT (Figure 10a) and TPOT (Figure 10a) distributions.

In these results, ShuntServe achieves a median TTFT of 1.37s and TPOT of 0.82s for Llama-3.1-70B, and 0.65s TTFT and 0.50s TPOT for Qwen3-32B. Compared to vLLM, which applies the intuitive even-partitioning placement, ShuntServe shows 19% lower median TTFT and 22% higher median TPOT for Llama-3.1-70B, and 2% lower TTFT and 6% higher TPOT for Qwen3-32B. At the P90 tail latency, ShuntServe reduces TTFT by 69% for Llama-3.1-70B (2.39s vs 7.70s) while maintaining comparable TPOT (0.93s vs 0.97s).

Among the baselines, HexGen exhibits the highest latency across both models. In particular, the median TPOT for Llama-3.1-70B reaches 1.31s, approximately 60% higher than ShuntServe. This stems from HexGen’s genetic algorithm, which favors expanding the pipeline parallelism dimension to equalize processing capability across stages. For instance, HexGen assigns the four L4 GPUs in a g6.12xlarge instance to four individual pipeline stages with TP1, whereas other baselines group them as a single stage with TP4. While this deep pipeline increases batch size and enables concurrent processing of multiple requests through micro-batching, it fails to exploit parallel computation through intra-node high-bandwidth interconnects, resulting in significant latency degradation.

AlpaServe achieves the lowest TPOT among all baselines for Qwen3-32B at 0.22s. AlpaServe selects the most aggressive replication among the baselines, which distributes requests across more pipelines with smaller batch sizes per pipeline. This is advantageous under online workloads that do not saturate the system. However, this strategy allocates more GPU memory to model weights rather than KV cache, which degrades throughput. Furthermore, as model size increases, the advantage of AlpaServe’s placement diminishes. As shown by the annotated multipliers in Figure 10, ShuntServe’s median TPOT increases by 1.6 $\times$ when scaling from Qwen3-32B to Llama-3.1-70B, whereas




Figure 11: Effect of beam search width k on algorithm execution time and the quality of the resulting model placement for Llama-3.1-70B and Qwen3-32B. The left y-axis represents the algorithm execution time on a log scale, and the right y-axis represents the estimated throughput of the resulting placement. (a) shows results on the evaluation cluster (3 GPU types, 9 instances, 24 GPUs), and (b) shows results on the large cluster (7 GPU types, 15 instances, 76 GPUs).

AlpaServe’s increases by 3.6 $\times$. This is because larger models impose tighter memory capacity constraints that limit the degree of replication, reducing the effectiveness of AlpaServe’s placement strategy. These results confirm that ShuntServe’s throughput-oriented placement does not sacrifice request latency, while exhibiting stable scaling behavior as model size increases.

7.1.4. Algorithm Execution Time of Beam Search

ShuntServe employs beam search to manage the combinatorially explosive search space of the DP-based optimizer, retaining only the top- k candidate placements at each DP entry. This experiment evaluates the trade-off between the algorithm execution time and the quality of the resulting model placement as the beam size k varies. Two cluster configurations are used: (1) the Evaluation cluster, consisting of 3 GPU types, 9 instances, and 24 GPUs used in our main experiments, and (2) the Large cluster, comprising 7 GPU types across 15 instances with 76 GPUs, configured with one instance of each of the 15 GPU instance types available on AWS. The algorithm execution time depends on the CPU used for optimization, with all measurements in this experiment conducted on an m8a.8xlarge instance.

Figure 11a presents the results on the evaluation cluster. In terms of execution time, the algorithm completes in approximately 15.90s for Llama-3.1-70B and 11.49s for Qwen3-32B at $k = 1$, increasing linearly up to $k = 8$. Beyond $k = 11$, execution time converges to approximately 94s for Llama-3.1-70B and 75s for Qwen3-32B, as the number of feasible candidate model placements in the evaluation cluster becomes smaller than k , eliminating the pruning of beam search. Figure 11b presents the results on the large cluster. Execution time grows linearly with k throughout the measured range up to $k = 32$, reaching 9,433s for Llama-3.1-70B and 6,418s for Qwen3-32B, as the number of possible stage configurations at each DP entry increases combinatorially with the number of GPU types and instances in the cluster.

In terms of placement quality, the throughput of the resulting placement is largely insensitive to k . For Llama-3.1-70B on the evaluation cluster, throughput increases from 2.382 req/s at $k = 1$ to 2.830 req/s at $k = 3$, where it reaches a plateau. On the large cluster, throughput varies by only approximately




Figure 12: Spot instance availability over time for three GPU instance types under the evaluation scenario.

0.1% across all k values. For Qwen3-32B, the optimizer finds identical placements across all k values on the evaluation cluster at 9.424 req/s, and exhibits similarly negligible variation on the large cluster. These results indicate that the DP-based optimizer of ShuntServe effectively identifies near-optimal placements even at small k values, as the partial placements in the DP table closely satisfy the optimal substructure property. Based on these findings, we set $k = 3$ in our main experiments to balance algorithm execution time and placement quality.

7.2. Enhanced Tolerance to Spot Interruptions

We evaluate ShuntServe’s fault tolerance mechanisms for handling spot instance interruptions, namely output-preserving request migration and concurrent initialization. To ensure consistent evaluation across all experiments, we extract a 50-minute spot availability scenario by tracking the availability of each instance type and simulating the corresponding spot events. Figure 12 illustrates this scenario. The scenario is extracted from a 6-day trace spanning March 13–18, 2026, which corresponds to the same observation period as Figure 1. To select the most challenging scenario, we assign each candidate window a composite score based on the frequency of availability change events and the magnitude of instances affected by each event, and select the window with the highest score. This ensures that the evaluation reflects worst-case conditions among the observed traces. We further observe that approximately 40.4% (688 of 1,701) of the candidate trace windows have a score of zero, indicating no availability changes. In such scenarios, using spot instances directly reduces cost without any interruption overhead.

Baselines. To the best of the authors’ knowledge, ShuntServe is the first system to leverage spot instances for LLM serving on heterogeneous GPU clusters. Although SpotServe [33] addresses LLM serving on spot instances, it targets homogeneous GPU clusters, and its core fault tolerance mechanism based on the Kuhn-Munkres algorithm for device mapping does not generalize to heterogeneous pipeline configurations. SkyServe [34] addresses spot instance utilization for AI serving in cloud environments. Yet, it considers higher-level concerns such as instance provisioning, placement, and scheduling, not the fault tolerance of distributed LLM serving. We therefore evaluate ShuntServe’s spot tolerance against the following baselines. *On-demand Only* uses exclusively on-demand instances without any




Figure 13: Offline serving throughput under a spot availability scenario. The y-axis represents throughput in Requests per Second (req/s). (a) shows results for Llama-3.1-70B, and (b) shows results for Qwen3-32B.

interruptions. *No Handle* uses spot instances but applies no fault tolerance mechanism upon interruption. *Request Migration* implements output-preserving request migration on vLLM without concurrent initialization. *Concurrent Initialization* implements concurrent initialization on vLLM without request migration. ShuntServe combines both mechanisms.

7.2.1. Serving Throughput under Spot Interruptions

We evaluate the impact of spot interruptions on offline workload. The entire dataset is dispatched to the system at once. After a 50-minute spot availability window, the system is halted and the serving metrics are collected from the requests completed within the window. Figure 13 presents the results.

For Llama-3.1-70B (Figure 13a), as expected, On-demand Only achieves the highest throughput at 1.23 req/s since it is not affected by spot interruptions, while No Handle yields the lowest at 0.85 req/s. Request Migration improves throughput by approximately 8% over No Handle, reaching 0.92 req/s, by recovering the computation of in-flight requests that would otherwise be discarded upon interruption. Concurrent Initialization yields a larger improvement of approximately 21% to 1.03 req/s by reducing the pipeline downtime through overlapping replacement node preparation with ongoing serving. ShuntServe, which combines both mechanisms, achieves 1.12 req/s, a 32% improvement over No Handle and the highest throughput among the baselines that utilize spot instances.

For Qwen3-32B (Figure 13b), the results follow a similar trend, but the gap between the baselines that utilize spot instances and On-demand Only narrows considerably. Specifically, the gap between On-demand Only and No Handle decreases from approximately 45% for Llama-3.1-70B to approximately 13% for Qwen3-32B (i.e., 5.19 req/s vs 4.60 req/s). This difference arises from the characteristic of distributed LLM serving where the interruption of a single instance causes the entire pipeline to stall. Larger models require pipelines with more nodes, making them more vulnerable to this cascading effect. Conversely, smaller models require fewer nodes per pipeline, resulting in more independent pipelines within the cluster and reducing the scope of each interruption. For example, ShuntServe constructs 2 pipelines for Llama-3.1-70B and 4 pipelines for Qwen3-32B on the evaluation cluster. This indicates that leveraging spot instances is particularly effective for

serving smaller models. Consequently, ShuntServe achieves 5.12 req/s for Qwen3-32B, with only approximately 1% throughput difference compared to On-demand Only.

7.2.2. Temporal Analysis of Latency under Spot Interruptions

We further analyze the impact of spot interruptions on real-time request processing by measuring the temporal evolution of end-to-end latency under online workloads. We evaluate both mean and P90 end-to-end latency to capture the average behavior and the tail behavior, respectively. As in the online evaluation of Section 7.1.3, the original arrival pattern is preserved while the overall request rate is scaled to keep all baselines below saturation. For Llama-3.1-70B, the arrival intervals are scaled by a factor of six, while for Qwen3-32B, the original trace is replayed unmodified, as the system throughput exceeds the original arrival rate. Figure 14 presents the temporal trends of mean and P90 end-to-end latency for both models across the five baselines.

For Llama-3.1-70B, No Handle exhibits the highest overall mean latency at 162.67s in Figure 14a. Tail latency is even more severely affected, with each baseline exhibiting pronounced peaks in Figure 14b in response to spot interruption events. For instance, the spot interruption at approximately 25 minutes drives the temporal P90 latency of No Handle beyond 500s. As a result, its overall P90 latency reaches 355.51s, approximately 81% higher than the 196.39s of On-demand Only. Among the baselines that incorporate fault tolerance, Request Migration achieves a mean latency of 155.88s and a P90 latency of 320.27s, while Concurrent Initialization achieves 126.10s and 251.95s, respectively.

For Qwen3-32B, where each pipeline consists of fewer instances and the cluster hosts a larger number of pipelines, the baselines exhibit a similar pattern to Llama-3.1-70B in Figures 14c and 14d. No Handle exhibits the highest mean and P90 latencies at 207.43s and 487.87s, followed by Request Migration at 188.70s and 467.86s, and Concurrent Initialization at 156.01s and 382.15s.

Across both models, Concurrent Initialization consistently achieves lower mean and P90 latency than Request Migration. This gap stems from the limitation of output-preserving request migration, which preserves the state of in-flight requests upon interruption but does not reduce the downtime itself. When a pipeline becomes unavailable during the downtime, incoming requests are concentrated on the remaining pipelines, increasing their batch sizes and inducing additional queueing delay. The resulting latency degradation persists until the affected pipeline is restored and the accumulated backlog is fully processed. These results indicate that minimizing downtime is critical for latency-sensitive workloads such as online serving. Consequently, ShuntServe, which combines both mechanisms, achieves the lowest latency among all baselines on both models, with mean and P90 latencies of 115.08s and 211.44s on Llama-3.1-70B and 138.00s and 289.09s on Qwen3-32B, corresponding to only 2.6%–7.7% increases over On-demand Only.




Figure 14: End-to-end latency trends over time under the spot availability scenario. Each data point represents a 5-minute trailing moving average of end-to-end request latency. (a) and (b) show results for Llama-3.1-70B, and (c) and (d) show results for Qwen3-32B. The x-axis represents elapsed time in seconds, and the y-axis represents end-to-end latency in seconds. The left column shows mean latency, and the right column shows P90 latency. Horizontal dashed lines indicate the overall metric for each baseline across the entire serving window.




Figure 15: Cost efficiency comparison across baselines, normalized to On-demand Only (lower is better). The percentage above each bar denotes the reduction relative to On-demand Only.

7.2.3. Cost Efficiency Comparison

The operating cost over the 50-minute window amounts to \$27.16 for On-demand Only, \$17.05 for Request Migration and No Handle, and \$18.15 for Concurrent Initialization and ShuntServe. Concurrent Initialization and ShuntServe incur approximately \$1.10 more than the other spot-based baselines because the interrupted node remains active until the new pipeline initialization completes, resulting in temporary overlap of billing between the interrupted and replacement nodes. To evaluate whether these cost reductions translate into actual cost efficiency

gains after accounting for the performance degradation caused by spot interruptions, Figure 15 presents the cost efficiency of each baseline normalized to On-demand Only. Even No Handle achieves lower cost per performance than On-demand Only in most cases, as the cost savings from spot pricing outweigh the performance loss. The exception arises in the P90 tail latency, where the impact of spot interruptions is most pronounced. For Llama-3.1-70B, No Handle exhibits a 2.5% higher cost per performance than On-demand Only.

For Llama-3.1-70B, Concurrent Initialization achieves lower cost per performance than Request Migration across all three workload settings (i.e., offline, online mean, and online P90) in Figure 15a. This stems from the fact that larger models occupy a greater portion of the cluster within each pipeline, expanding the scope affected by an interruption and amplifying the benefit of downtime reduction. In contrast, Qwen3-32B distributes the cluster across more pipelines, dispersing the impact of a single interruption. As a result, Request Migration achieves slightly lower cost per performance than Concurrent Initialization in the offline workload (Figure 15b). For online workloads, however, downtime affects request latency more severely, and Concurrent Initialization achieves lower cost per performance than Request Migration in both online mean and P90 metrics. ShuntServe, which combines both mechanisms, achieves the lowest cost per performance across all models and workloads, with cost efficiency improvements of up to 31.9% over On-demand Only.




Figure 16: Concurrent initialization overhead analysis, averaged across the measured runs.

7.2.4. Concurrent Initialization Overhead

We present a time breakdown across the components of the concurrent initialization process to quantify the duration of new pipeline preparation upon spot interruption. The node preparation time is measured six times for each of the three instance types in the evaluation cluster, yielding 18 measurements in total. The Shared Tensor Store and Inference Engine loading times are measured from the completion of node preparation until both components complete loading, repeated 10 times across both pipelines for Llama-3.1-70B, which incurs a longer Shared Tensor Store loading time than Qwen3-32B. Figure 16 presents the results.

The replacement node was provisioned and became ready for inference in 41.55s on average with a standard deviation of 7.54s. Once the node was ready, the Shared Tensor Store process was launched on it, loading the model partition from remote storage in 61.85s on average (std 9.59s, max 81.45s). In parallel, a new Inference Engine was initialized on both the unaffected node (i.e., Node 0) and the replacement node (i.e., New Node 1), taking 64.51s on average (std 9.25s, max 81.84s). The total initialization time amounts to approximately 111.3s on average. The downtime that ShuntServe incurs corresponds to the portion of the initialization time that exceeds the grace period. Since the grace period varies across cloud providers from 30s in Azure to 120s in AWS to 300s in Alibaba Cloud, the resulting downtime varies accordingly. In the AWS environment used in our evaluation, the average duration falls within the 120s grace period, allowing ShuntServe to achieve near-zero downtime upon spot interruption.

8. Discussion

The evaluation results demonstrate that ShuntServe achieves substantial cost efficiency on heterogeneous spot GPU clusters while maintaining comparable service quality to on-demand baselines. We discuss the limitations of the current design in Section 8.1 and outline directions for future extension in Section 8.2.

8.1. Limitation

Fidelity of the Estimator under Real-World Workloads. Optimization algorithms for identifying the best model placement

must evaluate and compare candidate placements throughout their search process; accurate prediction of the serving performance of a given placement configuration is therefore essential. ShuntServe achieves an average MAPE of 6.63% through roofline-based static latency estimation and throughput modeling, yet this accuracy is established under a synthetic dataset with fixed input/output lengths and a static batching environment. In real-world workloads, however, context length distributions are highly heterogeneous and request arrival patterns are dynamic. Moreover, the analytical formulation does not adequately capture non-linear factors such as kernel launch overhead, SM utilization, and cache effects. Prior systems such as Vidur [22] and Habitat [35] adopt ML-based approaches to model these non-linear characteristics, but at the cost of profiling that is sensitive to the combinatorially large configuration space. A promising direction for future work is to combine the strengths of these approaches. Integrating ShuntServe’s analytical formulation with an ML-based approach could reduce the profiling burden while capturing non-linear effects, and further integration with a simulation-based approach could enable adaptation to dynamic workload regimes.

Fault Tolerance for Long Context Workload. While the recomputation-based approach is sufficiently efficient for restoring the state of in-flight requests, it can incur higher overhead than KV cache transfer under very long context lengths (e.g., in the comparative measurements of Section 4.1, recomputation on L40S exhibits approximately 9.6% higher latency than transfer at a 64K context length). This overhead stems from the quadratic cost of computing attention scores in transformer layers, and can result in considerable tail latency degradation for workloads characterized by very long contexts. KV cache transfer, however, has a strict constraint in spot environments that the transfer, whose duration depends on the data size and the available network bandwidth, must complete within the grace period, making it unsuitable as a sole request recovery mechanism. A promising direction for future work is thus a hybrid recovery scheme that tracks the progress of in-flight requests and the remaining grace period, and selects an appropriate request recovery mechanism for each request individually. For instance, cache transfer may be elected for requests with very long contexts, while recomputation is retained for shorter ones.

8.2. Future Work

Adapting to Spot Volatility across Time and Regions. ShuntServe currently optimizes model placement over a static heterogeneous GPU cluster and responds to interruptions through auxiliary on-demand fallback. Spot environments, however, exhibit substantial fluctuations across both time and regions. Along the temporal dimension, the system could be extended to support the flexible membership of heterogeneous GPU instances in a manner that maximizes spot availability within the cluster, which would require a model placement algorithm capable of rapidly reconfiguring the pipeline topology in response to the dynamically changing cluster. Furthermore, while the current design assumes operation within a single region, extending

ShuntServe to multiple regions could yield additional cost savings and motivates a multi-region model placement optimizer tailored to the parallelization characteristics of LLM serving pipelines (e.g., replication incurs little traffic across replicas and is therefore well suited to placement across regions [36]).

Phase-Aware Model Placement. Prefill-decode disaggregation techniques [9, 10, 11, 8] can be integrated to maximize resource utilization in clusters composed of heterogeneous GPU instances. The distinct resource requirements of the prefill and decode phases in generative LLM inference expand the optimization space for model placement. Realizing this direction requires per-phase latency prediction and throughput modeling, as well as techniques for mitigating the KV cache transfer bottleneck that arises across heterogeneous instances, particularly under limited network bandwidth.

9. Related Work

Machine Learning System on Spot Instances. Several studies have investigated predicting and handling spot instance interruptions to achieve cost-efficient training and serving of machine learning models in cloud environments. Bamboo [37] provides fault-tolerant DNN training via redundant forward and backward pass computations across neighboring pipeline nodes. Parcae [38] predicts spot interruptions and proactively migrates pipelines to optimize liveput. SpotServe [33] pioneered distributed LLM serving on spot instances using dynamic re-parallelization and KV cache migration.

LLM Inference Latency Prediction on Various GPUs. Among existing LLM inference performance estimation approaches, Vidur [22] reports latency prediction errors below 9%, but requires extensive operator-level profiling across a combinatorially explosive configuration space that varies with models, GPU types, parallelization strategies, input/output sequence lengths, and batch sizes. MaverIQ [39] reduces profiling cost by introducing a *lightweight fingerprint* that contains only a few unique layers of the target LLM, and reports 1.3–1.7 $\times$ lower estimation error than Vidur. However, MaverIQ still requires profiling across the configuration space, which does not fundamentally resolve the scalability challenge of profiling-based approaches.

In contrast, ShuntServe’s estimator requires only a one-time, lightweight hardware calibration (FLOPS, memory bandwidth, and network bandwidth) for each GPU type, with negligible cost. The achieved MAPE of 6.63% is competitive with these profiling-based approaches and demonstrates sufficient accuracy as the foundational component of the model placement optimizer.

LLM Serving System on Heterogeneous GPUs. To address the scarcity of high-end GPUs, several studies have investigated maximizing cluster utilization by incorporating abundantly available commodity GPUs and accounting for the diverse performance characteristics of heterogeneous devices. Helix [18] models throughput as a network flow and optimizes model placement using max-flow and MILP algorithms. HEXGEN [19] models

inference and network costs, optimizing placement via genetic algorithms and DP. Thunderserve [10] optimizes disaggregated serving by specializing heterogeneous devices based on their compute-bound and memory-bound characteristics.

Deploying LLM serving systems on heterogeneous spot GPU clusters introduces challenges that are not addressed by any single line of prior work. Systems designed for spot instances [37, 38, 33] assume homogeneous GPU clusters, which are vulnerable to correlated failures where all instances of the same type are simultaneously interrupted. Systems designed for heterogeneous clusters [18, 19, 10] do not target spot instance environments and therefore lack fault tolerance mechanisms. For instance, to maximize cluster-wide throughput, these systems intertwine pipeline paths across instances, creating a single point of failure where a single instance interruption propagates to multiple pipelines. Latency prediction approaches [18, 22] rely on extensive profiling, the cost of which becomes intractable as the number of hardware types and configuration combinations increases.

The DP-based model placement optimizer of ShuntServe determines the optimal parallelization strategy for each pipeline in the heterogeneous GPU cluster while constraining each instance to a single pipeline, confining the impact of a spot interruption to a single pipeline. The profiling-free serving performance estimator based on the roofline model, combined with beam search, enables rapid and efficient exploration of the combinatorially explosive search space across a heterogeneous GPU cluster. Output-preserving request migration recovers interrupted requests through a recomputation-based approach, and concurrent initialization via the shared tensor store minimizes downtime by overlapping the initialization of a new pipeline with ongoing serving. These mechanisms enable ShuntServe to effectively handle spot interruptions and maximize the utilization of heterogeneous GPUs for cost-efficient LLM serving.

10. Conclusion

This paper presents ShuntServe, the first cost-efficient distributed LLM serving system designed for heterogeneous spot GPU clusters. ShuntServe achieves 1.42 $\times$ higher throughput on Llama-3.1-70B than the state-of-the-art homogeneous serving system and 1.35 $\times$ higher throughput on Qwen3-32B than state-of-the-art heterogeneous baselines through the DP-based model placement optimizer with a profiling-free roofline model-based serving performance estimator. ShuntServe further achieves 31.9% and 31.2% cost efficiency improvements over on-demand instances for offline and online serving, respectively, through output-preserving request migration and concurrent initialization via the shared tensor store.

Acknowledgements

This work was supported by Institute of Information & communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (RS-2022-00144309 & RS-2025-25441560 & RS-2026-25492200)

References

- [1] A. Grattafiori, A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Vaughan, A. Yang, A. Fan, A. Goyal, A. Hartshorn, A. Yang, A. Mitra, A. Sravankumar, A. Korenev, A. Hinsvark, A. Rao, A. Zhang, A. Rodriguez, A. Gregerson, A. Spataru, B. Roziere, B. Biron, B. Tang, B. Chern, C. Caucheteux, C. Nayak, C. Bi, C. Marra, C. McConnell, C. Keller, C. Touret, C. Wu, C. Wong, C. C. Ferrer, C. Nikolaidis, D. Allonsius, D. Song, D. Pintz, D. Livshits, D. Wyatt, D. Esiobu, D. Choudhary, D. Mahajan, D. Garcia-Olano, D. Perino, D. Hupkes, E. Lakomkin, E. AlBadawy, E. Lobanova, E. Dinan, E. M. Smith, F. Radenovic, F. Guzmán, F. Zhang, G. Synnaeve, G. Lee, G. L. Anderson, G. Thattai, G. Nail, G. Mialon, G. Pang, G. Cucurell, H. Nguyen, H. Korevaar, H. Xu, H. Touvron, I. Zarov, I. A. Ibarra, I. Kloumann, I. Misra, I. Evtimov, J. Zhang, J. Copet, J. Lee, J. Geffert, J. Vranes, J. Park, J. Mahadeokar, J. Shah, J. van der Linde, J. Billock, J. Hong, J. Lee, J. Fu, J. Chi, J. Huang, J. Liu, J. Wang, J. Yu, J. Bitton, J. Spisak, J. Park, J. Rocca, J. Johnston, J. Saxe, J. Jia, K. V. Alwala, K. Prasad, K. Upasani, K. Plawiak, K. Li, K. Heafield, K. Stone, K. El-Arini, K. Iyer, K. Malik, K. Chiu, K. Bhalla, K. Lakhotia, L. Rantala-Young, L. van der Maaten, L. Chen, L. Tan, L. Jenkins, L. Martin, L. Madaan, L. Malo, L. Blecher, L. Landzaat, L. de Oliveira, M. Muzzi, M. Pasupuleti, M. Singh, M. Paluri, M. Kardas, M. Tsimpoukelli, M. Oldham, M. Rita, M. Pavlova, M. Kambadur, M. Lewis, M. Si, M. K. Singh, M. Hassan, N. Goyal, N. Torabi, N. Bashlykov, N. Bogoychev, N. Chatterji, N. Zhang, O. Duchenne, O. Çelebi, P. Alrassy, P. Zhang, P. Li, P. Vasic, P. Weng, P. Bhargava, P. Dubal, P. Krishnan, P. S. Koura, P. Xu, Q. He, Q. Dong, R. Srinivasan, R. Ganapathy, R. Calderer, R. S. Cabral, R. Stojnic, R. Raileanu, R. Maheswari, R. Girdhar, R. Patel, R. Sauvestre, R. Polidoro, R. Sumbaly, R. Taylor, R. Silva, R. Hou, R. Wang, S. Hosseini, S. Chennabasappa, S. Singh, S. Bell, S. S. Kim, S. Edunov, S. Nie, S. Narang, S. Raparthy, S. Shen, S. Wan, S. Bhosale, S. Zhang, S. Vandenhennde, S. Batra, S. Whitman, S. Sootla, S. Collot, S. Gururangan, S. Borodinsky, T. Herman, T. Fowler, T. Sheasha, T. Georgiou, T. Scialom, T. Speckbacher, T. Mihaylov, T. Xiao, U. Karn, V. Goswami, V. Gupta, V. Ramanathan, V. Kerkez, V. Gonguet, V. Do, V. Vogeti, V. Albiero, V. Petrovic, W. Chu, W. Xiong, W. Fu, W. Meers, X. Martinet, X. Wang, X. Wang, X. E. Tan, X. Xia, X. Xie, X. Jia, X. Wang, Y. Goldschlag, Y. Gaur, Y. Babaei, Y. Wen, Y. Song, Y. Zhang, Y. Li, Y. Mao, Z. D. Coudert, Z. Yan, Z. Chen, Z. Papakipos, A. Singh, A. Srivastava, A. Jain, A. Kelsey, A. Shajnfeld, A. Gangidi, A. Victoria, A. Goldstand, A. Menon, A. Sharma, A. Boesenberg, A. Baevski, A. Feinstein, A. Kallet, A. Sangani, A. Teo, A. Yunus, A. Lupu, A. Alvarado, A. Caples, A. Gu, A. Ho, A. Poulton, A. Ryan, A. Ramchandani, A. Dong, A. Franco, A. Goyal, A. Saraf, A. Chowdhury, A. Gabriel, A. Bharambe, A. Eisenman, A. Yazdan, B. James, B. Maurer, B. Leonhardi, B. Huang, B. Loyd, B. D. Paola, B. Paranjape, B. Liu, B. Wu, B. Ni, B. Hancock, B. Wasti, B. Spence, B. Stojkovic, B. Gamido, B. Montalvo, C. Parker, C. Burton, C. Mejia, C. Liu, C. Wang, C. Kim, C. Zhou, C. Hu, C.-H. Chu, C. Cai, C. Tindal, C. Feichtenhofer, C. Gao, D. Civin, D. Beaty, D. Kreymer, D. Li, D. Adkins, D. Xu, D. Testuggine, D. David, D. Parikh, D. Liskovich, D. Foss, D. Wang, D. Le, D. Holland, E. Dowling, E. Jamil, E. Montgomery, E. Presani, E. Hahn, E. Wood, E.-T. Le, E. Brinkman, E. Arcaute, E. Dunbar, E. Smothers, F. Sun, F. Kreue, F. Tian, F. Kokinos, F. Ozgenel, F. Caggioni, F. Kanayet, F. Seide, G. M. Florez, G. Schwarz, G. Badeer, G. Swee, G. Halpern, G. Herman, G. Sizov, Guangyi, Zhang, G. Lakshminarayanan, H. Inan, H. Shojanazeri, H. Zou, H. Wang, H. Zha, H. Habeeb, H. Rudolph, H. Suk, H. Aspegren, H. Goldman, H. Zhan, I. Damlaj, I. Molybog, I. Tufanov, I. Leontiadis, I.-E. Veliche, I. Gat, J. Weissman, J. Geboski, J. Kohli, J. Lam, J. Asher, J.-B. Gaya, J. Marcus, J. Tang, J. Chan, J. Zhen, J. Reizenstein, J. Teboul, J. Zhong, J. Jin, J. Yang, J. Cummings, J. Carvill, J. Shepard, J. McPhie, J. Torres, J. Ginsburg, J. Wang, K. Wu, K. H. U. K. Saxena, K. Khandelwal, K. Zand, K. Matosich, K. Veeraraghavan, K. Michelena, K. Li, K. Jagadeesh, K. Huang, K. Chawla, K. Huang, L. Chen, L. Garg, L. A. L. Silva, L. Bell, L. Zhang, L. Guo, L. Yu, L. Moshkovich, L. Wehrstedt, M. Khabsa, M. Avalani, M. Bhatt, M. Mankus, M. Hasson, M. Lennie, M. Reso, M. Groshev, M. Naumov, M. Lathi, M. Keneally, M. Liu, M. L. Seltzer, M. Valko, M. Restrepo, M. Patel, M. Vyatskov, M. Samvelyan, M. Clark, M. Macey, M. Wang, M. J. Hermoso, M. Metanat, M. Rastegari, M. Bansal, N. Santhanam, N. Parks, N. White, N. Bawa, N. Singhal, N. Egebo, N. Usunier, N. Mehta, N. P. Laptev, N. Dong, N. Cheng, O. Chernoguz, O. Hart, O. Salpekar, O. Kalinli, P. Kent, P. Parekh, P. Saab, P. Balaji, P. Rittner, P. Bontrager, P. Roux, P. Dollar, P. Zvyagina, P. Ratanchandani, P. Yuvraj, Q. Liang, R. Alao, R. Rodriguez, R. Ayub, R. Murthy, R. Nayani, R. Mitra, R. Parthasarathy, R. Li, R. Hogan, R. Battey, R. Wang, R. Howes, R. Rinott, S. Mehta, S. Siby, S. J. Bondu, S. Datta, S. Chugh, S. Hunt, S. Dhillon, S. Sidorov, S. Pan, S. Mahajan, S. Verma, S. Yamamoto, S. Ramaswamy, S. Lindsay, S. Lindsay, S. Feng, S. Lin, S. C. Zha, S. Patil, S. Shankar, S. Zhang, S. Zhang, S. Wang, S. Agarwal, S. Sajuyigbe, S. Chintala, S. Max, S. Chen, S. Kehoe, S. Satterfield, S. Govindaprasad, S. Gupta, S. Deng, S. Cho, S. Virk, S. Subramanian, S. Choudhury, S. Goldman, T. Remez, T. Glaser, T. Best, T. Koehler, T. Robinson, T. Li, T. Zhang, T. Matthews, T. Chou, T. Shaked, V. Vontimitta, V. Ajayi, V. Montanez, V. Mohan, V. S. Kumar, V. Mangla, V. Ionescu, V. Poenaru, V. T. Mihailescu, V. Ivanov, W. Li, W. Wang, W. Jiang, W. Bouaziz, W. Constable, X. Tang, X. Wu, X. Wang, X. Wu, X. Gao, Y. Kleinman, Y. Chen, Y. Hu, Y. Jia, Y. Qi, Y. Li, Y. Zhang, Y. Zhang, Y. Adi, Y. Nam, Yu, Wang, Y. Zhao, Y. Hao, Y. Qian, Y. Li, Y. He, Z. Rait, Z. DeVito, Z. Rosnbrick, Z. Wen, Z. Yang, Z. Zhao, and Z. Ma, “The llama 3 herd of models,” 2024. [Online]. Available: <https://arxiv.org/abs/2407.21783>
- [2] C. Wang, Q. Liang, and B. Ugaonkar, “An empirical analysis of amazon ec2 spot instance features affecting cost-effective resource procurement,” in *Proceedings of the 8th ACM/SPEC on International Conference on Performance Engineering*, ser. ICPE ’17. New York, NY, USA: Association for Computing Machinery, 2017, p. 63–74. [Online]. Available: <https://doi.org/10.1145/3030207.3030210>
- [3] B. Javadi, R. K. Thulasiram, and R. Buyya, “Characterizing spot price dynamics in public cloud environments,” *Future Generation Computer Systems*, vol. 29, no. 4, pp. 988–999, 2013, special Section: Utility and Cloud Computing. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167739X12001483>
- [4] S. Cheon, K. Kim, K. Kim, M. Song, and K. Lee, “Multi-node spot instances availability score collection system,” in *Proceedings of the 34th International Symposium on High-Performance Parallel and Distributed Computing (HPDC ’25)*. ACM, July 2025.
- [5] K. Kim and K. Lee, “Making cloud spot instance interruption events visible,” in *Proceedings of the ACM on Web Conference 2024*, ser. WWW ’24. New York, NY, USA: Association for Computing Machinery, 2024, p. 2998–3009. [Online]. Available: <https://doi.org/10.1145/3589334.3645548>
- [6] S. Lee, J. Hwang, and K. Lee, “Spotlake: Diverse spot instance dataset archive service,” in *2022 IEEE International Symposium on Workload Characterization (IISWC)*. Los Alamitos, CA, USA: IEEE Computer Society, nov 2022, pp. 242–255. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/IISWC55918.2022.00029>
- [7] K. Kim, S. Park, J. Hwang, H. Lee, S. Kang, and K. Lee, “Public spot instance dataset archive service,” in *Companion Proceedings of the ACM Web Conference 2023*, ser. WWW ’23 Companion. New York, NY, USA: Association for Computing Machinery, 2023, p. 69–72. [Online]. Available: <https://doi.org/10.1145/3543873.3587314>
- [8] Y. Zhong, S. Liu, J. Chen, J. Hu, Y. Zhu, X. Liu, X. Jin, and H. Zhang, “DistServe: Disaggregating prefill and decoding for goodput-optimized large language model serving,” in *18th USENIX Symposium on Operating Systems Design and Implementation (OSDI 24)*. Santa Clara, CA: USENIX Association, Jul. 2024, pp. 193–210. [Online]. Available: <https://www.usenix.org/conference/osdi24/presentation/zhong-yinmin>
- [9] P. Patel, E. Choukse, C. Zhang, A. Shah, I. n. Goiri, S. Maleki, and R. Bianchini, “Splitwise: Efficient generative llm inference using phase splitting,” in *Proceedings of the 51st Annual International Symposium on Computer Architecture*, ser. ISCA ’24. IEEE Press, 2025, p. 118–132. [Online]. Available: <https://doi.org/10.1109/ISCA59077.2024.00019>
- [10] Y. Jiang, F. Fu, X. Yao, T. Wang, B. Cui, A. Klimovic, and E. Yoneki, “Thunderserve: High-performance and cost-efficient llm serving in cloud environments,” 2025. [Online]. Available: <https://arxiv.org/abs/2502.09334>
- [11] Y. Jiang, F. Fu, X. Yao, G. He, X. Miao, A. Klimovic, B. Cui, B. Yuan, and E. Yoneki, “Demystifying cost-efficiency in LLM serving over heterogeneous GPUs,” in *Proceedings of the 42nd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, A. Singh, M. Fazel, D. Hsu, S. Lacoste-Julien, F. Berkenkamp, T. Maharaj, K. Wagstaff, and J. Zhu, Eds., vol. 267. PMLR, 13–19 Jul 2025, pp. 27 534–27 552. [Online]. Available: <https://proceedings.mlr.press/v267/jiang25c.html>

- [12] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei, "Scaling laws for neural language models," 2020. [Online]. Available: <https://arxiv.org/abs/2001.08361>
- [13] T. Henighan, J. Kaplan, M. Katz, M. Chen, C. Hesse, J. Jackson, H. Jun, T. B. Brown, P. Dhariwal, S. Gray, C. Hallacy, B. Mann, A. Radford, A. Ramesh, N. Ryder, D. M. Ziegler, J. Schulman, D. Amodei, and S. McCandlish, "Scaling laws for autoregressive generative modeling," 2020. [Online]. Available: <https://arxiv.org/abs/2010.14701>
- [14] D. Narayanan, M. Shoenybi, J. Casper, P. LeGresley, M. Patwary, V. Korthikanti, D. Vainbrand, P. Kashinkunti, J. Bernauer, B. Catanzaro, A. Phanishayee, and M. Zaharia, "Efficient large-scale language model training on gpu clusters using megatron-lm," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '21. New York, NY, USA: Association for Computing Machinery, 2021. [Online]. Available: <https://doi.org/10.1145/3458817.3476209>
- [15] Y. Huang, Y. Cheng, A. Bapna, O. Firat, M. X. Chen, D. Chen, H. Lee, J. Ngiam, Q. V. Le, Y. Wu, and Z. Chen, *GPipe: efficient training of giant neural networks using pipeline parallelism*. Red Hook, NY, USA: Curran Associates Inc., 2019.
- [16] X. He, P. Shenoy, R. Sitaraman, and D. Irwin, "Cutting the cost of hosting online services using cloud spot markets," in *Proceedings of the 24th International Symposium on High-Performance Parallel and Distributed Computing*, ser. HPDC '15. New York, NY, USA: Association for Computing Machinery, 2015, p. 207–218. [Online]. Available: <https://doi.org/10.1145/2749246.2749275>
- [17] A. Ali-Eldin, J. Westin, B. Wang, P. Sharma, and P. Shenoy, "Spotweb: Running latency-sensitive distributed web services on transient cloud servers," in *Proceedings of the 28th International Symposium on High-Performance Parallel and Distributed Computing*, ser. HPDC '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 1–12. [Online]. Available: <https://doi.org/10.1145/3307681.3325397>
- [18] Y. Mei, Y. Zhuang, X. Miao, J. Yang, Z. Jia, and R. Vinayak, "Helix: Serving large language models over heterogeneous gpus and network via max-flow," in *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1*, ser. ASPLOS '25. New York, NY, USA: Association for Computing Machinery, 2025, p. 586–602. [Online]. Available: <https://doi.org/10.1145/3669940.3707215>
- [19] Y. Jiang, R. Yan, X. Yao, Y. Zhou, B. Chen, and B. Yuan, "Hexgen: generative inference of large language model over heterogeneous environment," in *Proceedings of the 41st International Conference on Machine Learning*, ser. ICML'24. JMLR.org, 2024.
- [20] W. Kwon, Z. Li, S. Zhuang, Y. Sheng, L. Zheng, C. H. Yu, J. Gonzalez, H. Zhang, and I. Stoica, "Efficient memory management for large language model serving with pagedattention," in *Proceedings of the 29th Symposium on Operating Systems Principles*, ser. SOSP '23. New York, NY, USA: Association for Computing Machinery, 2023, p. 611–626. [Online]. Available: <https://doi.org/10.1145/3600006.3613165>
- [21] NVIDIA, "TensorRT-LLM: A library for accelerating large language model inference," GitHub repository, 2023, accessed: 2025-10-30. [Online]. Available: <https://github.com/NVIDIA/TensorRT-LLM>
- [22] A. Agrawal, N. Kedia, J. Mohan, A. Panwar, N. Kwatra, B. S. Gulavani, R. Ramjee, and A. Tumanov, "Vidur: A large-scale simulation framework for llm inference," in *Proceedings of Machine Learning and Systems*, P. Gibbons, G. Pekhimenko, and C. D. Sa, Eds., vol. 6, 2024, pp. 351–366. [Online]. Available: https://proceedings.mlsys.org/paper_files/paper/2024/file/b74a8de47d2b3c928360e0a011f48351-Paper-Conference.pdf
- [23] Z. Li, L. Zheng, Y. Zhong, V. Liu, Y. Sheng, X. Jin, Y. Huang, Z. Chen, H. Zhang, J. E. Gonzalez, and I. Stoica, "AlpaServe: Statistical multiplexing with model parallelism for deep learning serving," in *17th USENIX Symposium on Operating Systems Design and Implementation (OSDI 23)*. Boston, MA: USENIX Association, Jul. 2023, pp. 663–679. [Online]. Available: <https://www.usenix.org/conference/osdi23/presentation/li-zhouhan>
- [24] S. Williams, A. Waterman, and D. Patterson, "Roofline: an insightful visual performance model for multicore architectures," *Commun. ACM*, vol. 52, no. 4, p. 65–76, Apr. 2009. [Online]. Available: <https://doi.org/10.1145/1498765.1498785>
- [25] C. Yang, Y. Wang, S. Farrell, T. Kurth, and S. Williams, "Hierarchical roofline performance analysis for deep learning applications," 2020. [Online]. Available: <https://arxiv.org/abs/2009.05257>
- [26] R. W. Hockney, "The communication challenge for mpp: Intel paragon and meiko cs-2," *Parallel Computing*, vol. 20, no. 3, pp. 389–398, 1994. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167819106800219>
- [27] R. Thakur and W. D. Gropp, "Improving the performance of collective operations in mpich," in *Recent Advances in Parallel Virtual Machine and Message Passing Interface*, J. Dongarra, D. Laforenza, and S. Orlando, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2003, pp. 257–267.
- [28] Z. Hu, S. Shen, T. Bonato, S. Jeaugey, C. Alexander, E. Spada, J. Dinan, J. Hammond, and T. Hoeftler, "Demystifying NCCL: An In-depth Analysis of GPU Communication Protocols and Algorithms," in *Proceedings of the 32nd Annual Symposium on High-Performance Interconnects (HOTI'25)*. IEEE Press, 08 2025.
- [29] R. Sanovar, S. Bharadwaj, R. S. Amant, V. Rühle, and S. Rajmohan, "Lean attention: Hardware-aware scalable attention mechanism for the decode-phase of transformers," in *Proceedings of Machine Learning and Systems*, vol. 7, 2025.
- [30] P. Patel, E. Choukse, C. Zhang, I. n. Goiri, B. Warriar, N. Mahalingam, and R. Bianchini, "Characterizing power management opportunities for llms in the cloud," in *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3*, ser. ASPLOS '24. New York, NY, USA: Association for Computing Machinery, 2024, p. 207–222. [Online]. Available: <https://doi.org/10.1145/3620666.3651329>
- [31] P. Moritz, R. Nishihara, S. Wang, A. Tumanov, R. Liaw, E. Liang, M. Elibol, Z. Yang, W. Paul, M. I. Jordan, and I. Stoica, "Ray: a distributed framework for emerging ai applications," in *Proceedings of the 13th USENIX Conference on Operating Systems Design and Implementation*, ser. OSDI'18. USA: USENIX Association, 2018, p. 561–577.
- [32] Q. Team, "Qwen3 technical report," 2025. [Online]. Available: <https://arxiv.org/abs/2505.09388>
- [33] X. Miao, C. Shi, J. Duan, X. Xi, D. Lin, B. Cui, and Z. Jia, "Spotserv: Serving generative large language models on preemptible instances," in *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, ser. ASPLOS '24. New York, NY, USA: Association for Computing Machinery, 2024, p. 1112–1127. [Online]. Available: <https://doi.org/10.1145/3620665.3640411>
- [34] Z. Mao, T. Xia, Z. Wu, W.-L. Chiang, T. Griggs, R. Bhardwaj, Z. Yang, S. Shenker, and I. Stoica, "Skyserve: Serving ai models across regions and clouds with spot instances," in *Proceedings of the Twentieth European Conference on Computer Systems*, ser. EuroSys '25. New York, NY, USA: Association for Computing Machinery, 2025, p. 159–175. [Online]. Available: <https://doi.org/10.1145/3689031.3717459>
- [35] G. X. Yu, Y. Gao, P. Golikov, and G. Pekhimenko, "Habitat: A Runtime-Based computational performance predictor for deep neural network training," in *2021 USENIX Annual Technical Conference (USENIX ATC 21)*. USENIX Association, Jul. 2021, pp. 503–521. [Online]. Available: <https://www.usenix.org/conference/atc21/presentation/you>
- [36] Z. Mao, T. Xia, Z. Wu, W.-L. Chiang, T. Griggs, R. Bhardwaj, Z. Yang, S. Shenker, and I. Stoica, "Skyserve: Serving ai models across regions and clouds with spot instances," in *Proceedings of the Twentieth European Conference on Computer Systems*, ser. EuroSys '25. New York, NY, USA: Association for Computing Machinery, 2025, p. 159–175. [Online]. Available: <https://doi.org/10.1145/3689031.3717459>
- [37] J. Thorpe, P. Zhao, J. Eyolfson, Y. Qiao, Z. Jia, M. Zhang, R. Netravali, and G. H. Xu, "Bamboo: Making preemptible instances resilient for affordable training of large DNNs," in *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*. Boston, MA: USENIX Association, Apr. 2023, pp. 497–513. [Online]. Available: <https://www.usenix.org/conference/nsdi23/presentation/thorpe>
- [38] J. Duan, Z. Song, X. Miao, X. Xi, D. Lin, H. Xu, M. Zhang, and Z. Jia, "Parcae: Proactive, Liveput-Optimized DNN training on preemptible instances," in *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24)*. Santa Clara, CA: USENIX Association, Apr. 2024, pp. 1121–1139. [Online]. Available: <https://www.usenix.org/conference/nsdi24/presentation/duan>

- [39] D. Liakopoulos, P. Sinha, T. Hu, M. Lee, and N. J. Yadwadkar, "Maveriq: Fingerprint-guided extrapolation and fragmentation-aware layering for intent-based llm serving," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '25. New York, NY, USA: Association for Computing Machinery, 2025, p. 1676–1696. [Online]. Available: <https://doi.org/10.1145/3712285.3759867>