



OrbitFlow: SLO-Aware Long-Context LLM Serving with Fine-Grained KV Cache Reconfiguration

Xinyue Ma*
POSTECH
xinyuema@postech.ac.kr

Heelim Hong*
UNIST
heelim@unist.ac.kr

Taegeon Um
Samsung Research
taegeon.um@samsung.com

Jongseop Lee
POSTECH
jslee202403@postech.ac.kr

Seoyeong Choy
POSTECH
sychoy0704@postech.ac.kr

Woo-Yeon Lee
Samsung Research
wooyeon0.lee@samsung.com

Myeongjae Jeon
POSTECH
mj.jeon@postech.ac.kr

ABSTRACT

Serving long-context LLMs is challenging because request lengths and batch composition vary during token generation, causing the memory footprint to fluctuate significantly at runtime. Offloading KV caches to host memory limits effective GPU memory usage, but existing static and predetermined offloading strategies cannot adapt to the rapidly shifting memory demands of long-context serving. This often leads to excessive CPU-to-GPU KV transfers that translate into latency spikes and frequent SLO violations.

To address these challenges, we introduce ORBITFLOW, a fine-grained and adaptive KV cache management system that meets latency SLOs in long-context LLM serving. ORBITFLOW employs a lightweight ILP solver to decide which layers' KV caches to retain on the GPU for each request, within memory capacity constraints. It continuously refines KV placements based on runtime feedback when the active plan becomes suboptimal during token generation. Under heavy load, ORBITFLOW invokes a fallback mechanism to temporarily defer in-flight requests with large memory footprints, preserving overall SLO attainment. Our experiments demonstrate that ORBITFLOW improves SLO attainment for TPOT and TBT by 62% and 66%, respectively, while reducing the 95th percentile (i.e., P95) latency by 38% and achieving up to 3.3× higher throughput compared to existing offloading methods.

PVLDB Reference Format:

Xinyue Ma, Heelim Hong, Taegeon Um, Jongseop Lee, Seoyeong Choy, Woo-Yeon Lee, and Myeongjae Jeon. OrbitFlow: SLO-Aware Long-Context LLM Serving with Fine-Grained KV Cache Reconfiguration. PVLDB, 19(5): 1046 - 1059, 2026.

doi:10.14778/3796195.3796214

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/omnia-postech/OrbitFlow>.

*These authors contributed equally to this work.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 19, No. 5 ISSN 2150-8097.
doi:10.14778/3796195.3796214

1 INTRODUCTION

Large language models (LLMs) have gained significant prominence across various natural language processing tasks, including text generation [5, 7, 9, 11, 17, 28], language translation [31, 34, 53], summarization [26, 38], and question answering [16, 22, 41]. Their capabilities have greatly improved over the past few years, largely empowered by dramatic increases in *context window size*—i.e., the maximum number of tokens in words a model can process at once. For instance, GPT-2 [42] supported a context window of 1K tokens, while its successor GPT-3 [7] doubled the window size to 2K tokens. More recent models such as GPT-4 [27] and Claude [3] have pushed the limits further, handling up to 128K or even 1M tokens.

To provide the best user experience, LLM serving systems must respond to requests within stringent latency requirements defined by service-level objectives (SLOs) [1, 40, 56]. Upon receiving an input sequence of tokens, these systems generate output tokens autoregressively, one token at a time, through consecutive *decode steps*. Many LLM services, such as chat agents, creative writing tools, and real-time translators, stream tokens as they are generated so that users can read responses incrementally. To keep decode latency perceptually invisible, these services typically set per-token latency SLOs (e.g., 200–300 ms for chatbots) to match natural reading speeds [36, 56]. Frequent violations of these latency SLOs can impair user experience and reduce revenue [4], regardless of response quality. As context windows expand and workloads become more decode-dominant, optimizing for per-token latency SLOs becomes ever more critical.

However, serving long-context LLMs within strict time bounds is challenging because GPU memory demands vary widely during runtime. Modern serving systems accelerate token generation by performing decode steps based on cached states of all preceding input-output tokens, termed *KV cache*. Rapid decoding thus requires storing all KV caches in GPU memory, which was manageable with small context windows. Such caching has now become nontrivial for two reasons: (1) although initially small, the KV cache grows steadily during lengthy decoding, creating a high likelihood of exhausting GPU memory (i.e., *token dimension*); and (2) LLM serving systems often dynamically batch multiple requests to process them concurrently on the GPU [20, 29], further straining the available

GPU memory (i.e., *batch dimension*). There are thriving efforts to reduce KV cache size itself during the decode step. For instance, state restoration [45] recomputes parts of the KV cache at runtime, and token pruning [13, 30, 47, 49, 54] discards presumably less critical KV entries. While effective, these approaches carry inherent limitations, as they either apply only to specific models with multi-head attention (MHA) [45] or risk degrading response quality.

This paper tackles GPU memory challenges in long-context serving through *KV offloading* [50], which moves parts of the KV cache to host memory and overlaps their retrieval with ongoing GPU computation. KV offloading does not compromise accuracy and, when scheduled concurrently, can achieve effective computation-communication overlap. The key insight behind this is that each layer uses only its own cached KV states and does not access caches from other layers. However, current common practices are largely SLO-unaware: they statically and monolithically offload and prefetch states for all batched requests layer by layer to save memory [30, 51]. Even with larger offload distances, i.e., evicting caches at regular but wider layer intervals within the GPU memory budget, static strategies fail to conceal escalating data-transfer overheads as the cache expands along both token and batch dimensions. This overhead markedly delays GPU computations and subsequent token decode steps for every request in a batch, making it difficult to consistently meet latency SLOs across requests.

Our system, ORBITFLOW, realizes effective KV offloading through *adaptive, fine-grained* offload plans. At its core, ORBITFLOW employs a lightweight ILP-based solver that dynamically optimizes GPU memory usage across batched requests to minimize SLO violations. The solver iteratively refines optimal offload plans using real-time feedback on varying state-transfer and GPU computation costs per layer, allowing ORBITFLOW to quickly adjust decisions under current GPU memory and compute resource constraints.

For fine-grained KV offloading, ORBITFLOW treats each request in a dynamic batch individually based on its current stage in the decoding phase. A key challenge is balancing GPU memory allocation among requests with distinct memory demands given limited GPU capacity. For instance, when serving a single request, ORBITFLOW configures the largest possible offload distance to fully utilize available GPU memory and achieve maximum token generation speed. When another request arrives later, the two requests form a batch at different stages in decoding, with non-identical KV cache sizes. Within such batches, the memory demands of requests often conflict: a newly arrived request requires memory for rapid token generation, whereas an older request with substantial decoding progress needs to retain a large portion of its KV cache on GPU to effectively overlap computation and memory prefetching. To reconcile this, our solver holistically examines various offload-reload plans within the GPU memory budget at the *request level* and selects the one that best prevents latency violations for each request.

It is worth noting that ORBITFLOW may fail to meet latency SLOs when the server is severely oversubscribed. Nonetheless, it offers two simple yet effective fallback mechanisms—*Token-Deposit* and *Pause-Resume*—to mitigate latency impacts during transient oversubscription. When a request produces output tokens while aggressively using GPU memory, ORBITFLOW does not always emit these tokens immediately. Instead, it emits tokens at a rate compliant with latency SLOs. Specifically, if LLM serving must adhere




Figure 1: (a) Prefill stage computes all input tokens in parallel. (b) Decode stage loads KV from the cache and computes only one token at a time. (c) MHA maps distinct KV per query head, while GQA shares KV across head groups to reduce the KV size.

to a constant time-between-tokens (TBT) [1, 40], ORBITFLOW first deposits generated tokens into a buffer and then releases them at a steady rate matching the predefined TBT to smooth the output rate. This mechanism does not stand alone; our solver uses it to temporarily pause suitable requests with high memory footprints and many buffered tokens, freeing GPU memory to better accommodate other requests in the batch. Since ORBITFLOW continues supplying users with deposited tokens during pauses, it maintains the perception of continuous request processing. These mechanisms are particularly beneficial in batches where some requests complete quickly, as resources freed by such requests can then be reclaimed to resume paused requests without disrupting the user experience.

We build ORBITFLOW atop vLLM v0.6.6 [29] and evaluate it against baseline offloading systems on ShareGPT-derived synthetic traces across different arrival rates. Our evaluation shows that ORBITFLOW improves time-per-output-token (TPOT) and TBT SLO attainment by 62% and 66%, respectively, while reducing P95 latency by 38% and achieving up to 3.3× higher throughput, over existing methods. Additionally, ORBITFLOW consistently provides performance improvements across different context lengths and batch sizes, and under distributed serving scenarios.

2 KV OFFLOADING: A UNIFIED VIEW

Service-level objectives (SLOs) typically target millisecond-level latency. Yet, the KV cache that grows linearly in batched long-context workloads and exhausts GPU memory makes these SLOs hard to meet with existing static KV offloading methods [50, 51]. In this section, we detail these problems and explain why dynamic offloading is particularly necessary for long-context LLM inference.

2.1 Background

Many contemporary LLM workloads involve decode-dominant, highly interactive tasks like chat agents, code assistants, and creative writing tools. Such tasks often require context lengths exceeding 32K tokens, sometimes up to millions, as seen in recent flagship models [3, 6, 8, 12, 17, 18, 24, 27, 48, 52]. The KV size grows with the input-output sequence length, so longer contexts translate directly into higher memory demand.

Inference Stages. In modern LLMs built from stacked Transformer layers, inference comprises two distinct stages, as shown in Fig. 1: *prefill* (compute-bound), where input tokens are processed in parallel to generate initial context information, and *decode* (memory-bound), where tokens are generated autoregressively.

At the core of both stages is the attention mechanism. In this mechanism, inputs are projected to query (Q), key (K), and value (V)




Figure 2: (a) The KV cache size of various sequence lengths with the LLaMA3-70B model. (b) A layer’s computation time and communication time to transfer its KV from CPU to GPU over sequence lengths with the LLaMA3-8B model. (c) With a fixed token-level SLO, tokens generated in early steps become the latency margin.

tensors, which are split into multiple heads for parallel processing. As shown in Fig. 1c, unlike multi-head attention (MHA) adopted by initial GPT models (e.g., GPT-2), recent models such as LLaMA3 adopt grouped-query attention (GQA), where the KV tensors are split into fewer heads than the Q tensor to reduce the KV cache size. Even with GQA, the KV size remains massive at sequence lengths from 128K to 1M tokens, as shown in Fig. 2a.

SLO Metrics. In LLM serving, time-to-first-token (TTFT) and time-per-output-token (TPOT) are common SLO metrics because they provide straightforward insights into user-perceived responsiveness and average throughput. On their own, however, these metrics may not fully capture critical aspects of latency in extended context scenarios. While TPOT accurately measures average latency, it does not reveal whether token generation is consistently steady, since slow tokens can be offset by faster ones. As a result, recent research highlights token-level latency metrics, such as time-between-tokens (TBT), alongside TPOT. For the same latency target (e.g., 50 ms per token), under a TPOT SLO, a request violates the SLO when its average per-token latency exceeds the target, whereas under a TBT SLO, each token exceeding that latency constitutes a violation. TBT thus explicitly pinpoints individual token delays that affect user experience in interactive workloads [1, 23, 40]. By adopting both TPOT and TBT, we gain a comprehensive view of system responsiveness, ensuring that both overall throughput and fine-grained latency variations are captured and evaluated.

Continuous Batching. Production LLM servers process multiple concurrent requests by grouping them into batches. Continuous batching [1, 20, 29, 44, 56] is a widely adopted strategy because it sustains high throughput without provisioning additional serving instances. Under continuous batching, when any request in a batch completes, a new request is admitted and its prefill begins, while the ongoing decoding is temporarily paused. Unless otherwise specified, we assume continuous batching for all subsequent analyses.

2.2 Drawbacks of Previous Offloading Methods

Previous offloading methods treat all requests within a batch uniformly [2, 35, 51]. They operate at the layer granularity: if a layer is offloaded, the KV entries for that layer are moved to the CPU for all requests; otherwise, all these KV entries stay resident on GPU. On top of this, the layers to offload are selected in a distance-driven manner, meaning they are evenly distributed across the model depth. For instance, if two out of six layers are offloaded, the

offload pattern would look like [1 1 0 1 1 0], where 1 denotes a layer kept on GPU and 0 denotes one offloaded to CPU.

In Transformer-based LLM inference, computation in a given layer can start only after its entire KV cache is present on the GPU. Therefore, each layer’s KV cache across batched requests share a single operational timeline. The computation time per layer is consistent, with the same amount of KV data to transfer if a layer is offloaded. This makes offloading every N-th layer of a model equivalent to repeating groups of N layers (e.g., two groups of [1 1 0] in the example above). Within each group, the computation, communication, and resulting stalls are identical, giving the runtime a steady and predictable pattern to exploit.

The distance N is usually chosen based on the worst-expected workload, so that KV caches at maximum sequence length still fit in GPU memory and offloaded layers can overlap their transfers with computation as much as possible. Some prior systems [35, 51] profile offloading distances over a grid of batch sizes and sequence lengths to select the best value for each workload. Fixing this distance ahead of time avoids runtime overhead but can be suboptimal if actual workloads deviate from the worst-case assumptions. This reduces compute-communication overlap and increases stall risk. We examine this effect in two representative cases: token-dimension drift and batch-dimension mismatch.

Token-Dimension Drift. Fig. 2b plots computation and data transfer times against sequence length for LLaMA3-8B on a single NVIDIA A6000 GPU. As decoding proceeds and sequence length grows, both costs grow linearly. However, the transfer time rises faster, making the optimal offloading strategy dependent on their relative ratio. For instance, at a sequence length of 2K, data transfer takes about 1.9× the time needed to compute one layer. In this case, offloading every third layer hides the transfer time under computation, yielding perfect overlap. As decoding continues and the sequence length reaches 16K, the ratio becomes 6.1×, and the offloading strategy that worked at 2K tokens no longer fully hides the transfer time and must be adjusted.

Notably, reducing decode latency always improves the TPOT SLO attainment, but does not always improve TBT SLO attainment. For example, if we apply an offloading strategy optimal for perfect overlap and the sequence is initially short, early tokens may be emitted faster than necessary. However, each subsequent token’s TBT SLO resets relative to the prior one, so early savings cannot be carried forward in improving TBT attainment. To illustrate, the green area in Fig. 2c represents the *latency margin* created while the TBT stays below the 52 ms SLO. As the sequence length increases, the margin is gone even though TPOT stays low, showing that early speed gains improve TPOT but not TBT. In this work, we track *latency margin* as a system-wide resource, harvesting it early and spending it to mask later stalls.

Batch-Dimension Mismatch. Real-world workloads vary substantially in request lengths, arrival times, and output sizes. As a result, requests in the same batch may be at different decoding stages with different KV cache sizes. Compute time grows with the total number of tokens in the batch, and data transfer time grows with the number of KV blocks that are offloaded. Under a uniform distance-driven offload policy, the KV cache of one layer

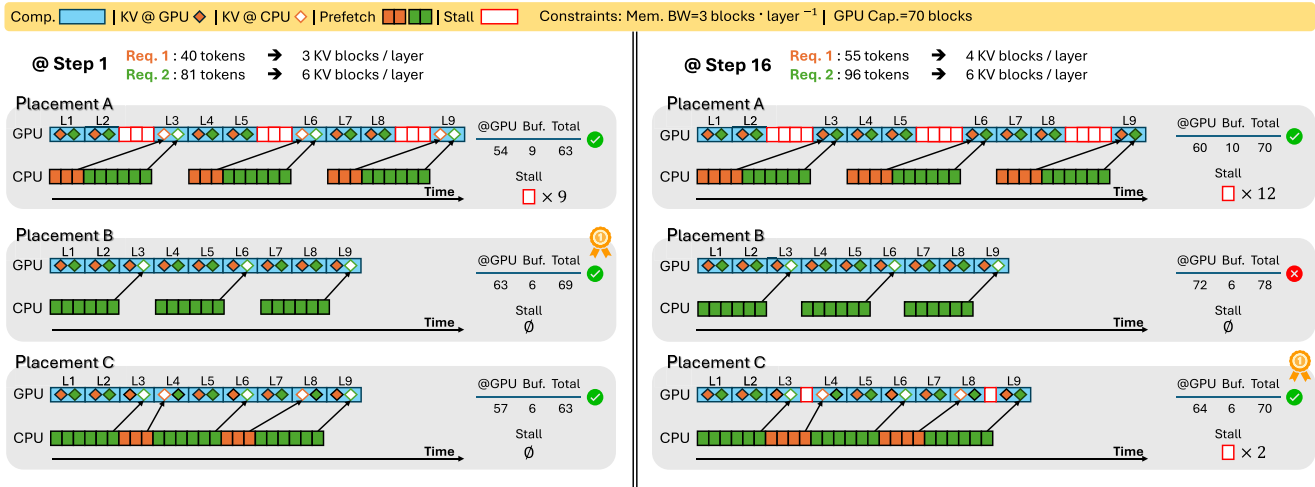


Figure 3: An illustrative example showing why a static, uniform offload policy is insufficient. The timelines track one decode step early (Step 1, left) and one much later (Step 16, right) for three placements. Each timeline is annotated with the resident KV blocks on GPU, the size of the per-layer prefetch buffer, and the resulting stalls.

is either entirely on GPU or CPU, so both costs scale with the total KV size rather than the individual KV cache sizes of each request. Short requests are therefore delayed, while the KV of a long request pushes up both compute and data transfer time for the batch.

This impact also depends on the batching strategy. Consider a batch with one long request processed alongside several short requests, with additional short requests waiting in the queue. Continuous batching allows the admission of one queued request whenever a short request in the batch finishes. The long request, however, continues to dominate compute and data transfer and slows down token generation. This is a poor trade-off: a single request with a large KV footprint can put all other short requests at risk of SLO violations, thereby degrading overall system performance.

2.3 An Illustrative Example

To ground the two limitations pinpointed in the previous section, we walk through the minimal scenario shown in Fig. 3. We assume the PagedAttention KV layout [29], where each layer’s KV cache is allocated in fixed-size blocks of 16 tokens. Two concurrent requests are served on a 9-layer LLM with Request 2 around twice as long, highlighting the *batch-dimension mismatch*. The left half of the figure captures this initial state (Step 1), while the right half shows the status after 15 decode steps (Step 16). By this point, Request 1’s KV cache has grown from three to four blocks per layer because its additional tokens demand a new 16-token block—each decode step adds one token, but allocation increases only when a 16-token block boundary is crossed. Request 2, however, still fits within its existing blocks, as shown in the figure’s upper panel. This growth illustrates *token-dimension drift*. In this example, we assume the GPU can store at most 70 KV blocks, and the PCIe bandwidth allows transferring at most 3 KV blocks within the computation time of one layer. When the combined KV cache size of all requests exceeds this budget, some blocks must stay in CPU memory and be prefetched into a small staging buffer on the GPU to process the corresponding layer.

The scenario in Fig. 3 compares three KV placement plans. **Placement A** serves as an example of the *uniform* offload method in

which both requests offload the evenly spaced layers L3, L6 and L9. **Placement B** attempts to mitigate the batch-dimension mismatch by keeping all KVs of the shorter Request 1 on the GPU and offloading 3 layers for the longer Request 2. **Placement C** offloads 2 layers of Request 1 and 3 layers of Request 2, trading a little more CPU-to-GPU traffic for extra headroom in GPU memory.

At Decoding Step 1. Placement A must move 9 blocks just in time for every offloaded layer. Since the bandwidth allows transferring 3 blocks per layer time, those transfers spill past the preceding layer’s compute and create 9 units of stalls. Placement A is already the *best* uniform choice: offloading one more layer leaves even less computation time to overlap with the data transfer, resulting in longer stalls; offloading one less layer requires keeping too many blocks on GPU, violating the 70-block memory limit. Both Placements B and C apply request-wise placement and achieve perfect overlap without stalls, but Placement B transfers fewer blocks and keeps higher GPU memory utilization.

Observation 1) A uniform offload policy that ignores the batch-dimension mismatch can miss the optimal placement and leave avoidable stalls. Allowing *fine-grained* per-request placements can better overlap computation with data transfer.

At Decoding Step 16. After 15 more decode steps, Request 2 still needs 6 blocks per layer, while Request 1 has grown to 4 blocks per layer. These additional blocks push Placement B’s memory usage to 78 blocks—*beyond* the 70-block GPU memory capacity—making it infeasible. Placement A remains feasible but must now prefetch 10 blocks per layer, causing 12 stalls in total. Placement C, which had sacrificed a bit of GPU residency earlier, still fits within capacity and incurs only 2 stalls, becoming the new optimal strategy.

Observation 2) A placement that was optimal early in decoding can later become suboptimal—or even infeasible—as KV caches grow. A single, static placement therefore cannot serve all decode steps; *dynamic reconfiguration* is required.

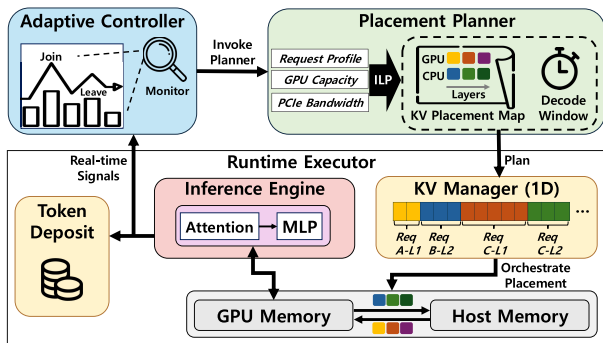


Figure 4: The overall architecture of ORBITFLOW.

3 SYSTEM DESIGN

We now introduce ORBITFLOW, an implementation of fine-grained, adaptive KV cache placement that optimizes KV offloading for long-context LLM serving under limited GPU memory.

3.1 Design Overview

The design of ORBITFLOW is guided by the following principles:

- **Fine-grained Placements.** To enable fine-grained offloading, ORBITFLOW treats each request-layer combination as a placement choice on either GPU or host memory and accounts for how transfer schedules interact with per-layer compute costs. Because decode steps last only tens of milliseconds, the solver must stay lightweight by exploring the search space efficiently without violating SLOs.
- **Dynamic Reconfiguration.** The system must decide when to reconfigure to balance stability and adaptivity. Reconfiguring after every step closely tracks workload drifts, but frequent data movement and plan switching can offset gains. In contrast, infrequent reconfiguration risks performance degradation and SLO violations.
- **System-level SLO Management.** To meet system-level SLOs, it must accumulate latency margins to mask later stalls and, in cases of memory pressure, pause a long-running, KV-dominant request to guard overall SLOs. It must navigate this global trade-off without inducing starvation, which would affect tail latency.

Guided by these principles, ORBITFLOW operates three key components that form a closed-loop control system, shown in Fig. 4.

Runtime Executor (§3.2) performs each decoding step. It executes computations, installs a KV placement, initiates asynchronous KV transfers based on it, and records compute time, transfer time, and GPU memory usage. To maintain steady user-visible latency in the presence of latency margins, tokens generated within the SLO are buffered in the *Token-Deposit* and released at the SLO rate. The Executor forwards recorded metrics to Adaptive Controller.

Adaptive Controller (§3.4) evaluates whether the placement remains valid. It monitors any batch-dimension mismatch that alters the KV cache size and the ratio between compute and communication, as well as the current placement’s expiry time provided by the Placement Planner, which indicates when token-dimension drift is expected to invalidate the current placement. If any of these conditions is triggered, it requests a new placement from the Planner. When the Planner’s best placement would still violate too many SLOs, the Controller falls back to *Pause-Resume*.




Figure 5: (a) Deposited tokens mask all the SLO violations for a single request. (b) Deposited tokens mask the transient latency spikes caused by a short-lived request with a long prompt. S1, E1, S2, and E2 mark the start and end of Request 1 and Request 2.

Placement Planner (§3.3) finds the best KV placement given the current request lengths, available GPU block capacity, and the present PCIe bandwidth. It balances SLO awareness, solver overhead, and adaptive call frequency. The Planner solves an ILP to choose per-request offload distances that minimize decode latency while respecting memory and SLO constraints. To keep the search space tractable, it applies principled pruning and prevents the exponential explosion that a completely free layout would cause. Solver latency is hidden by running the optimization in parallel with the computation, so the cost does not appear on the critical path. Alongside the new layer-wise KV placement, the Planner also estimates an expiry time—i.e., the specific decoding step at which the growing KV size changes the communication-to-computation ratio (Fig. 2b) or exceeds available GPU memory. This plan is then forwarded to the Executor to install.

3.2 Runtime Executor

Runtime Executor integrates an SLO-aware *Token-Deposit* that decouples token generation from its delivery and a *KV manager* that enforces the KV placements pulled from the *Planner*.

SLO-aware Token-Deposit. Rather than delivering each token as soon as it is generated, ORBITFLOW decouples token generation from its delivery. Doing so allows us to convert any early latency margins into a reserve that can be spent to hide SLO violations. Upon generation, the token is simply appended to *Token-Deposit*. While the deposit is non-empty, the Executor delivers exactly one token per SLO interval. An SLO violation is visible only when the deposit is empty. In this case, the next token is delivered as soon as it is generated. If a request is paused or preempted, the deposit continues to drain, hiding the SLO violations until it empties. Once a request finishes, no further violations can occur, and any remaining tokens are delivered in a single burst.

Fig. 5a illustrates how *Token-Deposit* helps hide SLO violations for a single request. During the early decode iterations, the KV cache is still small, so the generation latency margin remains well below the SLO, and the deposit balance grows steadily. As soon as the generation latency exceeds the SLO, the deposit is drawn down. Although the backend incurs many more SLO violations, the deposit masks a significant portion of them from the user, with no additional computation or memory traffic. The same buffering logic smooths over transient latency spikes, as shown in Fig. 5b. The arrival of a long-prompt, short-output request (i.e., Request 2)

pushes the decode latency over the SLO temporarily, but this spike is mostly absorbed by deposited tokens until it empties.

The Token-Deposit buffer is a CPU-resident FIFO that stores a 4-byte integer for each generated token. Its size is dynamic, but the footprint is negligible, typically only hundreds of KBs. Enqueue and dequeue operations are fast regardless of the buffer size.

KV Manager. The KV manager *enforces* the placement produced by the solver and keeps it consistent throughout decoding. It maintains up-to-date GPU-CPU block tables that record each request’s KV cache address and track the next write location for every decode iteration. When a new placement arrives, the manager computes the delta with respect to the current layout and performs the minimum data movement needed to realize it. For requests whose future layers remain offloaded, the manager writes each freshly generated KV entry directly to its designated slot in host memory, bypassing GPU capacity altogether. Finally, because a new placement can change the split between KV blocks on GPU and the *prefetch buffer*, the manager reallocates GPU memory so that the buffer always spans a *contiguous* range of blocks big enough to hold the worst-case layer transfer dictated by Eq. 1.

3.3 Solver-Based Placement Planner

We first formulate KV offloading as a constrained optimization problem and optimize it with an off-the-shelf solver [19]. Our formulation is guided by three considerations:

(1) **SLO-awareness.** We impose a tunable upper bound on the rate of token-level SLO violations. The solver reports *infeasible* when no placement satisfies this cap. ORBITFLOW then triggers a fallback mechanism that trims the batch and re-optimizes.

(2) **Solver overhead.** We run the solver one decode step ahead and on a separate thread, so it adds no stalls as long as its overhead stays below the decode latency. To keep it small, we shrink the search space to distance-driven placements for each request, without exhaustive searching of all layer-placement combinations.

(3) **Adaptive call frequency.** We let the solver lengthen the interval between invocations when the SLO margin is ample, and shorten it when it is tight, so it is executed only when necessary. Under these three principles, the KV offloading problem reduces to a compact ILP problem that, for each batch of requests, returns a KV placement that satisfies GPU memory limits and SLO guarantees with negligible runtime overhead.

3.3.1 Latency Modeling. We solve for *one* decode iteration for a batch of \mathcal{R} requests that traverse L Transformer layers. The solver’s immutable inputs are (i) request-specific KV block sizes b_r ($r \in \mathcal{R}$), (ii) the peak interconnect bandwidth B , and (iii) per-layer compute times $comp_\ell$ ($\ell = 1, \dots, L$). Its *decision variable* is the binary placement matrix $\mathbf{x} = [x_{r,\ell}] \in \{0, 1\}^{|\mathcal{R}| \times L}$, where $x_{r,\ell} = 1$ keeps the KV of request r at layer ℓ on GPU and $x_{r,\ell} = 0$ offloads it to CPU.

Given a placement \mathbf{x} , Algorithm 1 yields the batch latency $T_{\text{batch}}(\mathbf{x})$. The solver tackles the following optimization problem:

$$\min_{\mathbf{x} \in \{0,1\}^{|\mathcal{R}| \times L}} T_{\text{batch}}(\mathbf{x}),$$

subject to the GPU capacity, decode window, and token-level SLO constraints introduced next.

Algorithm 1 Batch-level decode latency given a KV placement

```

1: Input: KV placement  $x_{r,\ell}$ , block sizes  $b_r$ , peak memory bandwidth  $B$ , per layer compute times  $comp_\ell$ 
2: Output: batch latency  $T_{\text{batch}}$ 
3:  $All \leftarrow \{(r, \ell) \mid x_{r,\ell} = 0\}$  ▷ every required transfer
4:  $Finished \leftarrow \emptyset, Pending \leftarrow All, Cur \leftarrow \emptyset$ 
5:  $Stall_{\text{tot}} \leftarrow 0, Comp_{\text{tot}} \leftarrow 0, comp_{\text{prev}} \leftarrow 0$ 
6: for  $\ell = 1$  to  $L$  do ▷ boundary before layer  $\ell$ 
7:    $Q \leftarrow \text{CANSTART}(Pending, Cur)$ 
8:    $\text{PROMOTE}(Q, \text{from} = Pending, \text{to} = Cur)$ 
9:    $DestL \leftarrow \{e \in Cur \mid e.\text{layer} = \ell\}$  ▷ tasks blocking layer  $\ell$ 
10:  sort  $Cur$  ascending by size
11:   $elapsed \leftarrow 0; n \leftarrow |Cur|; stall_\ell \leftarrow 0; i \leftarrow 1$ 
12:  while  $DestL \neq \emptyset$  do
13:     $\beta \leftarrow B/n$ 
14:     $need \leftarrow Cur_i.\text{size}/\beta$ 
15:     $elapsed \leftarrow elapsed + need$ 
16:    if  $elapsed > comp_{\text{prev}}$  then
17:       $stall_\ell \leftarrow elapsed - comp_{\text{prev}}$ 
18:    end if
19:    if  $Cur_i \in DestL$  then
20:       $DestL \leftarrow DestL \setminus \{Cur_i\}$ 
21:    end if
22:    remove  $Cur_i$  from  $Cur$ ;  $n \leftarrow n - 1$ ;  $i \leftarrow i + 1$ 
23:  end while
24:   $Stall_{\text{tot}} \leftarrow Stall_{\text{tot}} + stall_\ell$ 
25:   $Comp_{\text{tot}} \leftarrow Comp_{\text{tot}} + comp_\ell$ 
26:   $comp_{\text{prev}} \leftarrow comp_\ell$ 
27: end for
28:  $T_{\text{batch}} \leftarrow Comp_{\text{tot}} + Stall_{\text{tot}}$ 

```

GPU Capacity Bound. The sum of all blocks resident on GPU is bounded by the total GPU budget, i.e., the total number of KV blocks that the GPU can hold at any time. We denote this budget by B_{GPU} , which consists of all blocks we keep on GPU and a dedicated *prefetch buffer* reserved for blocks that are fetched just-in-time for the next layer. Formally,

$$\sum_{r \in \mathcal{R}} \sum_{\ell=1}^L b_r x_{r,\ell} + B_{\text{buf}} \leq B_{\text{GPU}}, \quad (1)$$

where the prefetch buffer must be large enough to hold all offloaded blocks required by any *one* layer:

$$B_{\text{buf}} = \max_{\ell \in L} \sum_{r \in \mathcal{R}} (1 - x_{r,\ell}) b_r.$$

Bounded SLO Violations. We formulate the SLO as a hard feasibility check. Concretely, we require that the number of SLO violations of a decode step across a batch of requests \mathcal{R} be at most a tunable cap α :

$$\sum_{r \in \mathcal{R}} \text{sloFail}_r \leq \alpha,$$

where sloFail_r denotes whether a request r violates the SLO in the current decode step and α controls how many violations the solver can tolerate across the batch \mathcal{R} . Under this constraint, latency becomes a meaningful objective: many placements satisfy the same

SLO bound but incur markedly different latencies, with the lower one being more preferable for exploiting Token-Deposit. If we set α to 0, a plan that leads to any SLO violation across requests will be rejected. If no placement satisfies the constraint, the model is declared *infeasible*; the runtime then pauses a request and re-optimizes (see §3.4). Because pausing a request causes at most one violation in each decode step across the batch, we count each pause as one potential violation and set α to 1 in ORBITFLOW. Proactive pausing is preferable in this scenario as it frees both GPU memory and bandwidth, allowing the remaining requests to run under a placement that achieves strictly lower latency than any placement that tries to keep all requests alive.

Self-Determined Decode Window. Invoking the solver at every decode step would neutralize the exact latency we try to minimize. Instead, we let the solver itself decide how long it can safely run unchecked by introducing an integer variable $\Delta \in [\Delta_{\min}, \Delta_{\max}]$, which we denote as the *decode window*. Formally, the window length interacts directly with the SLO feasibility constraint. Instead of enforcing the token-level SLO constraint for the current step, we enforce it on the average over the next Δ steps:

$$\frac{1}{\Delta} \sum_t^{t+\Delta-1} \sum_{r \in \mathcal{R}} \text{sloFail}_r \leq \alpha.$$

Mechanically, if the SLO is tight, the solver chooses a small Δ so that the average is dominated by the imminent steps, the KV placement will be re-optimized sooner. On the other hand, if the latency margin is ample, the solver chooses a larger Δ and amortizes minor violations over a longer period, avoiding redundant re-optimizations. The lower bound Δ_{\min} is a tunable parameter which guarantees that the solver is never invoked more frequently than every Δ_{\min} iterations, as a way to keep the solver time bounded.

Scheduling KV Prefetching. To assess whether a batch meets its token-level SLO, we first need an accurate estimate of compute time and the data-transfer stalls introduced by a given placement. A placement matrix \mathbf{x} specifies, for every request $r \in \mathcal{R}$ and layer $\ell \in \{1, \dots, L\}$, whether that layer’s KV blocks reside on GPU ($x_{r,\ell} = 1$) or must be fetched from CPU ($x_{r,\ell} = 0$). The aggregate blocks that must cross the PCIe bus in one decode iteration are

$$\text{blocks_to_fetch} = \sum_{r \in \mathcal{R}} \sum_{\ell=1}^L b_r (1 - x_{r,\ell}).$$

Because LLM layers execute strictly sequentially, a layer can begin only when both of its inputs are ready: 1) the outputs produced by its previous layer, and 2) the KV cache for that layer. A *stall* occurs when the KV blocks are still in flight while the previous layer has finished its computation. The compute units must idle until the fetching of those blocks completes. Because layers are serialized, a stall at any layer pushes back the start time of all subsequent layers, propagating delay throughout the rest of the decode iteration. To best overlap data transfers with computation, we launch prefetches greedily as soon as their preconditions are satisfied. Because a request can have at most one layer’s KV in flight at any time, it needs only a single asynchronous data transfer stream. For any layer, the total number of active prefetch streams is bounded by the number of requests $|\mathcal{R}|$. A request’s stream is *active* only if all of the following hold: (1) **Current layer’s KV is on GPU**, i.e.

$x_{r,\ell} = 1$. Otherwise, the execution of layer ℓ must reuse the request’s prefetch buffer and launching another transfer could overwrite data in use. (2) **Some later layer’s KV is offloaded**. There exists an $\ell' > \ell$ with $x_{r,\ell'} = 0$; otherwise, all required transfers for this decode iteration are already complete. (3) **The transfer is still in flight**. The stream that fetches the KV cache for layer ℓ' has not finished when layer ℓ begins execution.

We assume all active prefetch streams divide the available bandwidth evenly. When the smallest transfer finishes, its share is redistributed equally among the remaining streams. Algorithm 1 applies this greedy bandwidth-sharing rule to compute the batch-level decode latency for a given KV placement. The batch-level decode latency T_{batch} is the optimization objective of the Planner.

3.3.2 Pruning the Search Space. A completely free layout would force the solver to explore $2^{|\mathcal{R}|L}$ combinations, one binary choice for every request-layer pair. An exhaustive search is prohibitively expensive, whether it is performed offline or online. Therefore, a practical system requires a principled way to prune the search space and focus only on candidates that matter. We achieve this in two steps. First, we drop the uniform requirement. As discussed in §2.2, enforcing a uniform offload strategy for the entire batch lets long requests monopolize GPU memory and push shorter or newer sequences into SLO violations. Second, we retain the distance-driven policy itself. Grouping layers this way aligns identical compute times with identical KV transfers: offloading every N -th layer turns execution into repeating N -layer groups whose compute, communication, and stall patterns are identical, yielding a steady, predictable runtime. The per-request alternative to the 2^L search is:

$$D(L) = \left| \{ \lfloor L/k \rfloor \mid k = 2, \dots, L \} \right| \leq 2 \lfloor \sqrt{L} \rfloor - 1,$$

and the search space size drops from $2^{|\mathcal{R}|L}$ to $D(L)^{|\mathcal{R}|}$. As an example, with $L = 32$ and $|\mathcal{R}| = 4$, we have $D(32) = 9$ possible KV placements. Dropping the distance constraint altogether would raise the space to $2^{128} \approx 3.4 \times 10^{38}$ placements, whereas keeping it but letting each of the $|\mathcal{R}| = 4$ requests choose its own distance yields only $9^4 = 6561$ possibilities (30 orders of magnitude smaller).

Solver Overhead. The ILP nature of the solver makes it NP-hard, which means the worst-case execution time grows exponentially. ORBITFLOW hides this overhead in practice with three measures: 1) The search space is kept small by considering only distance-driven placements for each request, without exhaustive searching of all layer-placement combinations. 2) The solver is launched on the CPU ahead of time to prevent GPU computation delays. We determine when to launch it by comparing the solver overhead to the time per decoding step. With the pruned search space, the solver overhead is usually less than one decoding step. Therefore, we trigger the solver one step ahead and adjust each request’s token count by adding one token to align with the target step. 3) The solver is invoked sparingly, based on the solver-determined window or by the *Adaptive Controller* when batch conditions change.

3.4 Adaptive Controller

The *Adaptive Controller* monitors the execution signals, re-invokes the *Planner* when necessary, and provides a fallback mechanism when the Planner deems the current batch infeasible.




Figure 6: Illustration of flattened 1D KV cache layout in ORBITFLOW modified from PagedAttention.




Figure 7: The length distribution of traces.

Solver Invocations. The solver-chosen decode window Δ already guards against token-dimension drift: it guarantees that the current KV placement remains optimal for at least the next Δ decode iterations *provided* (i) the requests in the current batch stay unchanged and (ii) the profiled performance model is accurate. We keep an on-line monitor that captures signals to re-invoke the solver whenever either assumption breaks earlier than expected:

(1) Profile-mismatch trigger.; After each decode iteration, ORBITFLOW compares the iteration latency predicted by Algorithm 1 with the latency it actually observes. If the absolute difference exceeds a configurable fraction of the predicted value, ORBITFLOW interprets the gap as evidence of unexpected GPU or bandwidth contention. It then updates the profiled per-layer compute times and effective bandwidth using an exponentially weighted moving average and immediately re-runs the solver with the refreshed profile.

(2) Batch-change trigger. The arrival of a new request or completion of an existing one alters the batch composition \mathcal{R} , causing a batch-dimension mismatch between the current placement and the new batch composition. The solver is re-run with the updated batch so that KV placement reflects the fresh load mix.

Both triggers are infrequent. Profile deviation surfaces only under contention, and batch changes can occur at most once per request completion or arrival.

Pause-Resume Fallback. When the solver reports *infeasible* under the GPU capacity and SLO bounds, ORBITFLOW abandons the “serve-all” stance and proactively *pauses* one request instead of forcing an extreme offload pattern that would cause a cascade of SLO violations. The victim is the request with the largest sum of KV blocks and deposited tokens. Selecting such a request prevents any single long request from monopolizing GPU memory while masking as many SLO violations as possible for the victim. Unlike reactive *preemption*, which is triggered when the GPU memory constraint is violated and therefore drops a request’s entire KV cache in one shot, ORBITFLOW pauses a request *before* memory becomes tight. Because the GPU still has headroom at the moment of pausing, ORBITFLOW evicts the request’s KV blocks *on demand and at the layer-level*: its KV blocks on GPU are dropped layer by layer when their space is actually needed. This on-demand eviction makes a later resume much cheaper than repatriating an all-or-nothing dump, requiring at most those evicted blocks. A request can be paused and resumed multiple times. ORBITFLOW resumes a paused request only after some in-flight requests have completed and a feasible plan exists (i.e., the transient overload has cleared). If the server becomes overloaded again, ORBITFLOW chooses a new

or previously paused request as a victim. This gating mechanism ensures no single request monopolizes resources and starves others.

3.5 Extension to Distributed Inference

Distributed inference becomes necessary when KV offloading on a single GPU cannot fully relieve memory pressure. We currently focus on single-node, multi-GPU serving where a single model is configured to use all GPUs within a node. In this scenario, tensor parallelism (TP) [46], which symmetrically shards weights and computations across the GPUs, is the most suitable among common parallelism strategies. Pipeline parallelism (PP) [21] can also be adopted to reduce per-GPU memory demand, but it depends on microbatches to hide pipeline bubbles. Unlike training, serving handles requests that arrive stochastically, so microbatches rarely stay full and pipeline stages can sit idle.

We extend ORBITFLOW to support TP and validate its effectiveness experimentally in Fig. 11. With symmetric sharding and shared memory bandwidth, a KV placement computed for one worker is expected to be optimal for all workers. ORBITFLOW runs the solver on one worker and broadcasts the resulting placement to the others. The Runtime Executor of each worker installs the placement synchronously, thereby maintaining consistent KV placement across workers. We note that PP and data parallelism (DP) [10] are typically more desirable choices than TP when serving scales to multiple nodes with slower interconnects. Since PP incurs lower communication costs than TP, a single model instance can be organized with PP across nodes and TP within each node. DP is used to scale out further by coordinating multiple model instances via a global scheduler. We leave extension to these scenarios as future work.

4 IMPLEMENTATION

We build ORBITFLOW on top of vLLM v0.6.6 [29] with three core changes. First, we enable request-wise KV placement by modifying vLLM’s cache engine to flatten the original per-layer KV layout into one contiguous 1D layout, indexed by a layer-aware 2D table. As illustrated in Fig. 6, this allows flexible allocation of a different set of layers for each request, unlike vLLM’s default KV cache management, which uniformly moves every N^{th} layer for all requests. We also embed the *KV manager* in vLLM’s cache engine, which pulls the placement decision from the Planner and issues necessary data movements to realize the placement. Second, we enable *Pause-Resume* by tweaking the scheduler to temporarily drop a running request. To optimize Pause-Resume, we implement the *removable cache* to manage paused requests efficiently. Rather

than immediately offloading all GPU-resident layers of paused requests, we incrementally move layers to CPU only when additional GPU space is needed, reducing data transfer overhead and resume latency. Third, the *Placement Planner* runs in a separate thread and always solves one decode iteration ahead. The cache engine pulls the finished plan in the next iteration and applies the changes, without device-wide synchronization.

5 EVALUATION

We evaluate ORBITFLOW against baseline offloading systems, including DeepSpeed [11], FlexGen [51] and SLO-aware Offloading [35], on ShareGPT-derived synthetic traces across different arrival rates. Our evaluation shows that ORBITFLOW outperforms the baselines, achieving 62% and 66% higher TPOT and TBT SLO attainment, and up to 3.3× higher throughput. Additionally, ORBITFLOW consistently provides performance improvements under various context lengths, batch sizes, and distributed serving scenarios.

5.1 Experimental Setup

Models and Hardware Settings. We evaluate ORBITFLOW using LLaMA3 models [17]. Our default configuration uses the 8B model on a single NVIDIA RTX A5000 (24GB) GPU with 384GB of host memory over PCIe 3.0×16. For longer contexts up to 128K tokens, we also report results on LLaMA3-70B using a 4-GPU node with RTX A6000 (48GB) and 256GB of host memory over PCIe 4.0×16.

Workloads. Because publicly available production traces for long-context serving are limited, we synthesize workloads by sampling requests from ShareGPT [43]. We use Poisson arrivals with varying arrival rates to emulate real-world LLM serving scenarios. The generated traces span a wide spectrum of short to long inputs and outputs, as illustrated in Fig. 7. Unless otherwise specified, we use a batch size of 4 and limit the total sequence length per batch to 32K tokens.¹ For design validations and sensitivity studies, we fix the arrival rate at 0.97 req/min and vary individual parameters including burstiness of arrivals, batch size, and context length.

Key Metrics. We focus on two complementary latency SLOs. TBT measures the latency between consecutive output tokens and thus captures user-perceived interactivity at the token level. TPOT measures the average per-token latency. We set the base SLO by profiling the decoding latency of the longest request that can be served without offloading, and scale it to create progressively looser budgets for fair comparison across latency constraints. Unless otherwise specified, experiments use an SLO scale of 1.5, representing a mid-level constraint. We also report throughput, requests completed per minute, and end-to-end (E2E) latency, and the average per-request completion time, to capture overall serving efficiency.

Baselines. We consider five baseline KV cache offloading strategies, all implemented on the same codebase for a fair comparison.

(1) **DeepSpeed-Inference** [2] offloads all KV caches to the CPU and retains only the current layer’s KV cache in GPU. It overlaps data transfers with computation by prefetching the next layer’s KV cache during the current layer’s execution.

(2) **FlexGen** [51] maximizes throughput by profiling a fixed, optimal KV placement offline, assuming peak GPU performance. It does not consider SLO constraints or reconfigure at runtime.

(3) **FlexGen+** is an enhanced version of FlexGen that dynamically reconfigures its KV placement when requests join and leave, allowing it to adapt to batch-dimension mismatches.

(4) **SLO-aware Offloading** [35] adjusts KV cache placement based on offline profiling while maximizing throughput within SLO constraints. It still employs uniform offload distances across requests, and reconfiguration is confined to batch boundaries, similar to FlexGen+.

(5) **Dynamic Heuristic** is a variant of ORBITFLOW that replaces the solver-based planner with a heuristic that minimizes offloaded layers to reduce communication costs. It reconfigures at the iteration level but lacks the solver’s fine-grained optimization.

DeepSpeed-Inference and FlexGen use *static* placements determined offline, while the remaining three baselines *adaptively* reconfigure at runtime.

5.2 Performance Analysis

SLO Attainment. We evaluate the effectiveness of ORBITFLOW by examining its performance in terms of TBT and TPOT SLO attainment compared to the baselines under varying request arrival rates and SLO scales, as illustrated in Fig. 8a.

We observe that ORBITFLOW consistently outperforms the baselines in both TBT and TPOT SLO attainment, maintaining over 75% TBT attainment under high request rates and a tight SLO scale. Among the *static* methods, DeepSpeed-Inference shows the weakest performance, as its layer-wise offloading stalls computation with frequent KV transfers that cannot be hidden within a single layer’s compute time. FlexGen does not reconfigure at runtime or account for SLO constraints, so its static placement becomes increasingly suboptimal as the KV cache grows during decoding. The *adaptive* baselines perform better but remain limited. FlexGen+ improves over FlexGen by reconfiguring whenever the batch composition changes, accounting for batch-dimension mismatch. SLO-aware Offloading adjusts offload distances in response to the SLO target—offloading more aggressively under loose constraints and reducing offload distances as the SLO tightens—but still applies uniform distances across requests, limiting its adaptability when requests in the same batch have different KV cache sizes. Dynamic Heuristic performs best among the baselines, yet its greedy approach lacks the precision of solver-based planning to sustain high attainment.

Under relaxed latency constraints (SLO scale=2.5), most baselines except DeepSpeed-Inference maintain stable attainment at lower arrival rates. However, attainment drops to around 50% as the arrival rate increases. This degradation accelerates when the SLO scale is tightened further. ORBITFLOW experiences the same trend, but demonstrates considerably better robustness. Its solver-driven, adaptive fine-grained approach adapts to both token-dimension drifts and batch-dimension mismatch, the sources of performance degradation that static policies cannot accommodate. TPOT attainment penalizes any per-token latency spike more heavily, as TPOT evaluates average latency at the request level.

¹On a single RTX A5000 (24GB), an 8B model uses about 16GB of GPU memory, and PyTorch runtime consumes ~6GB. The remaining ~2GB can hold roughly 16K tokens. With offloading, the system can accommodate up to 32K tokens per batch without memory oversubscription.

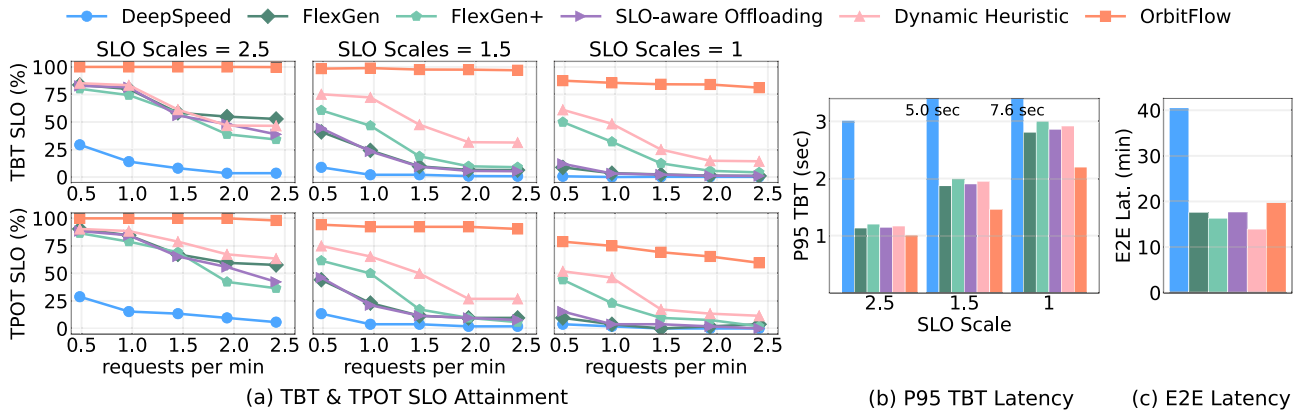


Figure 8: (a) TBT and TPOT SLO attainment across arrival rates and SLO scales. Y-axis indicates the SLO attainment, and X-axis indicates the arrival rate. (b) Comparison of P95 TBT latency across varying SLO scales. (c) Comparison of end-to-end latency.




Figure 9: Comparison of (a) P99 TBT latency. (b) P99 TPOT latency. (c) Throughput. (d) Average GPU memory utilization.

Tail Latency (P95 and P99). Beyond average SLO attainment, we examine how ORBITFLOW performs under worst-case conditions. Fig. 8b reports the P95 latency across SLO scales. ORBITFLOW consistently demonstrates the lowest latency. At an SLO scale of 1, ORBITFLOW achieves 21.7% and 68.6% lower P95 latency than FlexGen and DeepSpeed-Inference, respectively. This gap narrows as the SLO loosens, but ORBITFLOW maintains the lowest P95 latency across all scales. At the extreme-tail latencies (P99) under SLO scale of 1.5, ORBITFLOW achieves a 65% lower P99 TBT and a 64% lower P99 TPOT compared to DeepSpeed-Inference, as shown in Fig. 9a-b. Against other baselines, the improvement is a consistent 3–9% in both metrics. The margin narrows at P99 because rare, burst-induced stalls affect all methods similarly. Nonetheless, ORBITFLOW still delivers the lowest tail latency overall.

E2E Latency and Throughput. As shown in Fig. 8c, ORBITFLOW’s E2E latency is 12–21% higher than the adaptive baselines, a direct consequence of the Pause-Resume mechanism, which defers long requests to keep the batch feasible under memory and SLO constraints. Dynamic Heuristic—which does not use Pause-Resume—achieves the lowest E2E latency among all methods, confirming that the additional E2E cost is from the deferred requests. However, this is a trade-off for system-level efficiency. Fig. 9d shows that ORBITFLOW achieves the highest throughput, outperforming DeepSpeed-Inference by 3.3× and other baselines by 4–52%. By deferring stragglers that would otherwise slow down the entire batch, Pause-Resume allows the remaining requests to decode faster, improving both SLO attainment and throughput at the cost of longer completion time for individual deferred requests.

Memory Usage. ORBITFLOW’s throughput gains rely on effective use of limited GPU memory. Fig. 9c shows the average GPU utilization across methods. The static methods under-utilize GPU memory: DeepSpeed-Inference keeps at most one layer on the GPU at a time, whereas FlexGen employs a static offloading decision based on the maximum expected batch size. Among the adaptive methods: SLO-aware Offloading prioritizes CPU placement conservatively, resulting in utilization similar to FlexGen. FlexGen+ shows improvement over FlexGen by recomputing the offload distances upon batch composition changes. Dynamic Heuristic closely matches the high utilization achieved by ORBITFLOW, as it aggressively maximizes GPU memory usage.

5.3 Sensitivity Studies

We study ORBITFLOW’s sensitivity to different workloads and serving scenarios. We first evaluate longer contexts and multi-GPU setups under tensor parallelism, then vary arrival burstiness and batch size on the default single-GPU setup. Finally, we evaluate ORBITFLOW under an alternative request scheduling policy.

Context Length Variation. Increasing context length magnifies KV cache demands due to significant token-dimension drift and exacerbates batch-level imbalance, as requests with widely varying lengths are in the same batch. Fig. 10 compares ORBITFLOW with FlexGen+ and Dynamic Heuristic—the two strongest baselines from Fig. 8—at 8K, 32K, and 128K sequence lengths on LLaMA3-70B with 2-way TP. While FlexGen+ and Dynamic Heuristic achieve near 100% SLO attainment with 8K tokens at an SLO scale of 2.5, their performance degrades as context length grows, dropping to 69.2% and 32.9% at 128K tokens. In contrast, ORBITFLOW consistently achieves above 85% SLO attainment across all context lengths.

Tensor Parallelism. Tensor parallelism (TP) distributes model weights and KV caches across GPUs, requiring frequent all-reduce operations. Without NVLink, these operations rely on PCIe, directly competing with KV cache offloading and causing PCIe contention. Fig. 11 evaluates ORBITFLOW against FlexGen+, SLO-aware Offloading, and Dynamic Heuristic under 2-way and 4-way TP. Baseline methods degrade sharply under this contention due to their static placement strategies, whereas ORBITFLOW sustains higher TBT SLO




Figure 10: Comparison of TBT SLO attainment under various context lengths.




Figure 11: Comparison under 2-way and 4-way TP.




Figure 12: Sensitivity studies with varying (a) CV scales and (b) batch sizes.




Figure 13: Effect of SRTF scheduling on ORBITFLOW.

attainment even under stringent SLOs by dynamically adapting its placement to the reduced available bandwidth.

Burstiness Variation. As shown in Fig. 12a, we vary the coefficient of variation (CV) of request arrivals to control burstiness. Higher CV produces more bursty traffic, increasing queuing delay and the risk of SLO violations. ORBITFLOW maintains TBT SLO attainment above 90% as CV increases, while FlexGen+ and Dynamic Heuristic degrade steadily. When CV exceeds 6, all methods flatten as bursts become intense enough to saturate the GPU. Extra burstiness only lengthens the time requests wait in the queue.

Batch Size Variation. We evaluate ORBITFLOW under different batch sizes, as shown in Fig. 12b. Dynamic Heuristic and FlexGen+ fail to scale since larger batches amplify drifts in both dimensions. The token dimension grows more rapidly as the total KV cache size increases with batch size, and the batch dimension becomes more volatile due to frequent request completions and arrivals. In contrast, ORBITFLOW effectively manages such fluctuations, maintaining consistently high performance regardless of the batch size.

Effect of Request Scheduling. ORBITFLOW uses the first-come, first-served (FCFS) scheduling policy in vLLM by default. We further examine the performance of ORBITFLOW under an alternative shortest-remaining-time-first (SRTF) policy, which sorts the requests in the waiting queue by their predicted output length at each admission, following [14]². SRTF prioritizes requests with shorter outputs and tends to batch requests with similar output lengths, thereby helping alleviate the batch-dimension mismatch. This reduces the average and P95 TTFT of ORBITFLOW by 40% and 24%, respectively, compared to FCFS, as shown in Fig. 13a. However, SRTF inevitably clusters long requests together, forcing more KV caches to be offloaded and increasing data transfer costs, which eventually translate into long decode stalls. As a result, ORBITFLOW+SRTF achieves TPOT SLO attainment of 92.4%, which is 4.8% lower than ORBITFLOW. SRTF also leads to a 29% increase in P95 TPOT and an 8% drop in overall throughput. Despite these trade-offs, ORBITFLOW maintains over 90% TPOT SLO attainment,

²We use the same trace as [14], scaled by $6\times$ to align with our long-context setting.




Figure 14: (a) Impact of ORBITFLOW’s components. (b) Effectiveness of Token-Deposit mechanism (c) Comparison of fallback victim selection.

confirming that its runtime adaptation remains effective across different scheduling policies.

5.4 Design Validation and Overheads

Impact of Individual Techniques. To assess the contributions of ORBITFLOW’s core components, we evaluate ORBITFLOW by adding each component incrementally, as shown in Fig. 14a. *Batch-Uniform* applies a uniform offload distance to all requests within a batch, achieving 45.1% SLO attainment at the tightest SLO scale of 1. Building on this, *+Request-Wise* introduces fine-grained, per-request offloading decisions optimized by the solver, improving SLO attainment to 54.5% by better balancing KV cache placement across heterogeneous requests. Adding *Pause-Resume* further boosts SLO attainment to 71.0% by temporarily deferring requests with large KV caches to free GPU memory for others. Finally, ORBITFLOW adds *Token-Deposit*, further raising SLO attainment to 85.6%. The same trend persists at a looser SLO scale of 2.5, with each component contributing consistent improvements.

Token-Deposit on Baselines. Fig. 14b adds *Token-Deposit* to FlexGen+ and Dynamic Heuristic to evaluate its generality. At a loose SLO scale of 2.5, *Token-Deposit* improves TBT attainment by 18%, 11%, and 5% for FlexGen+, Dynamic Heuristic, and ORBITFLOW, respectively. As the SLO tightens, the benefit grows for Dynamic Heuristic and ORBITFLOW, which gain a further 35% and 20%, as their adaptive placement generates early tokens fast enough to accumulate a buffer, whereas FlexGen+ gains little benefit.

Fallback Strategy. When *Pause-Resume* is triggered, ORBITFLOW must select which request to pause. We evaluate three approaches: pausing the longest, shortest, or a randomly chosen request from the batch. We study these policies using a stress-test trace with an increased fraction of long requests to increase overload frequency and trigger more fallback events. Because KV sizes grow linearly with sequence length, pausing the longest request releases the most GPU memory, allowing the remaining requests to continue without stalls, as confirmed in Fig. 14c.

Trace	Calls/Req.	E2E (%)	TBT (%)
Both Static	1.25	0.27	23.47
Token Dynamic	11.92	0.37	31.65
Both Dynamic	48.33	0.64	28.49

Table 1: Solver overhead across different workloads.

Component-wise Runtime Overhead. We break down the total execution time into *scheduler* (5.28%), *prefill* (7.04%), *decode* (87.23%), *KV Manager* (0.45%), and *Placement Planner* (0.02%), with the last two representing overheads introduced by ORBITFLOW. Table 1 further details the Planner’s overhead across three workloads with increasing dynamism: *Both Static* uses static batching with short outputs, limiting changes in both dimensions; *Token Dynamic* uses static batching but with long outputs, so KV caches grow along the token dimension while the batch stays fixed; *Both Dynamic* combines long outputs with continuous batching, varying both dimensions simultaneously. We report the solver’s accumulated share of total execution time (E2E%) and its per-invocation wall-time relative to one decode step (TBT%). For *Both Dynamic*, the solver is triggered approximately 4× and 39× more often than the static workloads, reflecting frequent changes in batch composition and KV cache size. Even in this volatile trace, E2E% is only 0.64%. Per invocation, TBT% stays well below 100%. As we launch the solver one iteration ahead on a separate thread, this cost can be fully hidden under computation. In rare cases, however, the solver exceeds one decode step, visible as the long tail in Fig. 15c.

Search Space Pruning. The low overhead is enabled by pruning the search space. Our solver restricts placements to evenly spaced offloaded layers, parameterized by a single offload distance. For a 32-layer model, distance 4 offloads every 4th layer (8 total) and distance 5 offloads every 5th layer (6 total). Counts between these, such as 7 offloaded layers, require uneven spacing and are never explored. To evaluate the impact of this restriction, we compare our solver, *Pruned*, against *Full* which allows arbitrary KV placements. We sample 2000 solver inputs from the evaluation traces and compare the placements. Fig. 15a shows that 80.9% of *Full*’s placements match *Pruned* in decode latency, with 56% identical and 24.9% differing only in which layers are offloaded. Only 16.7% are strictly better. Moreover, 98% of the strictly better cases arise from intermediate layer counts that even spacing skips, as we expected. The median speedup in these cases is just 2.29% (Fig. 15b), while closing this gap would require 7.74× more solver time (Fig. 15c) on average, which would itself violate the SLO.

6 RELATED WORK

KV Offloading. InfiniGen [30] speculatively prefetches critical KV entries to minimize CPU-to-GPU transfer overhead. Neo [25] and FlexInfer [39] extend memory capacity by offloading KV caches to host memory or SSD. These methods primarily optimize the spatial placement and movement of KV data to reduce transfer latency. ORBITFLOW follows a similar motivation but performs placement decisions online, adapting KV allocations at runtime according to batch composition and token-level latency feedback.

Pruning and Compression. Structural pruning and quantization reduce KV cache footprints on GPUs. StreamingLLM [49] and




Figure 15: Effect of search space pruning of ORBITFLOW.

H2O [54] prune less critical KV entries, KIVI [33] employs asymmetric quantization, and MiniCache [32] merges similar KV states across layers. Although effective in reducing memory usage, these methods risk degrading model accuracy.

Recomputation. Recomputation rematerializes KV states rather than transferring them, but prior work has focused on the prefill stage for reusing historical prompt contexts. HCache [45] reconstructs KV states with activations that are lightweight in MHA models. CachedAttention [15] reuses or recomputes KV states across conversation turns to avoid redundant transfers. Although such recomputation could be applied to the decode stage, its effectiveness is still limited to MHA models. ORBITFLOW is primarily evaluated on modern GQA models, where activations are no longer lightweight. However, in MHA models, combining KV offloading with recomputation may yield synergistic effects.

Request Scheduling. Request scheduling techniques improve GPU utilization and batching efficiency during inference. Orca [20] introduces continuous batching to insert new requests dynamically as others finish, and Sarathi-Serve [1] integrates prefill and decode using stall-free chunked batching. Disaggregated systems such as DistServe [56] decouple prefill and decode across GPU clusters, while Helix [37] and Hetegen [55] explore heterogeneous GPU graphs for model- and tensor-parallel pipelines. ORBITFLOW complements these methods by optimizing latency within active batches through fine-grained, layer-wise KV cache placement with a solver that adapts to token- and batch-level drifts.

7 CONCLUSION

ORBITFLOW provides an effective framework for long-context LLM serving, dynamically reconfiguring KV cache placement per request using a solver to meet stringent latency SLOs. Evaluations on LLaMA3-8B show that ORBITFLOW improves TPOT and TBT SLO attainment by 62% and 66% under the most demanding conditions, reduces P95 TBT latency by 38% over the best competing systems, and achieves up to 3.3× higher throughput compared to state-of-the-art offloading methods. These gains come with less than 1% runtime overhead and persist under multi-GPU settings.

ACKNOWLEDGMENTS

This work was supported by SAMSUNG Research, Samsung Electronics Co., Ltd., and by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grants (No. RS-2025-02304554, Efficient and Scalable Framework for AI Heterogeneous Cluster Systems; No. RS-2019-II191906, Artificial Intelligence Graduate School Program (POSTECH)), the National Research Foundation of Korea (NRF) grant (No. RS-2024-00354947), and the ETRI grant (No. 26ZS1100, Development of Large-Scale Parallel Computing Technology for Generative AI), funded by the Korean government (MSIT).

REFERENCES

- [1] Agrawal Amey, Kedia Nitin, Panwar Ashish, Mohan Jayashree, Kwatra Nipun, Gulavani Bhargav S, Tumanov Alexey, and Ramjee Ramachandran. Taming Throughput-Latency Tradeoff in LLM Inference with Sarathi-Serve. In *OSDI*, pages 117–134, 2024.
- [2] Reza Yazdani Aminabadi, Samyam Rajbhandari, Minjia Zhang, Ammar Ahmad Awan, Cheng Li, Du Li, Elton Zheng, Jeff Rasley, Shaden Smith, Olatunji Ruwase, and Yuxiong He. DeepSpeed Inference: Enabling Efficient Inference of Transformer Models at Unprecedented Scale. *arXiv preprint arXiv:2207.00032*, 2022.
- [3] AI Anthropic. The Claude 3 Model Family: Opus, Sonnet, Haiku. *Claude-3 Model Card*, 2024.
- [4] Patke Archit, Reddy Dhemath, Jha Saurabh, Qiu Haoran, Pinto Christian, Narayanaswami Chandra, Kalbarczyk Zbigniew, and Iyer Ravishankar. Queue Management for SLO-Oriented Large Language Model Serving. In *Proceedings of the 2024 ACM Symposium on Cloud Computing*, SoCC '24, page 18–35, 2024.
- [5] Acharya Arkadeep, Singh Brijraj, and Onoe Naoyuki. LLM Based Generation of Item-Description for Recommendation System. In *Proceedings of the 17th ACM Conference on Recommender Systems*, RecSys '23, page 1204–1207, 2023.
- [6] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. Qwen Technical Report, 2023.
- [7] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language Models are Few-Shot Learners. *NeurIPS*, 33:1877–1901, 2020.
- [8] Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen, Keyu Chen, Xin Chen, Xun Chen, Zehui Chen, Zhi Chen, Pei Chu, Xiaoyi Dong, Haodong Duan, Qi Fan, Zhaoye Fei, Yang Gao, Jiaye Ge, Chenyu Gu, Yuzhe Gu, Tao Gui, Aijia Guo, Qipeng Guo, Conghui He, Yingfan Hu, Ting Huang, Tao Jiang, Penglong Jiao, Zhenjiang Jin, Zhikai Lei, Jiaxing Li, Jingwen Li, Linyang Li, Shuaibin Li, Wei Li, Yining Li, Hongwei Liu, Jiaoning Liu, Jiawei Hong, Kaiwen Liu, Kuikun Liu, Xiaoran Liu, Chengqi Lv, Haijun Lv, Kai Lv, Li Ma, Runyuan Ma, Zerun Ma, Wenchang Ning, Linke Ouyang, Jiantao Qiu, Yuan Qu, Fukai Shang, Yunfan Shao, Demin Song, Zifan Song, Zhihao Sui, Peng Sun, Yu Sun, Huanze Tang, Bin Wang, Guoteng Wang, Jiaqi Wang, Jiayu Wang, Rui Wang, Yudong Wang, Ziyi Wang, Xingjian Wei, Qizhen Weng, Fan Wu, Yingtong Xiong, Chao Xu, Ruijiang Xu, Hang Yan, Yirong Yan, Xiaogui Yang, Haochen Ye, Huaiyuan Ying, Jia Yu, Jing Yu, Yuhang Zang, Chuyu Zhang, Li Zhang, Pan Zhang, Peng Zhang, Ruijie Zhang, Shuo Zhang, Songyang Zhang, Wenjian Zhang, Wenwei Zhang, Xingcheng Zhang, Xinyue Zhang, Hui Zhao, Qian Zhao, Xiaomeng Zhao, Fengzhe Zhou, Zaida Zhou, Jingming Zhuo, Yicheng Zou, Xipeng Qiu, Yu Qiao, and Dahua Lin. InternLM2 Technical Report, 2024.
- [9] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. PaLM: Scaling Language Modeling with Pathways. *Journal of Machine Learning Research*, 24(240):1–113, 2023.
- [10] Jeffrey Dean, Greg S. Corrado, Rajat Monga, Kai Chen, Matthieu Devin, Quoc V. Le, Mark Z. Mao, Marc'Aurelio Ranzato, Andrew Senior, Paul Tucker, Ke Yang, and Andrew Y. Ng. Large Scale Distributed Deep Networks. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'12, 2012.
- [11] DeepSpeed Model Implementations for Inference. <https://github.com/microsoft/DeepSpeed-MII>.
- [12] Jiayu Ding, Shuming Ma, Li Dong, Xingxing Zhang, Shaohan Huang, Wenhui Wang, Nanning Zheng, and Furu Wei. LongNet: Scaling Transformers to 1,000,000,000 Tokens. *arXiv preprint arXiv:2307.02486*, 2023.
- [13] Qichen Fu, Minsik Cho, Thomas Merth, Sachin Mehta, Mohammad Rastegari, and Mahyar Najibi. LazyLLM: Dynamic token pruning for efficient long context LLM inference. In *Workshop on Efficient Systems for Foundation Models II @ ICML2024*, 2024.
- [14] Yichao Fu, Siqi Zhu, Runlong Su, Aurick Qiao, Ion Stoica, and Hao Zhang. Efficient llm scheduling by learning to rank. In *Proceedings of the 38th International Conference on Neural Information Processing Systems*, NIPS '24, 2024.
- [15] Bin Gao, Zhuomin He, Puru Sharma, Qingxuan Kang, Djordje Jevdjic, Junbo Deng, Xingkun Yang, Zhou Yu, and Pengfei Zuo. Cost-Efficient Large Language Model Serving for Multi-turn Conversations with CachedAttention. In *2024 USENIX Annual Technical Conference (USENIX ATC 24)*, pages 111–126, 2024.
- [16] Izacard Gautier and Grave Edouard. Leveraging Passage Retrieval with Generative Models for Open Domain Question Answering. In *EACL 2021 - 16th Conference of the European Chapter of the Association for Computational Linguistics*, pages 874–880, 2021.
- [17] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, et al. The Llama 3 Herd of Models. *arXiv preprint arXiv:2407.21783*, 2024.
- [18] Mandy Guo, Joshua Ainslie, David Uthus, Santiago Ontanon, Jianmo Ni, Yun-Hsuan Sung, and Yinfei Yang. LongT5: Efficient text-to-text transformer for long sequences. In Marine Carpuat and Ivan Vladimir Meza Ruiz, editors, *Findings of the Association for Computational Linguistics: NAACL 2022*, 2022.
- [19] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2024.
- [20] Yu Gyeong-In, Jeong Joo Seong, Kim Geon-Woo, Kim Soojeong, and Chun Byung-Gon. Orca: A Distributed Serving System for Transformer-Based Generative Models. In *OSDI*, pages 521–538, 2022.
- [21] Yanping Huang, Youlong Cheng, Ankur Bapna, Orhan Firat, Mia Xu Chen, Dehao Chen, HyoukJoong Lee, Jiquan Ngiam, Quoc V. Le, Yonghui Wu, and Zhifeng Chen. *GPipe: Efficient Training of Giant Neural Networks Using Pipeline Parallelism*. Curran Associates Inc., Red Hook, NY, USA, 2019.
- [22] Wei Jason, Wang Xuezhai, Schuurmans Dale, Bosma Maarten, ichter brian, Xia Fei, Chi Ed, Le Quoc V, and Zhou Denny. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837, 2022.
- [23] Siddharth Jha, Coleman Richard Charles Hooper, Xiaoxuan Liu, Sehoon Kim, and Kurt Keutzer. Learned Best-Effort LLM Serving. In *Workshop on Efficient Systems for Foundation Models II @ ICML2024*, 2024.
- [24] Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. Mistral 7B, 2023.
- [25] Xuanlin Jiang, Yang Zhou, Shiyi Cao, Ion Stoica, and Minlan Yu. NEO: Saving GPU Memory Crisis with CPU Offloading for Online LLM Inference. In *Proceedings of the 8th Conference on Machine Learning and Systems (MLSys 2025)*, 2025. Available at OpenReview.
- [26] Zhang Jingqing, Zhao Yao, Saleh Mohammad, and Liu Peter. PEGASUS: Pre-training with Extracted Gap-sentences for Abstractive Summarization. In *Proceedings of the 37th International Conference on Machine Learning*, Proceedings of Machine Learning Research, pages 11328–11339, 2020.
- [27] Achiam Josh, Adler Steven, Agarwal Sandhini, Ahmad Lama, Akkaya Ilge, Aleman Florencia Leoni, Almeida Diogo, Altschmidt Janko, Altman Sam, Anadkat Shyamal, et al. GPT-4 Technical Report. *arXiv preprint arXiv:2303.08774*, 2023.
- [28] Li Kunze and Zhang Yu. Planning First, Question Second: An LLM-Guided Method for Controllable Question Generation. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 4715–4729, 2024.
- [29] Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Haoteng Zhang, and Ion Stoica. Efficient Memory Management for Large Language Model Serving with PagedAttention. In *SOSP*, pages 611–626, 2023.
- [30] Wonbeom Lee, Jungi Lee, Junghwan Seo, and Jaewoong Sim. InfiniGen: Efficient Generative Inference of Large Language Models with Dynamic KV Cache Management. In *18th USENIX Symposium on Operating Systems Design and Implementation (OSDI 24)*, pages 155–172, 2024.
- [31] Xue Linting, Constant Noah, Roberts Adam, Kale Mihir, Al-Rfou Rami, Siddhant Aditya, Barua Aditya, and Raffel Colin. mT5: A Massively Multilingual Pre-trained Text-to-Text Transformer. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, 2021.
- [32] Akide Liu, Jing Liu, Zizheng Pan, Yefei He, Gholamreza Haffari, and Bohan Zhuang. MiniCache: KV Cache Compression in Depth Dimension for Large Language Models. In *Proceedings of the 38th International Conference on Neural Information Processing Systems*, NIPS '24, 2024.
- [33] Zirui Liu, Jiayi Yuan, Hongye Jin, Shaochen Zhong, Zhaozhuo Xu, Vladimir Braverman, Beidi Chen, and Xia Hu. KIVI: A tuning-free asymmetric 2bit quantization for KV cache. In *Proceedings of the 41st International Conference on Machine Learning*. PMLR, 2024.
- [34] Yinquan Lu, Wenhao Zhu, Lei Li, Yu Qiao, and Fei Yuan. LLaMAX: Scaling linguistic horizons of LLM by enhancing translation capabilities beyond 100 languages. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, 2024.
- [35] Chenxiang Ma, Zhisheng Ye, Hanyu Zhao, Zehua Yang, Tianhao Fu, Jiaxun Han, Jie Zhang, Yingwei Luo, Xiaolin Wang, Zhenlin Wang, Yong Li, and Diyu Zhou. Memory Offloading for Large Language Model Inference with Latency SLO Guarantees. *arXiv preprint arXiv:2502.08182*, 2025.
- [36] Brysbaert Marc. How many words do we read per minute? A review and meta-analysis of reading rate. *Journal of Memory and Language*, 109:104047, 2019.
- [37] Yixuan Mei, Yonghao Zhuang, Xupeng Miao, Juncheng Yang, Zhihao Jia, and Rashmi 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, ASPLOS '25, 2025.
- [38] Lewis Mike, Liu Yinhan, Goyal Naman, Ghazvininejad Marjan, Mohamed Abdelrahman, Levy Omer, Stoyanov Veselin, and Zettlemoyer Luke. BART: Denoising

- Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, 2020.
- [39] Seonjin Na, Geonhwa Jeong, Byung Hoon Ahn, Aaron Jezghani, Jeffrey Young, Christopher J. Hughes, Tushar Krishna, and Hyesoon Kim. FlexInfer: Flexible LLM Inference with CPU Computations. In *Proceedings of the 8th Conference on Machine Learning and Systems (MLSys 2025)*, 2025. Available at OpenReview.
- [40] Patel, Pratyush and Choukse, Esha and Zhang, Chaojie and Shah, Aashaka and Goiri, Ínigo and Maleki, Saeed and Bianchini, Ricardo. Splitwise: Efficient Generative LLM Inference Using Phase Splitting. In *2024 ACM/IEEE 51st Annual International Symposium on Computer Architecture (ISCA)*, 2024.
- [41] Lewis Patrick, Perez Ethan, Piktus Aleksandra, Petroni Fabio, Karpukhin Vladimir, Goyal Naman, Küttler Heinrich, Lewis Mike, Yih Wen-tau, Rocktäschel Tim, Riedel Sebastian, and Kiela Douwe. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474, 2020.
- [42] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language Models are Unsupervised Multitask Learners. In *OpenAI Blog*. OpenAI, 2019.
- [43] ShareGPT Contributors. ShareGPT_V3_unfiltered_cleaned_split conversations dataset. https://huggingface.co/datasets/anon8231489123/ShareGPT_Vicuna_unfiltered, 2023. Version V3 (unfiltered, cleaned, split). Apache-2.0 license. Accessed 2025-06-02.
- [44] Sheng, Ying and Cao, Shiyi and Li, Dacheng and Zhu, Banghua and Li, Zhuohan and Zhuo, Danyang and Gonzalez, Joseph E. and Stoica, Ion. Fairness in Serving Large Language Models. In *Proceedings of the 18th USENIX Conference on Operating Systems Design and Implementation, OSDI'24*, 2024.
- [45] Gao Shiwei, Chen Youmin, and Shu Jiwu. Fast State Restoration in LLM Serving with HCache. In *Proceedings of the Twentieth European Conference on Computer Systems, EuroSys '25*, pages 128–143, 2025.
- [46] Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. Megatron-LM: Training Multi-Billion Parameter Language Models Using Model Parallelism. *arXiv preprint arXiv:1909.08053*, 2019.
- [47] Jiaming Tang, Yilong Zhao, Kan Zhu, Guangxuan Xiao, Baris Kasicki, and Song Han. QUEST: Query-Aware Sparsity for Efficient Long-Context LLM Inference. In *International Conference on Machine Learning*, pages 47901–47911. PMLR, 2024.
- [48] Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*, 2024.
- [49] Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. Efficient Streaming Language Models with Attention Sinks. In *The Twelfth International Conference on Learning Representations*, 2024.
- [50] Yi Xiong, Hao Wu, Changxu Shao, Ziqing Wang, Rui Zhang, Yuhong Guo, Junping Zhao, Ke Zhang, and Zhenxuan Pan. LayerKV: Optimizing Large Language Model Serving with Layer-wise KV Cache Management. *arXiv preprint arXiv:2410.00428*, 2025.
- [51] Sheng Ying, Zheng Lianmin, Yuan Binhang, Li Zhuohan, Ryabinin Max, Chen Beidi, Liang Percy, Christopher Ré, Stoica Ion, and Zhang Ce. FlexGen: High-throughput Generative Inference of Large Language Models with a Single GPU. In *Proceedings of the 40th International Conference on Machine Learning, ICML'23*, 2023.
- [52] 01.AI Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Heng Li, Jiangcheng Zhu, Jianqun Chen, Jing Chang, Kaidong Yu, Peng Liu, Qiang Liu, Shawn Yue, Senbin Yang, Shiming Yang, Tao Yu, Wen Xie, Wenhao Huang, Xiaohui Hu, Xiaoyi Ren, Xinyao Niu, Pengcheng Nie, Yuchi Xu, Yudong Liu, Yue Wang, Yuxuan Cai, Zhenyu Gu, Zhiyuan Liu, and Zonghong Dai. Yi: Open Foundation Models by 01.AI. *arXiv preprint arXiv:2403.04652*, 2025.
- [53] Tang Yuqing, Tran Chau, Li Xian, Chen Peng-Jen, Goyal Naman, Chaudhary Vishrav, Gu Jiatao, and Fan Angela. Multilingual Translation from Denoising Pre-Training. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3450–3466, 2021.
- [54] Zhenyu Zhang, Ying Sheng, Tianyi Zhou, Tianlong Chen, Lianmin Zheng, Ruisi Cai, Zhao Song, Yuandong Tian, Christopher Re, Clark Barrett, Zhangyang Wang, and Beidi Chen. H2O: Heavy-Hitter Oracle for Efficient Generative Inference of Large Language Models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [55] Xuanlei Zhao, Bin Jia, Haotian Zhou, Ziming Liu, Shenggan Cheng, and Yang You. HeteGen: Efficient Heterogeneous Parallel Inference for Large Language Models on Resource-Constrained Devices. In *MLSys*, 2024.
- [56] Yinmin Zhong, Shengyu Liu, Junda Chen, Jianbo Hu, Yibo Zhu, Xuanzhe Liu, Xin Jin, and Hao Zhang. DistServe: Disaggregating Prefill and Decoding for Goodput-optimized Large Language Model Serving. In *OSDI*, pages 193–210, 2024.