

PRKV: Page Restructured KV Cache for High Accuracy and Efficiency LLM Generation

Fang Wu¹, Congming Gao^{1,†}, Weixi Zhu², Jiwu Shu³

¹Xiamen University, ²Huawei Technologies, ³Tsinghua University

Abstract

As the key-value(KV) cache size scales with context length, the substantial GPU memory demand and the overhead of accessing a large KV cache in each decoding step pose challenges for deploying LLMs with long contexts. Among various sparse attention methods, offloading-based KV retrieval method preserves entire KV cache in CPU memory and dynamically retrieves most relevant KV pairs for each decoding step, which performs higher quality and effectively reduces GPU memory footprint than other methods. However, existing offloading-based KV retrieval methods perform page-level to reduce estimation overhead, which introduces inaccurate KV selection and significant retrieval overhead. We propose PRKV, a framework that both-optimizes algorithm and system for page-level KV retrieval with KV offloading. On the algorithm side, PRKV introduces hybrid KV selection that combines static and dynamic KV selection strategies. On the system side, PRKV employs contiguous memory indexing and batched transfer optimizations to improve retrieval efficiency. Experiments demonstrate that PRKV **improves accuracy** across various scenarios and models, delivering up to **6.75×** speedup compared to SOTA KV retrieval methods.

1 Introduction

Recent advancements in large language models (LLMs) have significantly enhanced their ability to process long contexts(Wang et al., 2024; Jiang et al., 2024). However, the efficient serving of these long-context LLMs is hampered by two critical challenges related to the key-value (KV) cache, whose size scales with sequence length and inference batch size. The resulting memory-bound access not only slows down decoding but also imposes substantial GPU memory demands, often exceeding hardware limits.

To address these challenges, numerous sparse attention methods have been proposed. KV eviction reduces memory usage and access latency by discarding less important tokens (Xiao et al., 2024b; Zhang et al., 2023), but its inability to recover evicted information can degrade performance on tasks requiring long-range dependencies (Gao et al., 2025; Liu et al., 2025d). As an alternative, KV retrieval preserves the full cache in CPU memory and dynamically fetches relevant KV pairs during decoding. While this approach improves accuracy and alleviates GPU memory pressure, it introduces new challenges in efficient and accurate token selection.

Inaccurate KV selection. KV retrieval requires estimating the importance of each KV pair at every decoding step, token-level approach such as Mag-icpig and Infinigen (Chen et al., 2025b; Lee et al., 2024)is computationally expensive estimation, as illustrated in Figure 2. To mitigate this overhead, existing methods, such as Quest and Arkvale (Tang et al., 2024; Chen et al., 2024), perform retrieval at the page level. However, this coarser granularity results in the selection of less accurate KV pairs compared to the ground-truth top-k selection, as illustrated in Figure 9.

Significant retrieval overhead. Limited by the transfer efficiency between CPU memory and GPU memory, retrieval introduces significant latency when fetching sparse KV pairs back to the GPU during decoding.

In this paper, we aim to solve above two concerns and build an inference system for long-context LLM with high accuracy and low latency.

Reusing top-k KV pairs. Performing high-precision top-k KV selection at every decoding step is computationally expensive. However, we observe strong similarity among query vectors and high overlap in the top-k KV pairs of consecutive tokens, which enables the reuse of previously selected top-k KV pairs. This reuse strategy main-

[†]Corresponding author

tains high attention recall while incurring minimal additional overhead.

Hybrid KV selection. Although reusing previous top-k KV pairs improves selection accuracy, static selection methods fail to account for the dynamic nature of KV importance (Chen et al., 2024). To address this limitation, we propose a novel hybrid KV selection approach that combines both static and dynamic retrieval strategies during decoding. As illustrated in Figure 3b, our hybrid method achieves a higher recall rate than alternative selection strategies.

New retrieval pattern with low overhead. To reduce latency when retrieving KV pairs from CPU memory, we propose a system-level optimization that addresses two key inefficiencies. First, we note that existing head-wise retrieval methods often overlook the critical impact of the KV cache layout on access efficiency. Our analysis shows that the HND layout enables contiguous memory access for head-wise operations, making it ideal for this context. Second, we implement a batched KV transfer mechanism to maximize PCIe bandwidth utilization during retrieval. Together, these optimizations significantly reduce latency.

Based on the preceding analysis, we introduce PRKV, a framework that both-optimizes algorithm and system design for page-level KV retrieval with KV offloading. As illustrated in Figure 4, the algorithmic component of PRKV features a hybrid KV selection method for decoding steps. On the system side, the framework employs contiguous memory indexing and a batched transfer mechanism to maximize retrieval efficiency.

We conducted extensive experiments and ablation studies to evaluate PRKV. In Section 4.1, we assess various LLMs on benchmarks including LongBench V2 (Bai et al., 2024), Needle In A Haystack (Kamradt, 2023), and the reasoning benchmarks AIME (Jia, 2025) and MATH500 (Hendrycks et al., 2021). The results show that PRKV achieves the **highest accuracy** among other methods and even exceeds the accuracy of the full cache. In Section 4.2, the end-to-end results demonstrate that PRKV achieves a speedup of up to $6.75\times$ over SOTA KV retrieval methods.

2 Observation

2.1 Recall Rate Definition

Attention mechanisms are core of transformer-based LLMs. For each attention head i , the general

attention computation formula is as follows:

$$W_i = \text{softmax}\left(\frac{Q_i K_i^T}{\sqrt{d}}\right), \quad O_i = W_i V_i \quad (1)$$

Here, W_i is the attention scores across tokens, and O_i is the attention output from the i -th attention head. Sparse attention algorithms aim to retain the most important tokens and achieve the highest attention scores with a fixed budget k through different estimation methods, which can be presented as:

$$I_i = \text{argtopk}(f(Q_i, K_i), k) \quad (2)$$

where I_i denotes the selected subset tokens' indices in i -th attention head, and the function f is an efficient estimation method to identify the importance of the tokens. The **Attention Recall Rate** is defined as the proportion of the ground truth attention scores that the selected subset covers:


$$\mathcal{R}_i = \sum_{j \in I_i} W_{ij} \quad (3)$$

For a given KV budget k , ground truth top-k means select k KV pairs with highest ground truth attention scores, which guarantees the highest recall rate in return. A high attention recall rate demonstrates that the selected tokens encompass the most significant attention scores, which is essential to maintain high accuracy.


2.2 Long estimation time of Token-Level methods

In long-context scenarios, KV cache is often offloaded to CPU memory due to limited GPU capacity. Token-level retrieval methods, such as MagicPIG (Chen et al., 2025b), estimate token importance by scanning the full KV cache, resulting in substantial computational overhead. In contrast, group-level approaches like Arkvale (Chen et al., 2024) leverage page-level metadata stored in GPU memory to perform coarse-grained importance estimation more efficiently.

To quantify this efficiency gap, we measure the estimation overhead under a 64K context length. As shown in Figure 2, MagicPIG introduces significantly higher estimation latency, reaching up to $15\times$ the cost of full-cache attention implemented with FlashAttention2 (Dao, 2023), whereas Arkvale incurs considerably lower overhead. These results highlight the scalability limitations of token-level retrieval and motivate the need for more efficient KV selection strategies.



(a) Attention recall rate of different approaches using a 1K token budget on a 37K NIAH task.



(b) Attention recall rate of different approaches using a 4K token budget on an AIME task.

Figure 1: Analysis of attention recall rate between Arkvale and ground-truth Top- k on different tasks. Figure 1a shows attention recall rates of layer 10, layer 20, and layer 30 on a NIAH task with a 1K token budget and 37K input length using Llama3.1-8B-Instruct. Figure 1b shows the average attention recall rate across all layers on an AIME problem with a 4K token budget and generation length up to 32K tokens using Qwen3-8B.




Figure 2: Estimation time of MagicPIG and Arkvale with 64K sequence length

2.3 Limitation of page-level selection

Page-level KV selection reduces estimation overhead by grouping consecutive tokens into coarse-grained pages. However, such approaches overlook the tokens within the same page can exhibit vastly different importance though they have high semantic similarities. As shown in Figure 9, important tokens selected by ground truth topk are distributed throughout the sequence, while Arkvale’s page grained nature fails to select these important individual tokens, directly resulting in a **lower accuracy**.

Empirical analysis further reveals the Attention Recall Rate between page-level selection and ground-truth top- k across long-context tasks. As shown in Figure 1a, Arkvale exhibits an average recall rate about **0.25** in NIAH and around **0.8** in AIME2024. Additionally, such discrepancies can introduce accumulated bias during sparse decoding, ultimately degrading model performance.

2.4 Opportunity: Locality-Aware Top-k Reuse

Prior studies show that query vectors of consecutive tokens exhibit strong temporal similarity (Su et al., 2024; Yuan et al., 2025). Leveraging this property, methods such as FreeKV (Liu et al., 2025b) reuse top- k KV pairs from previous decoding steps. Empirical analysis (Figure 3a) shows that reusing only the previous top- k tokens yields an attention recall rate of **0.65** under a KV budget of 1024. However, augmenting the reused set with a small local window (e.g., 64 tokens) increases the recall rate to a level comparable to ground-truth top- k selection. This observation indicates that combining top- k reuse with local context can substantially improve selection accuracy.

2.5 Opportunity: Hybrid sparse token selection

While top- k reuse provides fine-grained token selection, page-level retrieval offers dynamic KV selection. We explore a hybrid strategy that integrates both without token eviction. Given a KV budget k , we partition it into a static **token set** and a dynamic **page set**. The static set occupies a ratio S of the budget and consists of the top kS KV pairs reused from the previous step, while the remaining $1 - S$ budget is allocated to dynamically selected pages based on their importance during each decoding step. We evaluate the attention recall rate of this hybrid design on LongBench tasks (Figure 3b), starting selection from relative position 0. We compare against fully dynamic page-level




Figure 3: Average attention recall rate on sampled LongBench tasks with 32K input length. (a) “reuse w/o local KV pairs” reuses only top- k ($k = 1024$) KV pairs starting from relative position 0; “reuse w/ local KV pairs” additionally includes the last 64 local KV pairs. (b) Attention recall rate under different $\mathcal{S}$ settings with a 1K budget.

selection ($\mathcal{S} = 0$, representing Arkvale) and fully static selection ($\mathcal{S} = 1$, representing SnapKV). For hybrid selection, we take $\mathcal{S} = 0.25$ as an example. Results show that hybrid selection consistently achieves higher recall than both extremes, supporting the effectiveness of combining static token reuse with dynamic page retrieval.

3 PRKV

Based on the above analysis and observations, we propose Page Restructured KV Cache (PRKV), a method combines static and dynamic KV pair retrieval in page granularity, achieving higher Attention Recall Rate and lower retrieval overhead.

3.1 Hybrid token selection

Static token-level selection. The detailed mechanism of static token-level selection algorithm is formalized in Algorithm 1 lines 8-19. As discussed in Section 2.4, the local tokens are consistently important in sparse attention computation, so PRKV keeps the local tokens at the end of context. PRKV uses an observation window at the end of sequence (i.e., the local window). It then computes the attention scores between the queries in the observation window and the keys of all preceding tokens. Finally, PRKV selects the top- k tokens with the highest attention scores as the static tokens. After identifying static tokens, PRKV relocates their corresponding KV to the beginning of KV cache to make them contiguous, while preserving the relative order of the remaining tokens and still organizing KV cache to pages. PRKV recomputes and updates page representations since

the KV cache pages have been restructured (line 23-26).

To maintain compatibility with Grouped-Query Attention (GQA) and achieve more precise token selection, PRKV modifies the process of computing observation-based attention scores. PRKV first repeats the group’s Key cache for each query head. Then computes the attention score for each head individually and averages these scores within the group (line 8-10).

Principles of choosing $\mathcal{S}$. Reuse of previous topk tokens provides a fine-grained selection mechanism, the choice of static token ratio $\mathcal{S}$ should strike a balance that maximizes the likelihood of overlap between important tokens. Based on this intuition, we design an online algorithm for determining $\mathcal{S}$.

In the prefill phase, after computing the observation-window attention scores S (line 10), Let $S^{(1)}$ and $S^{(T)}$ denote the attention score vectors of the first and last query tokens in the prefilling phase observe window, respectively. For each head h , we compute the overlap of their topk key indices using

$$C_h = \left| \text{topk}\left(S_h^{(1)}, k\right) \cap \text{topk}\left(S_h^{(T)}, k\right) \right| \quad (4)$$

The static ratio is then obtained by averaging the overlap across all heads, P denotes the page size. Then the static ratio is used in the subsequent selection process.

$$\mathcal{S} = \frac{P \cdot \left\lfloor \frac{E_h[C_h]}{P} \right\rfloor}{k} \quad (5)$$

Table 1 shows the static token ratio of different models, input lengths and KV budgets.




Figure 4: Algorithm design overview of PRKV.

Algorithm 1 Static Token Selection

```

1: Input: queries of observation window  $Q_{obs}$ ,
    $K$  and  $V$ , token budget  $k$ , static ratio  $\mathcal{S}$ , de-
   code step  $ds$ , update interval  $T$ 
2: Constants: query group size  $gs$ , local window
   size  $nl$ 
3: Output:
4: if  $ds \% T \neq 0$  then
5:   return  $K, V, None$ 
6: end if
7: # compute token importance with observation
   window
8:  $K'_{pre} \leftarrow \text{repeat\_kv}(K[: -nl], gs)$ 
9:  $S \leftarrow \text{compute\_score}(Q_{obs}, K'_{pre})$ 
10:  $S \leftarrow \text{Average}(S)$ 
11: if  $S$  is None then
12:   # estimate static ratio during prefill
13:    $S^{(1)}, S^{(T)} \leftarrow \text{FirstLast}(S)$ 
14:    $C_h \leftarrow \text{TopKOverlap}(S_h^{(1)}, S_h^{(T)}, k)$ 
15:    $C_{avg} \leftarrow \text{Average}(C_h)$ 
16:    $\mathcal{S} \leftarrow \frac{P \cdot \lfloor C_{avg}/P \rfloor}{k}$ 
17: end if
18: # select static important tokens
19:  $I_{\text{topk}} \leftarrow \text{TopKIndices}(S_{\text{mean}}, k\mathcal{S})$ 
20: # concat static tokens
21:  $I_{\text{rem}} \leftarrow \text{AllIndices} \setminus I_{\text{topk}}$ 
22:  $P \leftarrow \text{Concat}(I_{\text{topk}}, I_{\text{rem}})$ 
23:  $K[: -nl] \leftarrow \text{gather}(K[: -nl], P)$ 
24:  $V[: -nl] \leftarrow \text{gather}(V[: -nl], P)$ 
25: # update page representation
26:  $\text{Rep}_{\text{page}} \leftarrow \text{update\_page\_representation}(K[: -nl])$ 
27: # asynchronous rewrite KV
28:  $KV_{\text{CPU}} \leftarrow \text{Copy}(K[: k], V[: k])$ 
29: return  $K, V, \text{Rep}_{\text{page}}$ 

```

KV Budget	Qwen2.5-14B-Instruct			Llama3.1-8B-Instruct		
	32K	64K	128K	32K	64K	128K
$k=512$	0.25	0.25	0.25	0.625	0.5625	0.5
$k=1024$	0.4062	0.375	0.3438	0.6875	0.625	0.5625
$k=2048$	0.5	0.4844	0.4219	0.7344	0.7031	0.625
$k=4096$	0.6016	0.5781	0.5	0.7891	0.75	0.6875

Table 1: $\mathcal{S}$ across input lengths and KV budgets

Hybrid Sparse Attention. PRKV offloads all KV cache to CPU, keeping page representations and KV space with budget k (named K', V') in GPU. Page representations are min and max vectors of keys in a page, allowing to approximate the maximum token importance in a page. During each decoding step, KV of static and local tokens are always in GPU. For a KV budget k , PRKV dynamically chosen a set of $\frac{k\mathcal{S}}{\text{pagesize}}$ pages and retrieval them (Section 3.2). These selected tokens' KV then participate in the sparse paged-attention computation. The detail is shown in Algorithm 2.

Algorithm 2 Hybrid Sparse Attention

```

1: Input:  $Q, K'$  and  $V'$ , Page Representation
    $P$ , token budget  $k$ , static ratio  $\mathcal{S}$ , In GPU KV
   pages  $GP$ 
2: Constants: local window size  $nl$ 
3: Output: attention out  $O$ 
4: # estimate page importance
5:  $EScores \leftarrow QP$ 
6:  $IPages \leftarrow \text{TopKIndices}(EScores, k(1-\mathcal{S}))$ 
7: # KV pages not in GPU
8:  $Rpage \leftarrow IPages \setminus GP$ 
9: # Retrieval pages not in GPU
10:  $K', V' \leftarrow \text{Retrieval}(RPages)$ 
11:  $O \leftarrow \text{PagedAttention}(Q, K', V')$ 

```

Periodic static token update. The importance of initially selected static tokens can decay over long decoding steps. To mitigate this, we introduce a periodic update strategy for the static token set, formalized in Algorithm 1. Given an interval T , when $ds \bmod T = 0$, we recompute the static tokens using the local window, update the page representation, and asynchronously write the updated set back to CPU memory Figure 23.

3.2 Efficient retrieval Optimization

Retrieving KV pages from CPU introduces two key overheads: inefficient indexing and data transfer. Per-head KV pages are stored non-contiguously in CPU memory, leading to costly indexing, while transferring needed pages requires multiple kernel launches and serial transfers. To mitigate these bottlenecks, we design CUDA kernels for contiguous memory indexing and batched KV transfer.

Contiguous Memory Indexing. When managing KV pages in CPU memory, two common layouts are NHD and HND, with tensor shapes of $(2, \text{page_size}, \text{num_head}, \text{head_dim})$ and $(\text{num_head}, 2, \text{page_size}, \text{head_dim})$, respectively. These layouts differ significantly in indexing efficiency for page retrieval, which is performed per head. Retrieving a page for head i under the NHD layout requires indexing all KV pairs within a page to copy those for head i . In contrast, the HND layout allows direct indexing of all KV pairs for a specific head without search overhead. Consequently, PRKV adopts the HND layout for both CPU and GPU KV caches, utilizing contiguous memory indexing. The efficiency of this design is demonstrated in Section 4.3.

Batched KV Transfer. As shown on the left of Figure 5, the naive page-to-page transfer mode moves KV data directly from CPU memory to GPU memory. This approach requires multiple kernel launches, which is inefficient. To address this, we introduce Batched KV Transfer (right side of Figure. 5). Our method first consolidates the required KV pages into a pinned CPU buffer, then performs a single, bulk transfer to GPU memory. Finally, the data is dispatched into the GPU’s KV pool.




Figure 5: KV transfer optimization

4 Experiments

4.1 Accuracy Evaluation

Setup. We evaluate PRKV across various models and tasks. We conduct evaluation on LongBench V2 (Bai et al., 2024), Needle In A Haystack (Kamradt, 2023), and long reasoning tasks, including MATH500 (Hendrycks et al., 2021) and AIME24 (Jia, 2025). For LongBench V2 and NIAH, we use general models Llama-3.1-8B-Instruct (Meta Llama, 2024) and Qwen-2.5-14B-Instruct (Qwen et al., 2025). For reasoning tasks, we use Qwen3-8B and Qwen3-4B (Yang et al., 2025).

Baselines We compare PRKV against four dynamic sparse attention methods, including Quest (Tang et al., 2024), Arkvale, ShadowKV (Sun et al., 2024) and SnapKV. All methods perform dynamic sparse attention during decoding. For Quest, Arkvale and PRKV, we set their page size to 32 and local window size to 64. For PRKV the static set ratio $\mathcal{S} = 0.25$ and update interval $T = 128$.

LongBench V2. LongBench V2 covers various difficulty levels and context lengths, ranging from 8K to 2M tokens. The average scores is shown in Figure 6a and Figure 6b. PRKV consistently achieves the highest score. While other methods experience serious performance degradation under a low KV budget, PRKV is more robust and even outperforms original full attention when tested on Qwen-2.5-14B-Instruct. The detail results of all kind of tasks are shown in Appendix B.1.

Needle In A Haystack. We evaluate retrieval accuracy on the NIAH, where a single critical sentence (“needle”) is inserted into a long irrelevant context. The average results are shown in Figure 6c and Figure 6d. On Llama3.1-8B-Instruct, PRKV and SnapKV achieve the best performance, closely matching the full cache. On Qwen-2.5-14B-Instruct, PRKV consistently outperforms the full cache. Detailed heatmaps in Appendix B.2.

Reasoning Tasks. We evaluate reasoning performance on MATH500 and AIME24 using Qwen3-8B and Qwen3-4B. The maximum generation length is set to 32K with temperature 0.6 and top- p 0.95. Following prior work (DeepSeek-AI et al., 2025), we report pass@ k and avg@ k (Chen et al., 2021), with $k = 4$. All KV retrieval methods use a KV budget of 1024. Since SnapKV is originally designed for prefilling, we adapt it to decoding using same static token update interval as our method.

As shown in Table 2, PRKV achieves accuracy comparable to the full KV cache and consistently




Figure 6: Average score of PRKV and baselines. (a) LongBench V2 on Llama3.1-8B-Instruct. (b) LongBench V2 on Qwen-2.5-14B-Instruct. (c) NIAH on Llama3.1-8B-Instruct. (d) NIAH on Qwen-2.5-14B-Instruct.

Methods	AIME24		MATH500	
	<i>pass@k</i>	<i>avg@k</i>	<i>pass@k</i>	<i>avg@k</i>
<i>Qwen3-4B</i>	86.67	74.17	95.40	94.20
Quest	70.00	56.67	84.80	78.50
Arkvale	76.67	62.50	85.40	84.30
ShadowKV	76.67	59.17	85.00	81.10
SnapKV	83.33	67.50	93.60	90.00
PRKV	86.67	75.00	95.00	93.80
<i>Qwen3-8B</i>	83.33	74.17	96.20	95.10
Quest	76.67	65.83	86.80	78.60
Arkvale	80.00	68.33	87.00	86.90
ShadowKV	76.67	67.50	86.80	82.60
SnapKV	83.33	67.50	93.60	93.60
PRKV	83.33	73.33	95.40	94.40

Table 2: Accuracy results of long reasoning tasks.

outperforms other compression methods across all datasets, demonstrating the effectiveness of our hybrid KV selection strategy.

4.2 Efficiency Evaluation

The efficiency experiments were conducted on an Nvidia A100 80GB GPU connected with Intel Xeon Silver 4310 CPU via PCIe Gen4. We compare PRKV with offloading methods Arkvale and ShadowKV under long-input and long-generation scenarios. We set $k = 1024$. The end-to-end latency is in the left of Figure 7. PRKV achieves a multi-fold speedup over Arkvale and ShadowKV across multiple models, demonstrating significant efficiency and achieving maximum speedup of up to $6.75\times$ compared to Arkvale.

4.3 Ablation Results

System optimization. Since PRKV’s system optimizations including contiguous memory indexing

(CI) and batched KV transfer (BT) are designed to reduce recall latency, we evaluate the Time Per Output Token (TPOT) using Llama3.1-8B-Instruct and Qwen2.5-14B-Instruct with 512 generated tokens. As shown right of Figure 7, contiguous memory indexing achieve $1.5\times$ reduction in TPOT. Batched KV transfer contributes most to the improvements, achieving up to a $7.7\times$ TPOT reduction.

Attention breakdown. Figure 8 presents Attention breakdown of Full Attention and PRKV across sequence lengths from 32K to 128K with batchsize=4. While PRKV introduces minor computational overheads from criticality estimation and top-K filtering, its sparse attention significantly reduces latency. When the sequence length reaches 128K, Full Attention requires over 500 ms to retrieve the KV cache from CPU, whereas PRKV takes only 1 ms to retrieve important KV pages.

Update interval. We study the impact of the static token update interval T on the attention recall rate. As shown in Table 3, smaller update intervals consistently yield higher recall across models and input lengths, indicating that more frequent updates better track the evolution of token importance during decoding. However, performance degrades precipitously when $T > 128$, especially under long-context settings. Therefore, we set $T = 128$ as our default configuration to achieve an optimal balance between attention accuracy and update overhead.

T	Qwen2.5-14B-Instruct			Llama3.1-8B-Instruct		
	32K	64K	128K	32K	64K	128K
$T = 32$	0.947	0.933	0.919	0.956	0.935	0.913
$T = 64$	0.939	0.927	0.911	0.943	0.929	0.910
$T = 128$	0.927	0.915	0.902	0.934	0.916	0.901
$T = 256$	0.819	0.773	0.721	0.804	0.758	0.6875

Table 3: Update interval T ablation. Attention Recall Rate of different T across input lengths and models




Figure 7: **Left:** End-to-end latency under long-input and long-generation scenarios. **Right:** TPOT of different system optimization designs.




Figure 8: PRKV self-attention time breakdown compared to FlashInfer Full Attention.

5 Related Work

The KV cache enables efficient autoregressive inference but incurs substantial memory overhead as context length grows. Prior work mitigates this issue through three main directions: quantization, eviction, and retrieval-based sparse attention.

KV Cache Quantization. Quantization methods reduce KV memory footprint by approximating high-precision values with low-bit representations (Liu et al., 2024; Zhang et al., 2024; Hooper et al., 2025). For example, KIVI (Liu et al., 2024) adopts asymmetric quantization, applying channel-wise quantization for keys and token-wise quantization for values. These approaches lower storage cost and are orthogonal to sparse attention methods.

KV Cache Eviction. Eviction-based approaches maintain a fixed-size KV cache by retaining important tokens and discarding others. StreamingLLM (Xiao et al., 2024b) preserves sink tokens and a lo-

cal window, while SnapKV (Li et al., 2024) identifies important tokens using an observation window at the prompt end. H2O (Zhang et al., 2023) estimates token importance from accumulated historical attention scores. However, premature eviction may degrade accuracy, as tokens initially deemed unimportant can become relevant later in decoding.

Retrieval-based Sparse Attention. Retrieval methods retain the full KV cache in CPU memory and dynamically select important KV pairs for attention computation. *Token-level retrieval* methods (Zhang et al., 2025, 2026) estimate token importance directly. RetrievalAttention (Liu et al., 2025a) and MagicPIG (Chen et al., 2025b) employ approximate nearest neighbor search and locality-sensitive hashing on CPU, respectively. InfiniGen (Lee et al., 2024) computes partial query-key interactions on GPU to guide retrieval. While precise, their importance estimation cost scales linearly with context length. *Group-level retrieval* methods operate on token groups with precomputed metadata. Quest, Arkvale, and InfLLM (Tang et al., 2024; Chen et al., 2024; Xiao et al., 2024a) partition sequences into fixed-size pages, whereas ClusterKV (Liu et al., 2025c) and RetroInfer (Chen et al., 2025a) form semantic clusters via k-means.

6 Conclusion

We propose PRKV, a hybrid sparse attention framework for efficient page-level KV retrieval with offloading. By jointly improving selection accuracy and retrieval efficiency, PRKV enables scalable long-context inference with consistent accuracy gains and substantial decoding speedups.

7 Limitations

First, the method introduces additional hyperparameters (e.g., update interval), which may require re-tuning across models, context lengths, and hardware platforms, and the best configuration may not transfer reliably. Second, our evaluation primarily focuses on the decode phase. Finally, our experiments are conducted on a specific GPU and software stack; performance and memory behavior may vary across other accelerators and deployment environments.

References

- Yushi Bai, Shangqing Tu, Jiajie Zhang, Hao Peng, Xiaozhi Wang, Xin Lv, Shulin Cao, Jiazheng Xu, Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. 2024. Longbench v2: Towards deeper understanding and reasoning on realistic long-context multitasks. *arXiv preprint arXiv:2412.15204*.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, and 39 others. 2021. [Evaluating large language models trained on code](#). *Preprint*, arXiv:2107.03374.
- Renze Chen, Zhuofeng Wang, BeiQuan Cao, Tong Wu, Size Zheng, Xiuhong Li, Xuechao Wei, Shengen Yan, Meng Li, and Yun Liang. 2024. [Arkval: Efficient generative LLM inference with recallable key-value eviction](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Yaoqi Chen, Jinkai Zhang, Baotong Lu, Qianxi Zhang, Chengruidong Zhang, Jingjia Luo, Di Liu, Huiqiang Jiang, Qi Chen, Jing Liu, Bailu Ding, Xiao Yan, Jiawei Jiang, Chen Chen, Mingxing Zhang, Yuqing Yang, Fan Yang, and Mao Yang. 2025a. [Retroinfer: A vector-storage approach for scalable long-context llm inference](#). *Preprint*, arXiv:2505.02922.
- Zhuoming Chen, Ranajoy Sadhukhan, Zihao Ye, Yang Zhou, Jianyu Zhang, Niklas Nolte, Yuandong Tian, Matthijs Douze, Leon Bottou, Zhihao Jia, and Beidi Chen. 2025b. [MagicPIG: LSH sampling for efficient LLM generation](#). In *The Thirteenth International Conference on Learning Representations*.
- Tri Dao. 2023. [Flashattention-2: Faster attention with better parallelism and work partitioning](#). *Preprint*, arXiv:2307.08691.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, and 181 others. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *Preprint*, arXiv:2501.12948.
- Wei Gao, Xinyu Zhou, Peng Sun, Tianwei Zhang, and Yonggang Wen. 2025. [Rethinking key-value cache compression techniques for large language model serving](#). *Preprint*, arXiv:2503.24000.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. [Measuring mathematical problem solving with the math dataset](#). *Preprint*, arXiv:2103.03874.
- Coleman Hooper, Sehoon Kim, Hiva Mohammadzadeh, Michael W. Mahoney, Yakun Sophia Shao, Kurt Keutzer, and Amir Gholami. 2025. [Kvquant: Towards 10 million context length llm inference with kv cache quantization](#). *Preprint*, arXiv:2401.18079.
- Maxwell Jia. 2025. [Maxwell-jia/aime_2024](#). https://huggingface.co/datasets/Maxwell-Jia/AIME_2024. Accessed: 2025-09-16.
- Juyong Jiang, Fan Wang, Jiasi Shen, Sungju Kim, and Sunghun Kim. 2024. [A survey on large language models for code generation](#). *Preprint*, arXiv:2406.00515.
- Greg Kamradt. 2023. Needle in a Haystack: Pressure Testing LLMs. https://github.com/gkamradt/LLMTest_NeedleInAHaystack. Accessed: 2025-09-16.
- Wonbeom Lee, Jungi Lee, Junghwan Seo, and Jaewoong Sim. 2024. Infinigen: efficient generative inference of large language models with dynamic kv cache management. In *Proceedings of the 18th USENIX Conference on Operating Systems Design and Implementation*, OSDI'24. USENIX Association.
- Yuhong Li, Yingbing Huang, Bowen Yang, Bharat Venkitesh, Acyr Locatelli, Hanchen Ye, Tianle Cai, Patrick Lewis, and Deming Chen. 2024. [SnapKV: LLM knows what you are looking for before generation](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Di Liu, Meng Chen, Baotong Lu, Huiqiang Jiang, Zhenhua Han, Qianxi Zhang, Qi Chen, Chengruidong Zhang, Bailu Ding, Kai Zhang, Chen Chen, Fan Yang, Yuqing Yang, and Lili Qiu. 2025a. [Retrievalattention: Accelerating long-context LLM inference via vector retrieval](#).
- Guangda Liu, Chengwei Li, Zhenyu Ning, Minyi Guo, and Jieru Zhao. 2025b. [Freekv: Boosting kv cache retrieval for efficient llm inference](#). *Preprint*, arXiv:2505.13109.
- Guangda Liu, Chengwei Li, Jieru Zhao, Chenqi Zhang, and Minyi Guo. 2025c. [Clusterkv: Manipulating llm kv cache in semantic space for recallable compression](#). *Preprint*, arXiv:2412.03213.

- Xiang Liu, Zhenheng Tang, Hong Chen, Peijie Dong, Zeyu Li, Xiuze Zhou, Bo Li, Xuming Hu, and Xiaowen Chu. 2025d. [Can llms maintain fundamental abilities under kv cache compression?](#) *Preprint*, arXiv:2502.01941.
- Zirui Liu, Jiayi Yuan, Hongye Jin, Shaochen Zhong, Zhaozhuo Xu, Vladimir Braverman, Beidi Chen, and Xia Hu. 2024. Kivi: A tuning-free asymmetric 2bit quantization for kv cache. *arXiv preprint arXiv:2402.02750*.
- Meta Llama. 2024. [Llama-3.1-8B-Instruct](#).
- Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, and 25 others. 2025. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115.
- Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. 2024. [Roformer: Enhanced transformer with rotary position embedding](#). *Neurocomput.*, 568(C).
- Hanshi Sun, Li-Wen Chang, Wenlei Bao, Size Zheng, Ningxin Zheng, Xin Liu, Harry Dong, Yuejie Chi, and Beidi Chen. 2024. Shadowkv: Kv cache in shadows for high-throughput long-context llm inference. *arXiv preprint arXiv:2410.21465*.
- Jiaming Tang, Yilong Zhao, Kan Zhu, Guangxuan Xiao, Baris Kasikci, and Song Han. 2024. Quest: query-aware sparsity for efficient long-context llm inference. In *Proceedings of the 41st International Conference on Machine Learning, ICML'24*. JMLR.org.
- Minzheng Wang, Longze Chen, Cheng Fu, Shengyi Liao, Xinghua Zhang, Bingli Wu, Haiyang Yu, Nan Xu, Lei Zhang, Run Luo, Yunshui Li, Min Yang, Fei Huang, and Yongbin Li. 2024. [Leave no document behind: Benchmarking long-context llms with extended multi-doc qa](#). *Preprint*, arXiv:2406.17419.
- Chaojun Xiao, Pengle Zhang, Xu Han, Guangxuan Xiao, Yankai Lin, Zhengyan Zhang, Zhiyuan Liu, and Maosong Sun. 2024a. [InfLLM: Training-free long-context extrapolation for LLMs with an efficient context memory](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. 2024b. [Efficient streaming language models with attention sinks](#). In *The Twelfth International Conference on Learning Representations*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- Jingyang Yuan, Huazuo Gao, Damai Dai, Junyu Luo, Liang Zhao, Zhengyan Zhang, Zhenda Xie, Y. X. Wei, Lean Wang, Zhiping Xiao, Yuqing Wang, Chong Ruan, Ming Zhang, Wenfeng Liang, and Wangding Zeng. 2025. [Native sparse attention: Hardware-aligned and natively trainable sparse attention](#). *Preprint*, arXiv:2502.11089.
- Yu Zhang, Dong Guo, Fang Wu, Guoliang Zhu, Dian Ding, and Yiming Zhang. 2025. [Anchorattention: Difference-aware sparse attention with stripe granularity](#). *Preprint*, arXiv:2505.23520.
- Yu Zhang, Songwei Liu, Chenqian Yan, Sheng Lin, Beichen Ning, Fangmin Chen, and Xing Wang. 2026. [S2o: Early stopping for sparse attention via online permutation](#). *Preprint*, arXiv:2602.22575.
- Zhenyu Zhang, Shiwei Liu, Runjin Chen, Bhavya Kailkhura, Beidi Chen, and Atlas Wang. 2024. [Q-hitter: A better token oracle for efficient llm inference via sparse-quantized kv cache](#). In *MLSys*.
- Zhenyu Zhang, Ying Sheng, Tianyi Zhou, Tianlong Chen, Lianmin Zheng, Ruisi Cai, Zhao Song, Yuandong Tian, Christopher Ré, Clark Barrett, Zhangyang Wang, and Beidi Chen. 2023. H2o: heavy-hitter oracle for efficient generative inference of large language models. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23*, Red Hook, NY, USA. Curran Associates Inc.

A Picture of Observation

Figure 9 shows the important token selection distribution of ground truth topk and Arkvale on a 37K NIAH task, using a 1K token KV budget. The result is from layer 5 of the 14-th generation token.

B Detail results in Evaluation

In this section, we provide detail results in evaluation.

B.1 LongBench V2

We report accuracy under the context length categories of short, medium and long, which has easy and hard two level problems, as well as the overall accuracy. As shown in Table 4 and Table 7

B.2 NIAH

We show the detail results of NIAH across Llama-3.1-8B-Instruct and Qwen-2.5-14B-Instruct. We test the KV budget from 256, 512, 1024, 2048, 4096. As shown from Figure 11 to Figure 19

C NHD and HND KV layouts GPU efficiency

As shown in Table 6

D Additional evaluation on Qwen2.5-32B-Ins

In this section, we provide an additional evaluation on Qwen2.5-32B-Instruct using two A100 GPUs, which represents a larger-scale model. As shown in Table 7 and Figure 20, 21, 22.




Figure 9: Important token selection distribution of ground truth top-k and Arkvale on a NIAH task.

Table 4: Accuracy results of LongBench v2

		Overall	Easy	Hard	Short	Medium	Long
<i>Llama-3.1-8B-Instruct (base)</i>		30.8	30.7	30.9	38.3	27.0	25.9
<i>k=4096</i>	Quest	24.85	22.92	26.05	30.0	21.86	22.22
	ArkVale	30.62	28.65	31.83	36.11	27.91	26.85
	ShadowKV	29.62	29.17	29.9	33.89	28.84	24.07
	SnapKV	30.42	28.65	31.51	33.89	30.23	25.0
	PRKV	30.62	28.12	32.15	36.67	28.84	24.07
<i>k=2048</i>	Quest	25.25	22.92	26.69	31.11	23.72	18.52
	ArkVale	30.13	31.49	29.27	36.05	27.04	26.0
	ShadowKV	30.62	29.69	31.19	36.11	28.37	25.93
	SnapKV	30.42	28.12	31.83	36.67	28.37	24.07
	PRKV	30.62	28.65	31.83	37.22	28.84	23.15
<i>k=1024</i>	Quest	18.09	19.27	17.36	18.89	17.67	17.59
	ArkVale	29.56	31.29	28.51	37.12	25.43	26.19
	ShadowKV	30.22	27.6	31.83	33.9	29.3	25.9
	SnapKV	29.62	28.12	30.55	35.56	27.44	24.07
	PRKV	30.82	30.21	31.19	35.56	29.77	25.0
<i>k=512</i>	Quest	10.54	9.9	10.93	12.22	9.77	9.26
	ArkVale	29.22	28.65	29.58	33.89	27.91	24.07
	ShadowKV	30.0	27.6	31.5	34.4	28.4	25.9
	SnapKV	29.03	28.12	29.58	36.67	26.05	22.22
	PRKV	30.62	31.25	30.23	36.11	29.77	23.15
<i>k=256</i>	Quest	5.96	3.65	7.4	8.33	5.58	2.78
	ArkVale	29.22	27.6	30.23	33.89	27.91	24.07
	ShadowKV	29.82	29.17	30.23	36.67	28.37	21.3
	SnapKV	28.43	25.52	30.23	34.44	25.58	24.07
	PRKV	30.42	28.65	31.51	33.89	30.23	25.0

Table 5: Accuracy results of LongBench v2

		Overall	Easy	Hard	Short	Medium	Long
	<i>Qwen-2.5-14B-Instruct (base)</i>	31.6	35.4	29.3	37.2	30.2	25.0
<i>k=4096</i>	Quest	28.48	30.9	26.95	33.73	25.25	26.04
	ArkVale	34.19	41.15	29.9	37.78	34.88	26.85
	ShadowKV	31.81	33.85	30.55	41.67	27.44	24.07
	SnapKV	31.21	33.33	29.9	38.33	29.3	23.15
	PRKV	34.2	40.1	30.5	38.9	34.4	25.9
<i>k=2048</i>	Quest	27.69	28.27	27.33	35.75	23.26	23.15
	ArkVale	34.39	41.67	29.9	38.89	34.42	26.85
	ShadowKV	31.61	34.9	29.58	39.44	28.37	25.0
	SnapKV	31.21	33.33	29.9	38.33	29.3	23.15
	PRKV	35.59	42.71	31.19	38.89	35.81	29.63
<i>k=1024</i>	Quest	25.5	25.65	25.4	29.61	26.98	15.74
	ArkVale	34.59	40.62	30.87	39.44	34.42	26.85
	ShadowKV	32.41	35.94	30.23	39.44	29.77	25.93
	SnapKV	31.21	33.33	29.9	38.33	29.3	23.15
	PRKV	35.39	43.75	30.23	38.33	35.81	29.63
<i>k=512</i>	Quest	22.91	26.18	20.9	26.82	23.26	15.74
	ArkVale	33.4	38.02	30.55	37.22	33.49	26.85
	ShadowKV	29.22	31.77	27.65	36.11	27.91	20.37
	SnapKV	31.21	33.33	29.9	38.33	29.3	23.15
	PRKV	35.19	41.67	31.19	39.44	34.42	29.63
<i>k=256</i>	Quest	20.22	24.72	17.38	22.89	20.2	15.62
	ArkVale	33.4	40.62	28.94	38.33	33.02	25.93
	ShadowKV	30.82	33.33	29.26	38.89	28.37	22.22
	SnapKV	31.21	33.33	29.9	38.33	29.3	23.15
	PRKV	34.0	42.19	28.94	36.11	34.42	29.63




Figure 10: Compare all KV budget in PRKV with full cache in Llama-3.1-8B-Instruct




Figure 11: Compare all KV budget in PRKV with full cache in Qwen2.5-14B-Instruct




Figure 12: Compare all KV budget in Quest with full cache in Llama-3.1-8B-Instruct




Figure 13: Compare all KV budget in Quest with full cache in Qwen2.5-14B-Instruct




Figure 14: Compare all KV budget in Arkvale with full cache in Llama-3.1-8B-Instruct



Figure 15: Compare all KV budget in Arkvale with full cache in Qwen2.5-14B-Instruct




Figure 16: Compare all KV budget in ShadowKV with full cache in Llama-3.1-8B-Instruct




Figure 17: Compare all KV budget in ShadowKV with full cache in Qwen2.5-14B-Instruct




Figure 18: Compare all KV budget in SnapKV with full cache in Llama-3.1-8B-Instruct




Figure 19: Compare all KV budget in SnapKV with full cache in Qwen2.5-14B-Instruct

Table 6: Latency comparison of NHD vs. HND

KV Size	Prefill NHD (ms)	Prefill HND (ms)	Decode NHD (ms)	Decode HND (ms)
4,096	1.0358	1.2567	0.0359	0.0320
8,192	4.9863	4.8339	0.0331	0.0320
16,384	19.1479	19.0742	0.0568	0.0565
32,768	76.4648	77.2803	0.0971	0.0956
65,536	325.6366	317.4488	0.1809	0.1777
131,072	1548.4492	1546.6578	0.3658	0.3981

Table 7: Larger size model results of LongBench v2

		Overall	Easy	Hard	Short	Medium	Long
	<i>Qwen-2.5-32B-Instruct (base)</i>	39.36	40.62	38.59	42.22	38.6	36.11
<i>k=4096</i>	Quest	35.39	30.9	26.95	33.73	25.25	26.04
	ArkVale	37.77	38.54	37.3	41.11	37.67	32.41
	PRKV	38.77	40.1	37.94	42.22	38.14	34.26
<i>k=2048</i>	Quest	33.8	35.42	32.8	37.78	33.95	26.85
	ArkVale	37.38	38.02	36.98	41.11	36.74	32.41
	PRKV	37.77	38.54	37.3	41.11	37.21	33.33
<i>k=1024</i>	Quest	33.4	34.9	32.48	37.78	33.02	26.85
	ArkVale	36.18	36.98	35.69	41.11	34.88	30.56
	PRKV	36.78	36.98	36.66	41.11	35.81	31.48




Figure 20: Compare all KV budget in PRKV with full cache in Qwen2.5-32B-Instruct




Figure 21: Compare all KV budget in Quest with full cache in Qwen2.5-32B-Instruct




Figure 22: Compare all KV budget in Arkvale with full cache in Qwen2.5-32B-Instruct

E Rebuttal

E.1 CPU Memory Footprint

PRKV preserves the full KV cache on CPU, so its CPU memory footprint is the actual KV cache size:

$$M_{\text{CPU}} = 2 \cdot L \cdot B \cdot N \cdot H_{\text{kv}} \cdot D \cdot b,$$

where L is the number of layers, H_{kv} is the number of KV heads, D is the head dimension, N is the context length, B is the batch size, and b is the number of bytes per KV element. The leading factor 2 accounts for both the key and value caches.

During PRKV system initialization, we pre-allocate sufficient CPU memory for storing the offloaded KV cache, following a strategy similar to vLLM, which by default allocates up to 80% of the available CPU memory. In our experiments, the server has 1TB CPU memory, using 16 DDR4 memory modules of 64GB each.

Table 8 reports the actual CPU-side KV cache footprint of PRKV. With BF16 KV cache, batch size $B = 1$, and 128K context length, the full CPU-side KV cache requires 16.0GB for Llama3.1-8B-Instruct and 24.0GB for Qwen2.5-14B-Instruct.

Model	L	H_{kv}	D	KV Size
Llama3.1-8B-Instruct	32	8	128	16.0GB
Qwen2.5-14B-Instruct	48	8	128	24.0GB

Table 8: CPU-side KV cache footprint of PRKV with BF16 KV cache, batch size $B = 1$, and 128K context length.

E.2 clearer description of the asynchronous update

The asynchronous KV rewrite is triggered every T decoding steps when PRKV updates the static token set. At such steps, PRKV first gathers and restructures the selected static KV pages on GPU, as in Lines 23–24 of Algorithm 1, and then writes these updated static KV pages back to CPU memory asynchronously.

PRKV maintains two mapping tables to locate the CPU write-back destination: `gc2cc` maps each GPU cache page id to its CPU cache page id, and `c2p` maps the CPU cache page id to its physical location in the CPU KV pool. With these mappings, PRKV can asynchronously copy the restructured K/V pages to their corresponding CPU locations on a background stream. Decoding does not block on this write-back; the updated CPU layout is used after the asynchronous copy completes. The detailed

procedure is summarized in Algorithm 3, and the left part of Figure 23 illustrates how the restructured static KV pages are written back to the CPU KV pool.

Algorithm 3 Asynchronous KV Write-back

```

1: Input: restructured GPU static KV pages  $K'_{\text{GPU}}$  and  $V'_{\text{GPU}}$ , CPU KV pool  $K_{\text{CPU}}$  and  $V_{\text{CPU}}$ , mapping tables gc2cc and c2p, static GPU page set  $\mathcal{P}_{\text{static}}$ , batch index  $b$ , group index  $g$ 
2: Output: updated CPU KV cache
3: for each GPU static page id  $p_g \in \mathcal{P}_{\text{static}}$  do
4:   # GPU cache page id  $\rightarrow$  CPU cache page id
5:    $p_c \leftarrow \text{gc2cc}[b, g, p_g]$ 
6:   # CPU cache page id  $\rightarrow$  CPU pool page id
7:    $p_{\text{pool}} \leftarrow \text{c2p}[b, p_c]$ 
8:   # copy kernel runs on background CUDA stream
9:    $\text{ASYNCCOPY}(K_{\text{CPU}}[p_{\text{pool}}], K'_{\text{GPU}}[p_g])$ 
10:   $\text{ASYNCCOPY}(V_{\text{CPU}}[p_{\text{pool}}], V'_{\text{GPU}}[p_g])$ 
11: end for

```

E.3 Robustness of Online Static Ratio Estimation

We clarify that the variation of S in Table 1 reflects different temporal stability of important tokens across models and context lengths, rather than instability of the estimator. PRKV estimates S using the top- k overlap between the first and last query tokens in the observation window, which provides a conservative signal of how persistent the important KV set is.

We further validate this with a diagnostic analysis. Let C_0 be the static set selected after prefill and C_t be the ground-truth top- k set at decoding step t . We compute

$$C_{\text{future}} = \frac{1}{T_{\text{dec}}} \sum_{t=1}^{T_{\text{dec}}} \frac{|C_0 \cap C_t|}{|C_0|}.$$

As shown in Table 9, across 50 sampled LongBench V2 and NIAH examples, C_{future} remains high (0.72–0.81), indicating that the selected static set continues to cover important tokens during decoding.

PRKV is also robust to moderate estimation errors. If S is overestimated, dynamic page retrieval and the local window still provide complementary coverage; if S is underestimated, more budget is




Figure 23: Detail design of recomputes the static set and restructures pages

assigned to dynamic retrieval, mainly affecting efficiency rather than accuracy. For highly heterogeneous prompts, the first-last overlap tends to decrease, leading PRKV to choose a smaller S and rely more on dynamic retrieval. We will add this robustness discussion in the revision.

Dataset	Model	Budget	C_{future}
LongBench V2	Llama3.1-8B-Instruct	1024	0.79
LongBench V2	Qwen2.5-14B-Instruct	1024	0.72
NIAH	Llama3.1-8B-Instruct	1024	0.81
NIAH	Qwen2.5-14B-Instruct	1024	0.80

Table 9: Diagnostic analysis of the online static ratio estimation on 50 sampled examples from LongBench V2 and NIAH. C_{future} measures the average overlap between the selected static token set after prefill and the ground-truth top- k token set during decoding.

E.4 Overhead of Periodic Static KV Update

CPU-side restructuring overhead. We clarify that PRKV performs page restructuring mainly on the GPU side. When the static token set is updated, the selected static KV pages are gathered and reorganized into a contiguous layout on GPU. The CPU side does not redo the full restructuring; instead, it asynchronously writes the already restructured static KV pages back to the CPU KV cache. Therefore, the main CPU-side cost is asynchronous KV rewrite rather than gathering scattered tokens.

To quantify this cost, we measure the breakdown of page restructuring and asynchronous update overhead. The experiment uses 64K input length,

Model	Restr. /layer	Async /layer	TPOT Common	TPOT Update
Llama3.1-8B	0.152	10.934	37.64	46.71
Qwen2.5-14B	0.184	11.123	67.79	76.05

Table 10: Breakdown of periodic static KV update overhead on LongBench V2. The input length is 64K, the update interval is $T = 128$, and the KV budget is $k = 4096$. All values are reported in ms.

update interval $T = 128$, and top- $k = 4096$. All values are reported in milliseconds. TPOT (common step) denotes the TPOT of a normal decoding step, while TPOT (update step) denotes the TPOT when static-token update is triggered.

As shown in Table 10, the page restructuring cost itself is small. Although the per-layer asynchronous update takes about 11 ms, it runs in the background and overlaps with the computation of subsequent layers, so it is not fully exposed on the critical path. Consequently, the update step only increases TPOT by 9.07 ms and 8.26 ms for Llama3.1-8B and Qwen2.5-14B, respectively. Since the update is triggered only once every $T = 128$ steps, the amortized overhead is only 0.071 ms/token and 0.065 ms/token, respectively. Therefore, the CPU-side asynchronous rewrite introduces negligible amortized overhead and does not significantly affect end-to-end latency.

E.5 Model Consistency

We use Qwen3 for the reasoning evaluation because it has stronger long-chain reasoning capability and can generate sufficiently long reasoning traces on AIME24 and MATH500. In contrast, Llama3.1-8B and Qwen2.5-14B did not consistently produce long outputs, so they are used for general long-context tasks, where the focus is long-context understanding and retrieval. To further address model consistency, we additionally include results on DeepSeek-R1-Llama-8B for AIME24 and MATH500 in the appendix. As shown in Table 11, PRKV maintains high accuracy on DeepSeek-R1-Llama-8B and consistently outperforms other sparse KV methods.

Methods	AIME24	AIME24	MATH500	MATH500
	pass@k	avg@k	pass@k	avg@k
<i>DeepSeek-R1-Llama-8B</i>	85.00	77.83	96.80	96.00
Quest	72.33	67.67	84.20	77.80
ArkVale	76.00	66.33	88.00	87.90
ShadowKV	78.67	68.50	87.00	84.20
SnapKV	79.67	68.33	92.80	92.60
PRKV	80.67	72.00	95.00	93.80

Table 11: Reasoning results on DeepSeek-R1-Llama-8B for AIME24 and MATH500. PRKV maintains high accuracy while outperforming other sparse KV methods.

E.6 Ablation on T

T controls the static-token update frequency. $T = 1$ updates the static set at every decoding step, which is an extreme rather than practical setting, since each update requires recomputing token importance, reorganizing static KV tokens, updating page representations, and asynchronously writing the updated layout back to CPU memory. PRKV instead updates only when $ds \bmod T = 0$, amortizing this cost across decoding steps.

T	Qwen2.5-14B			Llama3.1-8B		
	32K	64K	128K	32K	64K	128K
$T = 1$	361.85	394.07	440.70	203.94	249.38	295.69
$T = 32$	112.35	130.12	162.45	72.18	79.40	94.85
$T = 64$	88.40	102.55	124.20	58.92	62.15	71.30
$T = 128$	55.95	69.90	86.98	32.01	34.21	40.46
$T = 256$	53.83	68.15	84.22	30.79	33.76	39.03

Table 12: Latency ablation of different static-token update intervals. The output length is fixed to 2K tokens. All values are reported in milliseconds, and 32K/64K/128K denote input lengths.

Our default $T = 128$ balances recall and latency. Table 12 shows the latency ablation with a fixed

2K output length. All values are reported in milliseconds, and 32K/64K/128K denote input lengths. The results show that $T = 1$ introduces substantial overhead, while $T = 256$ only marginally reduces latency over $T = 128$ but significantly hurts recall, as shown in Table 3. Thus, $T = 128$ is a better default configuration.