

FaTRQ: Tiered Residual Quantization for LLM Vector Search in Far-Memory-Aware ANNS Systems

Tianqi Zhang
UC San Diego
La Jolla, CA, USA
tiz014@ucsd.edu

Flavio Ponzina
San Diego State University
San Diego, CA, USA
fponzina@sdsu.edu

Tajana Rosing
UC San Diego
La Jolla, CA, USA
tajana@ucsd.edu

Abstract—Approximate Nearest-Neighbor Search (ANNS) is a key technique in retrieval-augmented generation (RAG), enabling rapid identification of the most relevant high-dimensional embeddings from massive vector databases. Modern ANNS engines accelerate this process using prebuilt indexes and store compressed vector-quantized representations in fast memory. However, they still rely on a costly second-pass refinement stage that reads full-precision vectors from slower storage like SSDs. For modern text and multimodal embeddings, these reads now dominate the latency of the entire query. We propose FaTRQ, a far-memory-aware refinement system using tiered memory that eliminates the need to fetch full vectors from storage. It introduces a progressive distance estimator that refines coarse scores using compact residuals streamed from far memory. Refinement stops early once a candidate is provably outside the top-k. To support this, we propose tiered residual quantization, which encodes residuals as ternary values stored efficiently in far memory. A custom accelerator is deployed in a CXL Type-2 device to perform low-latency refinement locally. Together, FaTRQ improves the storage efficiency by 2.4× and improves the throughput by up to 9× than SOTA GPU ANNS system.

Index Terms—Approximate Nearest-Neighbor Search, Residual Quantization, Tiered Memory, CXL

I. INTRODUCTION

Nearest-neighbor search (NNS) is the key component of modern data systems supporting tasks like semantic retrieval, recommendations, fraud detection, and large-scale ranking. While exact NNS becomes intractable at scale, approximate nearest-neighbor search (ANNS) offers high recall at much lower cost, and is now standard in production-scale systems.

A leading use case of ANNS is retrieval-augmented generation (RAG). As shown in Figure 1, RAG pipelines embed each document chunk once, cache the resulting vectors, and, at query time, embed the user prompt and launch an ANNS search whose results are passed to the language model. Modern dense embedding models such as OpenAI’s models [1] generate 1536-dimensional vectors, corresponding to about 6 kB per vector in full precision. For knowledge bases containing millions or even billions of entries [2], this quickly exceeds the capacity of main memory and renders exact search infeasible.

To reduce memory pressure and accelerate search, modern engines combine indexing data structures (e.g., IVF [3]) with compact vector codes, such as the product quantization family [4]–[6]. These methods quantize high-dimensional vectors into short codes, shrinking, for instance, a 6 kB floating-point vector to about 200B, so that the full index can fit in main memory. At query time, the graph guides the search process, and distances are estimated with the compact codes.




Fig. 1. ANNS is the key component of the RAG pipeline.

However, quantization comes at a cost: to recover full accuracy, systems re-rank the long candidate list by fetching original full-precision vectors from storage and recomputing exact distances. For modern dense embeddings [1], [7], this refinement step now dominates query latency. Our profiling detailed in Section II-A shows that over 90% of query time can be spent on reading vectors from storage. Worse, the widening gap between memory and storage bandwidth makes this bottleneck increasingly severe.

At the same time, the tiered memory technologies such as CXL-based memory expander and storage-class memory (SCM) have opened opportunities to rethink this refinement stage. While compact codes are much smaller than full-precision vectors, the aggregate working set of these codes can still exceed local DRAM (fast memory) capacity, especially for large-scale databases or multi-tenant workloads. In such cases, these tiered memories (far memory) offer an attractive option: they sit between local DRAM and SSD in both capacity and latency, making them well-suited for storing compact intermediate data structures rather than full vectors. Yet, existing ANNS systems [8]–[11] rarely exploit these tiers and continue to treat refinement I/O as an unavoidable bottleneck.

To leverage this opportunity, we propose **FaTRQ**, a **F**ar-memory-aware refinement system with **T**iered **R**esidual **Q**uantization that eliminates most expensive I/O in high-accuracy ANNS. The key idea is progressive distance estimation: rather than fetching full-precision vectors, FaTRQ incrementally refines coarse distances using compact residual codes stored in far memory. To support this, FaTRQ introduces *tiered residual quantization* (TRQ), where the “T” reflects both its *tier-aware* memory design and its *ternary* codebook. TRQ encodes the residual between a record and its coarse quantization into ternary values, stored in far memory and streamed during query time. The estimator builds on the vector decomposition and a light-weight distance estimation

calibration model, allowing residual contributions to be accumulated without reconstructing full vectors. Operating directly on ternary codes, it refines distances using only additions and subtractions, avoiding multiplications while preserving high accuracy. Though FaTRQ is applicable in software on tiered memory systems, we also built a proof-of-concept CXL Type-2 accelerator with lightweight custom logic, showing how FaTRQ can be integrated into emerging far-memory hardware.

We summarize our contributions as follows:

- Multi-level vector quantization for tiered memory: We design a hierarchical compression format where coarse quantized vectors reside in fast memory, and compact residual codes are stored in slower memory tiers, enabling fine-grained accuracy control with minimal I/O.
- Progressive distance estimation: We propose an incremental scoring mechanism that updates coarse distances without reconstructing vectors, reducing memory traffic from hundreds of bytes to 4 bytes.
- Accelerated refinement pipeline: We prototype a far-memory accelerator for residual refinements, reducing data movement and host CPU overhead.
- Key Results: Improves the ANNS throughput by $2.6\times$ to $9.4\times$ compared to GPU ANNS system [9], and enhances storage efficiency by $2.4\times$ than the refinement scheme in the SoTA pipeline [12].

II. BACKGROUND AND RELATED WORKS

A. ANNS System

Modern ANNS engines combine indexes (e.g., IVF [3], CAGRA [13]) with vector quantization to keep large-scale collections in memory while supporting fast distance estimation [3], [11]–[15]. The index prunes the search space, while quantized codes enable low-cost scoring. During index traversal, the system fetches quantized codes and identifies a subset of candidates. In RAG workloads, however, modern high-dimensional embeddings (e.g., 768 dimensions [7], [16]) require aggressive quantization to fit into memory, which reduces recall and necessitates a second-pass refinement that retrieves full-precision vectors from SSD [14], [17]. Figure 2 shows this refinement dominates latency on GPU-accelerated pipeline [9] where the index and quantized vectors reside in GPU memory while full-precision vectors are mmaped on the host. On an A10 GPU with a 40-thread CPU, index traversal accounts for only 2–15% of query time due to GPU acceleration, whereas refinement dominates due to random SSD I/O and distance computation. If all vectors could reside in main memory, an impractical assumption at large scale, performance would improve by up to $14\times$. This unattainable upper bound motivates our progressive distance estimation approach, which incrementally refines candidates using compact residual codes: reducing full-vector fetches, cutting refinement overhead, and avoiding the capacity ceilings of hardware-centric solutions.

B. Vector Quantization and Distance Estimation

Vector quantization compresses high-dimensional vectors into compact codes, reducing storage cost and enabling efficient




Fig. 2. Runtime breakdown of IVF-refinement ANNS system

distance computation. The simplest form, *scalar quantization* (SQ), discretizes each dimension independently, but scales poorly with vector length. *Product quantization* (PQ) [3], [4], [18] improves scalability by partitioning a vector into subspaces and by quantizing each with a separate codebook, enabling efficient table-based distance lookups and GPU acceleration [3], [12]–[14]. *Residual quantization* (RQ) [19], [20] further enhances accuracy by encoding a vector as a coarse approximation plus successive residual corrections, often using PQ at each stage. To estimate distances, most systems employ *asymmetric distance computation* (ADC) [21], which compares a full-precision query with quantized database codes:

$$\hat{d}(q, x) = d\left(q, \sum_{i=1}^L Q^{-1}(c_i)\right),$$

where c_i is the i -th layer RQ code. While effective, existing systems [3], [9] decode all quantization levels for every vector during traversal, even though most candidates are filtered out shortly thereafter. As a result, 10–100 \times more candidates are decoded than necessary, wasting bandwidth on vectors that are immediately discarded. FaTRQ addresses this inefficiency by enabling progressive ADC over residual quantization: early RQ levels are stored in fast memory while deeper residual codes are tiered into the slower far memory. Distance estimation proceeds incrementally, allowing early pruning of unpromising candidates and avoiding full reconstruction for most vectors.

C. Related Works

Recent work reduces ANNS memory cost by leveraging high-bandwidth memory systems such as CXL [8], persistent memory [10], near-data-processing (NDP) [22]–[25], and GPUs [12]–[14]. They either accelerate table lookups on quantized codes or assume full-precision vectors are directly available, leaving refinement to slow storage reads. CXL-ANNS [8] stores full vectors in a CXL memory and offloads distance computation to FPGA, but its scalability collapses once the original dataset cannot fit the memory capacity. HM-ANN [10] maps HNSW [26] index entry layers to fast DRAM–NVM tiers, but for high-dimensional vectors, additional layers provide little pruning benefit while increasing management overhead [27]. NDP designs [22]–[25] integrate distance computation with quantized vectors directly into memory devices, but leave refinement unaddressed. NDP accelerations are complementary to our work. We focus on progressive, tier-aware refinement that minimizes SSD access and can naturally benefit from or be combined with NDP acceleration across fast and far memory.

III. FATRQ: ESTIMATE THE DISTANCE PROGRESSIVELY

To reduce the costly I/O and computation of second-pass refinement, we propose FaTRQ, a tier-aware framework that




Fig. 3. Memory-tiered layout of the FaTRQ framework

incrementally sharpens coarse distance estimates. As shown in Figure 3, coarse PQ codes and the codebook remain in fast memory, while compact residual codes and metadata are stored in far memory, reducing high-latency SSD accesses. The refinement builds on a decomposition of the L2 distance into coarse and residual terms, which can be progressively estimated without vector reconstruction. Residuals are encoded into compact ternary codes that eliminate multiplications and pack efficiently into far memory. This compact format enables progressive refinement: first by residual distance estimation, then by multiplication-free encoding, and finally by an enhanced estimator that integrates all components into accurate L2 distance. Together, these steps allow FaTRQ to prune non-promising candidates early and cut SSD traffic.

A. L2 Distance Decomposition

Unlike prior systems [11], [12], [17] that perform reranking using a separate quantization code built solely for refinement, FaTRQ progressively reuses the coarse approximation already computed in earlier stages. This reuse avoids discarding useful information from the coarse code and enables refinement to be expressed cleanly through an L2 distance decomposition:

$$\|x - q\|_2^2 = \|x_c - q\|_2^2 + \|x_c - x\|_2^2 - 2\langle q - x_c, x - x_c \rangle,$$

where q is the query, x_c is the reconstructed vector of original vector x from the coarse code, and $\langle \cdot, \cdot \rangle$ is the inner product. The distance is split into three terms: the coarse approximation, the compression distortion, and a residual inner product.

Figure 4 shows this decomposition on the Wiki dataset [28]. For each record vector, we align its reconstruction of coarse quantization x_c to the origin (circle center) and scale the residuals $x - x_c$ to lie on the unit circle. The query offset $q - x_c$ is placed along the x-axis, with length proportional to $\frac{\|q - x_c\|}{\|x - x_c\|}$. In this representation, the original records x (blue points) lie on the unit circle, and queries q (red points) are located along the x-axis. It highlights that the residual is nearly orthogonal to the query offset, so their inner product is small.

This leads to a *first-order approximation*:

$$\hat{d}_1(q, x) = \|x_c - q\|_2^2 + \|x_c - x\|_2^2 = \hat{d}_0(q, x) + \|x_c - x\|_2^2,$$

where $\hat{d}_0(q, x) = \|x_c - q\|_2^2$. Note that $\|x_c - x\|_2^2$ is a scalar value associated with each record vector and can be precomputed offline. To fully recover the true distance, we must estimate residual item $\langle q - x_c, x - x_c \rangle$.

Letting $\delta = x - x_c$ denote the residual between the record and its coarse approximation, we have

$$\|x - q\|_2^2 = \|q - x_c\|_2^2 + \|\delta\|_2^2 + 2\langle x_c, \delta \rangle - 2\langle q, \delta \rangle.$$




Fig. 4. Visualization of the residual vector and query vector

The first three terms can be computed using only the coarse quantization code and precomputed scalars. The last term, $\langle q, \delta \rangle$, is estimated via quantized residuals without decoding the base vector. Compared to storing full raw vectors, residuals are smaller, since the base code already captures most of the vector structure. Storing them in far memory offers a better trade-off between latency and capacity. We refer to this refinement using quantized residual vectors as the *second-order approximation*. Since residual quantization is naturally stackable, distance estimates can be progressively refined. For example, we can first encode the residual on top of the coarse code, and then refine it further by encoding finer residuals on the remaining error, enabling progressively tighter distance estimates. Without loss of generality, we focus on second-order estimation in the following discussion. The next section details how δ is quantized and used to obtain an unbiased estimate of the inner product term.

B. FaTRQ Residual Distance Estimation

To estimate the inner product $\langle q, \delta \rangle$, we first quantize the residual vector δ to a codeword δ_c , and have

$$\langle q, \delta \rangle = \|q\|_2 \|\delta\|_2 \langle e_q, e_\delta \rangle,$$

where $e_q = q/\|q\|_2$ and $e_\delta = \delta/\|\delta\|_2$ are the normalized directions of the query q and the residual δ . Our goal is then to obtain an unbiased estimator of the directional term $\langle e_q, e_\delta \rangle$.

Observe that by adding and subtracting the projection of e_q onto e_{δ_c} , the direction of δ_c , we can write

$$\begin{aligned} \langle e_q, e_\delta \rangle &= \langle e_q, e_\delta \rangle + \langle e_q, e_{\delta_c} \rangle \langle e_\delta, e_{\delta_c} \rangle - \langle e_q, e_{\delta_c} \rangle \langle e_\delta, e_{\delta_c} \rangle \\ &= \langle e_q, e_{\delta_c} \rangle \langle e_{\delta_c}, e_\delta \rangle + \langle e_q - \langle e_q, e_{\delta_c} \rangle e_{\delta_c}, e_\delta \rangle \\ &= \langle e_q, e_{\delta_c} \rangle \langle e_{\delta_c}, e_\delta \rangle + \|e_q - \langle e_q, e_{\delta_c} \rangle e_{\delta_c}\|_2 \langle e_\perp, e_\delta \rangle. \end{aligned} \quad (1)$$

This equation decomposes e_q into components aligned with and orthogonal to e_{δ_c} . e_\perp is the unit vector orthogonal to e_{δ_c} in the plane spanned by e_{δ_c} and e_q .

This can be viewed as the dual form of the query disaggregation of previous work [29] that disaggregates it to the original direction e_δ , instead of the quantized direction e_{δ_c} in our work, and corresponding orthogonal basis. That work shows the orthogonal term (second term of Equation 1) is concentrated and its expectation is zero under mild conditions, specifically, when residual directions are evenly distributed and uncorrelated with the query. In our case, because the coarse quantization already captures most of the vector similarity, the residual vectors δ are nearly isotropic and uncorrelated with the query. This removes the need for an additional random projection step, allowing us to exploit the information encoded

in each residual code well. Consequently, we estimate the first term while treating the second as an error term having zero expectation, yielding an unbiased residual distance estimator.

C. FaTRQ Multiplication-Free Encoding

To enable efficient estimation of $\langle q, \delta \rangle$, FaTRQ encodes the residual direction with a sparse ternary vector. This compact representation can avoid multiplications during refinement and reduce the memory footprint. Concretely, given a residual vector $e_\delta \in \mathbb{R}^D$, we seek the code \bar{e}_{δ_c} in the codebase $\mathcal{C} = \{-1, 0, 1\}^D$, so that its normalized version

$$e_{\delta_c} = \frac{\bar{e}_{\delta_c}}{\|\bar{e}_{\delta_c}\|_2}$$

minimizes $\|e_\delta - e_{\delta_c}\|_2$. Observing that

$$\|e_\delta - e_{\delta_c}\|_2^2 = 2 - 2 \langle e_{\delta_c}, e_\delta \rangle = 2 - 2 \left\langle \frac{\bar{e}_{\delta_c}}{\|\bar{e}_{\delta_c}\|_2}, e_\delta \right\rangle,$$

we equivalently choose

$$\bar{e}_{\delta_c} = \arg \max_{c \in \mathcal{C}} \left\langle \frac{c}{\|c\|_2}, e_\delta \right\rangle = \arg \max_{c \in \mathcal{C}} \frac{\sum_{i=1}^D c_i e_{\delta,i}}{\|c\|_2}.$$

Direct enumeration of all 3^D candidates is infeasible. We solve this optimization by observing that any optimal c must select exactly k nonzero entries (each set to the sign of $e_{\delta,i}$) among the k largest magnitudes of e_δ . It would strictly decrease the numerator of the inner product if including any smaller-magnitude $e_{\delta,i}$ in place of a larger one, while changing the total count of nonzero entries k only affects the denominator \sqrt{k} . Thus, the best choice for a given k is to take the first k entries of the sorted list. Concretely, by sorting the absolute values into $x_1 \geq x_2 \geq \dots \geq x_D$, where $x_i = |e_{\delta,(i)}|$, we reduce the inner-product term to

$$\left\langle \frac{c}{\|c\|_2}, x \right\rangle = \frac{\sum_{i=1}^k x_i}{\sqrt{k}} \quad \text{whenever } \sum_i c_i^2 = k.$$

Therefore, the global maximization over \mathcal{C} collapses to one-dimensional search

$$\max_{c \in \mathcal{C}} \left\langle \frac{c}{\|c\|_2}, x \right\rangle = \max_{1 \leq k \leq D} \frac{\sum_{i=1}^k x_i}{\sqrt{k}}.$$

The proposed approach requires only $O(D \log D)$ time to sort x_i values, $O(D)$ time to build the prefix sums of x_i , and another $O(D)$ pass to evaluate each ratio S_k/\sqrt{k} and identify the maximizer k^* . Once k^* is known, we form \bar{e}_{δ_c} by assigning $\text{sign}(e_{\delta,(i)})$ to the top k^* entries and zeros elsewhere, then normalize by $\sqrt{k^*}$. This procedure yields the exact optimal ternary codeword in $O(D \log D)$ time without ever enumerating the full codebook \mathcal{C} .

D. Compact FaTRQ Code

Figure 3 illustrates the data layout used by FaTRQ to store per-record information in far memory during the refinement phase. Each record includes two scalar values: $\langle x_c, \delta \rangle$ and $\|\delta\|_2$. They serve as scaling metadata for computing the final distance estimate, along with one quantized residual vector \bar{e}_{δ_c} that captures the remaining error after level-1 approximation.

The residual vector \bar{e}_{δ_c} contains elements from a 3-level quantization set $\{-1, 0, 1\}$. Each element, therefore, requires more than 1 bit to encode, but it is unnecessary and wasteful to use a full 2 bits per value, as that would encode 4 states while only 3 are needed. The entropy of a uniformly distributed 3-level variable is $\log_2 3 \approx 1.58$ bits, which sets the theoretical limit for lossless compression.

To approach this bound in practice, FaTRQ uses a fixed-length compact encoding where every 5 dimensions of the residual vector are packed into a single byte. This is achieved by first mapping each value $x_i \in \{-1, 0, 1\}$ to the set $\{0, 1, 2\}$ using the transformation $x_i + 1$, and then encoding the resulting 5-digit base-3 number into a byte using the formula below. This encoding yields a per-dimension storage cost of 1.6 bits.

$$y = \sum_{i=0}^4 3^i (x_i + 1).$$

E. FaTRQ Enhanced Refinement Distance Estimator

With the decomposition of the distance function and the residual encoding already introduced, the final step is to wrap these components into the refinement estimator. Specifically, the L2 distance can be expressed as a combination of the coarse approximation \hat{d}_0 , precomputed scalar $\|\delta\|_2^2$ and $2\langle x_c, \delta \rangle$, and the estimated residual term $-2\langle q, \delta \rangle$.

An important insight emerges when evaluating the effectiveness of these estimations. Lower mean squared error (MSE) in distance approximation does not always lead to higher ANNS recall. This is because recall is primarily determined by how accurately candidates are ranked near the top- k decision boundary. Large estimation errors for very close or very distant points have little impact as they do not affect ranking within the critical region. Thus, improving global distance accuracy does not guarantee better recall; what matters is local precision around the decision boundary.

Based on this observation, FaTRQ learns a lightweight linear calibration model offline that directly optimizes for recall in the boundary region by subsampling a small calibration set of sample-neighbor pairs. Specifically, it randomly draws a small number of database vectors (about 0.3% in practice, which we found sufficient to saturate the model). For each sampled vector x , it leverages the existing index to identify approximate neighbors without requiring a costly exact kNN search: if the index is graph-based (e.g., CAGRA [13] or HNSW [26]), it uses the graph-adjacent vertices of x ; if the index is IVF-based, it uses the vectors in the same inverted list. These points naturally cluster near x in the original space and provide dense coverage of its local decision boundary.

For each sample-neighbor pair, we construct feature vector:

$$A = \left[\hat{d}_0, \hat{d}_{ip}, \|\delta\|_2^2, \langle x_c, \delta \rangle \right],$$

where $\hat{d}_0 = \|q - x_c\|_2^2$ is the coarse approximation, \hat{d}_{ip} is the refined estimate of $-2\langle q, \delta \rangle$ via ternary coding, $\|\delta\|_2^2$ is the residual norm, $\langle x_c, \delta \rangle$ is the precomputed cross-term. Let D denote the ground-truth distance. The calibration model learns in the offline stage by solving

$$\hat{W} = \arg \min_W \|D - AW\|_2^2,$$

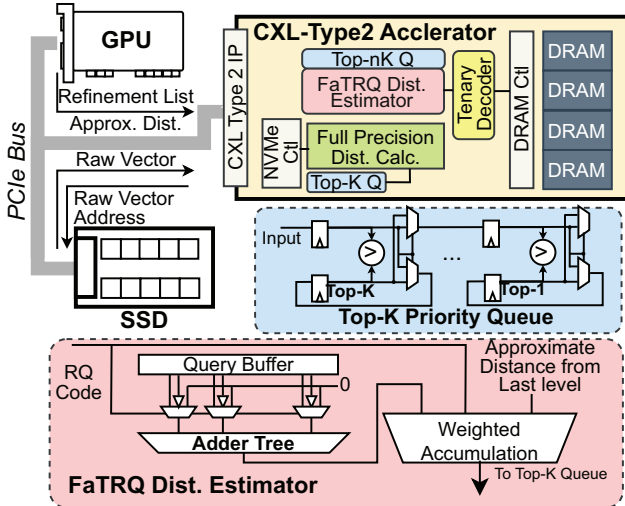


Fig. 5. System architecture of the FaTRQ-augmented ANNS pipeline.

via ordinary least squares. At query time, refinement distance estimation reduces to the lightweight computation of $A\hat{W}$.

IV. FATRQ-AUGMENTED ANNS SYSTEM

While FaTRQ targets a broad range of tiered memory systems, we also explore its benefits on modern far-memory devices. A CXL Type-2 module offers both the large capacity of far-memory (e.g. SCM) and, additionally, lightweight computation. Our prototype integrates customized refinement logic into such a device, pushing refinement directly into far memory to reduce host-device data movement. As compute-enabled large memory devices have already been commercialized [30], this demonstrates how FaTRQ can exploit the emerging hardware model, while our software implementation shows it is not tied to specialized far-memory hardware.

As shown in Figure 5, the front-stage GPU processes the index using coarse quantization codes to estimate approximate distances and generates a large candidate list. Instead of transferring hundreds of bytes of coarse code, only 4-byte coarse distance values per candidate are sent to the far-memory device equipped with the FaTRQ refinement logic. Unlike prior approaches [19], [20] that reconstruct record vectors from both coarse and residual codes, FaTRQ refines directly on compact residual codes, avoiding reconstruction. The far-memory accelerator reranks the candidates using its on-device distance estimator and trims the list based on updated scores. Only the top results are then fetched from SSD for final computation, cutting refinement cost by 85% (Section V-B).

Figure 5 also shows the detailed implementation. Two hardware priority queues, implemented using registers and comparators, are used to track the top- K nearest neighbors during refinement. One queue maintains the top- K candidates ranked by estimated distances from FaTRQ’s residual quantization path, while the other is used for final ranking based on full-precision vector distances. Each entry in the queue stores a distance value and a pointer to the corresponding vector. New candidates are inserted by comparing their distance to those in the queue, and bubbling smaller values forward through the pipeline of comparators. Each queue supports up to 1024 entries. A 256-entry

TABLE I
PARAMETERS FOR FATRQ EVALUATION

Parameter	Value
DRAM Configuration	8Gb x16 DDR5-4800
Timing (tRCD-tCAS-tRP)	34-34-34
Channels / Ranks per Channel	8 / 8
SSD Latency / Throughput [31]	45 μ s / 1200K IOPS
CXL Latency / Throughput [30]	271 ns : 22 GB/s

lookup table implements the ternary decoder, which unpacks residual codes into their ternary representation. Because the residual codes in FaTRQ are ternary, the inner product $\langle e_q, e_{\delta_c} \rangle$ between the query and residual vector can be computed using simple adders and multiplexers. A weighted accumulation unit, implemented as a MAC array, combines these results to produce the final estimated distance, as described in Section III-E.

V. EVALUATION

A. Experimental Setup

Software and Baselines: We implement FaTRQ as an extension to the cuVS [9] and FAISS [3] libraries. Our design stacks the proposed residual codes on top of product quantization (PQ) as a coarse quantizer, and integrates them with both IVF [3] and CAGRA [13] index structures to form the complete ANNS system. For evaluation, we use the FAISS GPU release for IVF and the GPU-oriented graph index CAGRA from cuVS. All parameters are tuned via grid search [13]. We compare these FaTRQ-enhanced systems against the baseline pipelines provided in cuVS [13], [17].

Datasets: We use two pre-embedded ANNS datasets. a) The *Wiki* dataset [28] (251 GB) contains 88M SBERT [7] multilingual sentence embeddings (768-D). b) The *LAION* dataset [32], the largest pre-embedded collection from VDBBench [33], contains 100M CLIP [16] embeddings derived from LAION-5B [32], totaling 286 GB. Both datasets use Euclidean distance as the similarity metric with 10k queries.

Platform: We extend Ramulator [34] to simulate far-memory access patterns of a CXL Type-2 device, with DRAM parameters listed in Table I. Hardware overhead is assessed by synthesizing our accelerator in Verilog at 1 GHz using the ASAP7 [35]. On-chip SRAM is modeled with FinCACTI [36]. All GPU-side index traversal and coarse distance computations run on an NVIDIA A10 GPU (24 GB VRAM), while refinement in the baseline systems executes on a 40-thread Intel Xeon Gold 6230 CPU with 128 GB DRAM and a 1 TB SSD.

B. Overall Performance

Figure 6 compares end-to-end throughput of FaTRQ against SoTA GPU pipelines, IVF-FAISS [3] and CAGRA-cuVS [9], at 85%, 90%, and 95% recall for top-10 queries. On the LAION, recall saturates at 94% when PQ codes fit in GPU VRAM, so results for LAION-95 are omitted. The baselines store quantized vectors in GPU memory and perform refinement on the CPU by fetching full-precision vectors from SSDs. We evaluate FaTRQ with residual codes stored in far memory: in the software mode (-SW), codes reside in CXL memory, but




Fig. 6. FaTRQ normalized throughput evaluation for different top-10 query recall rates on both IVF and CAGRA front-stage indexes.

filtering is done on the CPU, while in the hardware mode (HW), filtering is offloaded to a CXL Type-2 accelerator.

FaTRQ-HW delivers $3.1\times$ – $9.4\times$ speedup over IVF-FAISS and $2.6\times$ – $4.9\times$ over CAGRA-cuVS. The benefit is larger with IVF because it requires more refinements. For example, at 90% recall on Wiki, IVF refines 320 candidates per query, compared to 120 for CAGRA. In the baseline, this means 320 vs. 120 SSD fetches. With FaTRQ, those become 28 SSD plus 320 CXL accesses for IVF, and 17 SSD plus 120 CXL accesses for CAGRA, shifting most refinement traffic off SSD and yielding bigger improvements for IVF. The speedup narrows at 95% recall, as deeper traversal and coarse filtering required for high accuracy begin to dominate runtime. We also observe the hardware variant adds $1.2\times$ – $1.5\times$ higher throughput over the software path, with candidate filtering up to $3.7\times$ faster due to direct far-memory access and removal of host data movement.

C. Evaluation on Distance Estimation Distortion

As discussed in Section III-E, the accuracy of ANNS depends on the distance estimation between the query and a small subset of the collection that is close to the query. To evaluate the distance distortion introduced by FaTRQ, we calculate the distance between each query and the top-100 ground truth results obtained via exhaustive search. Figure 7 compares distance distortion on the Wiki dataset for three cases: INT8 quantization (w/o RQ), PQ with scalar quantization (SQ) residual codes [12], and PQ with FaTRQ residual codes. Both SQ and FaTRQ reuse the same PQ codebook from the earlier experiment, ensuring consistency across methods.

Unlike SQ-based residuals, which reconstruct record vectors, FaTRQ refines distances directly via enhanced estimation without transferring upper-level codes. As shown, FaTRQ stays much closer to the oracle line (full-precision residual vectors) than 3-bit SQ, with an MSE of 0.0159 vs. 0.258.

Moreover, FaTRQ achieves better storage efficiency. For a 768-dimensional vector, FaTRQ requires only $768/5 + 8 = 162$ bytes (packing five ternary values into a byte, plus four bytes for precomputed values), compared to $768 \times 4/8 = 384$ bytes needed by 4-bit SQ for comparable MSE (0.0134).

D. Refinement Reduction Analysis

We studied how the filtering ratio impacts candidate reranking. For each query, we collected the true top-100 based on PQ distances and examined reranking behavior. Figure 8




Fig. 7. The distortion in squared Euclidean distance estimation relative to the top-100 ground truth results.




Fig. 8. Recall@10 versus refinement ratio under varying proportions of candidates filtered with FaTRQ.

plots recall against the refinement ratio, defined as the SSD I/O normalized by the final top- $k = 10$. Each blue curve corresponds to a filtering rate, where only the top- $X\%$ of the FaTRQ-ranked queue accesses full-precision vectors; the yellow curve is the baseline where the entire list requires SSD I/O. Without FaTRQ, recovering the true top-10 with 99% probability requires scanning up to 70 full-precision vectors out of 100 candidates. With FaTRQ, the same guarantee is reached within 25, reducing the refinement by $2.8\times$.

E. Overhead Analysis

The customized processing unit integrated into the CXL Type-2 device adds only 0.729 mm^2 area and 897 mW power. The FaTRQ distance estimator accounts for 29% of the area and 27% of the power, while the priority queue adds 6% and 8%. Compared to a CXL memory controller [37] with 16 ARM Neoverse V2 cores (2.5 mm^2 , 1.4 W each [38]), the overhead is under 1.8% in area and 4% in power, showing FaTRQ a lightweight addition to the whole memory expander.

The offline overhead of FaTRQ is also minimal. Training the lightweight calibration model and constructing residual codes requires only a single parallel pass per vector, adding about 10 minutes in total, compared to roughly 3 hours to build the CAGRA index for either dataset in our experiments.

VI. CONCLUSION

We introduce FaTRQ, a tier-aware refinement layer that compresses residuals into ternary codes and incrementally refines coarse distances without reconstructing full vectors. FaTRQ targets a key bottleneck in modern ANNS systems: second-pass refinement that fetches full-precision vectors from slow storage. As embedding dimensionality and dataset scale continue to grow, this I/O cost increasingly dominates query latency. By streaming compact residuals from far memory and updating distances on the fly, FaTRQ enables early candidate pruning and cuts unnecessary fetches, boosting throughput by up to $9\times$ compared to GPU [13].

ACKNOWLEDGMENT

This work was supported in part by PRISM and Co-CoSys—centers in JUMP 2.0, an SRC program sponsored by DARPA; by the U.S. DOE DeCoDe Project No. 84245 at PNNL; and by the NSF under Grants No. 2112665, 2112167, 2052809, and 2211386.

REFERENCES

- [1] [Online]. Available: <https://platform.openai.com/docs/guides/embeddings>
- [2] J. Wang, X. Yi, R. Guo, H. Jin, P. Xu, S. Li, X. Wang, X. Guo, C. Li, X. Xu *et al.*, “Milvus: A purpose-built vector data management system,” in *Proceedings of the 2021 International Conference on Management of Data*, 2021, pp. 2614–2627.
- [3] M. Douze, A. Guzhva, C. Deng, J. Johnson, G. Szilvasy, P.-E. Mazaré, M. Lomeli, L. Hosseini, and H. Jégou, “The faiss library,” 2024.
- [4] T. Ge, K. He, Q. Ke, and J. Sun, “Optimized product quantization,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 36, no. 4, pp. 744–755, 2013.
- [5] R. Wang and D. Deng, “Deltapq: lossless product quantization code compression for high dimensional similarity search,” *Proceedings of the VLDB Endowment*, vol. 13, no. 13, pp. 3603–3616, 2020.
- [6] Y. Kalantidis and Y. Avrithis, “Locally optimized product quantization for approximate nearest neighbor search,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2014, pp. 2321–2328.
- [7] B. Wang and C.-C. J. Kuo, “Sbert-wk: A sentence embedding method by dissecting bert-based word models,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2146–2157, 2020.
- [8] J. Jang, H. Choi, H. Bae, S. Lee, M. Kwon, and M. Jung, “Cxl-anns: Software-hardware collaborative memory disaggregation and computation for billion-scale approximate nearest neighbor search,” in *2023 USENIX Annual Technical Conference (USENIX ATC 23)*, 2023, pp. 585–600.
- [9] “cuvs: Vector search and clustering on the gpu.” [Online]. Available: <https://github.com/rapidsai/cuvs>
- [10] J. Ren, M. Zhang, and D. Li, “Hm-ann: Efficient billion-point nearest neighbor search on heterogeneous memory,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 10672–10684, 2020.
- [11] S. Jayaram Subramanya, F. Devvrit, H. V. Simhadri, R. Krishnawamy, and R. Kadekodi, “Diskann: Fast accurate billion-point nearest neighbor search on a single node,” *Advances in neural information processing systems*, vol. 32, 2019.
- [12] V. Karthik, S. Khan, S. Singh, H. V. Simhadri, and J. Vedurada, “Bang: Billion-scale approximate nearest neighbor search using a single gpu,” *arXiv preprint arxiv:2401.11324*, 2024.
- [13] H. Ootomo, A. Naruse, C. Nolet, R. Wang, T. Feher, and Y. Wang, “Cagra: Highly parallel graph construction and approximate nearest neighbor search for gpus,” in *2024 IEEE 40th International Conference on Data Engineering (ICDE)*. IEEE, 2024, pp. 4236–4247.
- [14] B. Tian, H. Liu, Y. Tang, S. Xiao, Z. Duan, X. Liao, X. Zhang, J. Zhu, and Y. Zhang, “Fusionanns: An efficient cpu/gpu cooperative processing architecture for billion-scale approximate nearest neighbor search,” *arXiv preprint arxiv:2409.16576*, 2024.
- [15] B. Tian, H. Liu, Y. Tang, S. Xiao, Z. Duan, X. Liao, H. Jin, X. Zhang, J. Zhu, and Y. Zhang, “Towards high-throughput and low-latency billion-scale vector search via {CPU/GPU} collaborative filtering and re-ranking,” in *23rd USENIX Conference on File and Storage Technologies (FAST 25)*, 2025, pp. 171–185.
- [16] G. Ilharco, M. Wortsman, R. Wightman, C. Gordon, N. Carlini, R. Taori, A. Dave, V. Shankar, H. Namkoong, J. Miller, H. Hajishirzi, A. Farhadi, and L. Schmidt, “Openclip,” Jul. 2021, if you use this software, please cite it as below. [Online]. Available: <https://doi.org/10.5281/zenodo.5143773>
- [17] C. Nolet, “[Keynote] Accelerating vector search on the GPU with RAPIDS RAFT,” NeurIPS 2023, Big ANN Challenge Keynote, 2023, online at <https://neurips.cc/virtual/2023/83767> (accessed Sept 2, 2025).
- [18] H. Jegou, M. Douze, and C. Schmid, “Product quantization for nearest neighbor search,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 33, no. 1, pp. 117–128, 2010.
- [19] S. Liu, H. Lu, and J. Shao, “Improved residual vector quantization for high-dimensional approximate nearest neighbor search,” *arXiv preprint arXiv:1509.05195*, 2015.
- [20] J. Yuan and X. Liu, “Transformed residual quantization for approximate nearest neighbor search,” *arXiv preprint arXiv:1512.06925*, 2015.
- [21] H. Jegou, M. Douze, and C. Schmid, “Product quantization for nearest neighbor search,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 33, no. 1, pp. 117–128, 2010.
- [22] M. Chen, T. Han, C. Liu, S. Liang, K. Yu, L. Dai, Z. Yuan, Y. Wang, L. Zhang, H. Li *et al.*, “Drim-ann: An approximate nearest neighbor search engine based on commercial dram-pims,” *arXiv preprint arXiv:2410.15621*, 2024.
- [23] T. Yu, B. Wu, K. Chen, C. Yan, G. Zhang, and W. Liu, “Hdanns: In-memory hyperdimensional computing for billion-scale approximate nearest neighbour search acceleration,” *IEEE Transactions on Circuits and Systems for Artificial Intelligence*, 2025.
- [24] Z. Zhu, J. Liu, G. Dai, S. Zeng, B. Li, H. Yang, and Y. Wang, “Processing-in-hierarchical-memory architecture for billion-scale approximate nearest neighbor search,” in *2023 60th ACM/IEEE Design Automation Conference (DAC)*. IEEE, 2023, pp. 1–6.
- [25] T. Zhang, F. Ponzina, and T. Rosing, “SpANNs: Optimizing Approximate Nearest Neighbor Search for Sparse Vectors Using Near Memory Processing,” in *31st Asia and South Pacific Design Automation Conference*, 2026.
- [26] Y. A. Malkov and D. A. Yashunin, “Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 42, no. 4, pp. 824–836, 2018.
- [27] B. Munyampirwa, V. Lakshman, and B. Coleman, “Down with the hierarchy: The ‘h’ in hnsw stands for ‘hubs’,” *arXiv preprint arXiv:2412.01940*, 2024.
- [28] “Cohere/wikipedia-22-12.” [Online]. Available: <https://huggingface.co/datasets/Cohere/wikipedia-22-12>
- [29] J. Gao and C. Long, “Rabitq: quantizing high-dimensional vectors with a theoretical error bound for approximate nearest neighbor search,” *Proceedings of the ACM on Management of Data*, vol. 2, no. 3, pp. 1–27, 2024.
- [30] “Cxl near-memory accelerators and memory-expansion controllers.” [Online]. Available: “<https://www.marvell.com/products/cxl.html>”
- [31] “Samsung v-nand ssd 990 pro 2022 data sheet.” [Online]. Available: “https://download.semiconductor.samsung.com/resources/datasheet/Samsung_NVMe_SSD_990_PRO_Datasheet_Rev.1.0.pdf”
- [32] C. Schuhmann, R. Beaumont, R. Vencu, C. Gordon, R. Wightman, M. Cherti, T. Coombes, A. Katta, C. Mullis, M. Wortsman *et al.*, “Laion-5b: An open large-scale dataset for training next generation image-text models,” *Advances in neural information processing systems*, vol. 35, pp. 25278–25294, 2022.
- [33] “Vectordbbench(vdbbench): A benchmark tool for vectordb.” [Online]. Available: <https://github.com/zilliztech/VectorDBBench>
- [34] Y. Kim, W. Yang, and O. Mutlu, “Ramulator: A fast and extensible dram simulator,” *IEEE Computer architecture letters*, vol. 15, no. 1, pp. 45–49, 2015.
- [35] L. T. Clark, V. Vashishtha, L. Shifren, A. Gujja, S. Sinha, B. Cline, C. Ramamurthy, and G. Yeric, “Asap7: A 7-nm finfet predictive process design kit,” *Microelectronics Journal*, vol. 53, pp. 105–115, 2016.
- [36] A. Shafaei, Y. Wang, X. Lin, and M. Pedram, “Fincacti: Architectural analysis and modeling of caches with deeply-scaled finfet devices,” in *2014 IEEE Computer Society Annual Symposium on VLSI*. IEEE, 2014, pp. 290–295.
- [37] “Structera a 2504 memory-expansion controller.” [Online]. Available: “<https://www.marvell.com/content/dam/marvell/en/public-collateral/assets/marvell-structera-a-2504-near-memory-accelerator-product-brief.pdf>”
- [38] M. Bruce, “Arm neoverse v2 platform: Leadership performance and power efficiency for next-generation cloud computing, ml and hpc workloads,” in *2023 IEEE Hot Chips 35 Symposium (HCS)*, 2023, pp. 1–25.