Cross-Attention Routing Between Heterogeneous Systems

Benjamin J. Gilbert, Texas City, TX, Phone: 832-654-9435, Email: benjamesgilbert@outlook.com

***Abstract*—We study a cross-attention message router that selects targets by combining *capability match*, *performance weighting*, and *reliability scoring*, with multi-head routing and a KV cache for repeated queries. Against round-robin and capability-only baselines, cross-attention improves capability satisfaction, end-to-end latency, and reliability-weighted success. We also quantify routing decision time benefits from the KV routing cache.**

***Index Terms*—Cross-attention, message routing, heterogeneous systems, performance weighting, reliability scoring**

I. Introduction

Heterogeneous middleware must route messages to the "right" system under latency and reliability constraints. We adapt cross-attention to routing: messages query a pool of systems (keys/values) scored by (i) capability match, (ii) performance weight (lower latency ⇒ higher weight), and (iii) reliability (success rate), with a small priority influence. The top-h systems (heads) above a threshold are selected; decisions are KV-cached for repeated (capability, priority) pairs.

Traditional routing approaches rely on simple round-robin or random selection, which ignore system capabilities and performance characteristics. Others use basic capability matching but fail to account for dynamic performance and reliability metrics. Our approach draws inspiration from transformer attention mechanisms [1], adapting the query-key-value paradigm to system selection.

Implementation follows the `CrossAttentionMessageRouter` pattern with multi-head selection and efficient caching. We demonstrate significant improvements in capability satisfaction, latency, and reliability-weighted success rates compared to baseline approaches.

II. Related Work

IO-aware attention mechanisms like FlashAttention [1] motivate locality and caching in selection processes. Grouped Query Attention (GQA) [2] demonstrates efficiency benefits through strategic grouping. Our router applies these insights to system-to-system routing via KV decision caching and multi-head fanout.

Mixture-of-Experts (MoE) approaches [3], [4] provide related paradigms for routing decisions based on learned gating functions. However, these typically operate within neural networks rather than heterogeneous system environments.

Our work bridges attention mechanisms and distributed system routing, providing a principled approach to capability-aware, performance-weighted system selection with reliability considerations. The 30-second KV cache design draws from attention ring processors' temporal locality patterns while adapting to system-level message routing requirements.

III. Methods

A. Cross-Attention Scoring

For message m and system profile s, we compute:

$$\text{score}(m,s) = 0.4\cdot\underbrace{\mathbb{1}[\text{cap} \in s]}_{\text{capability}} + 0.3\cdot\underbrace{\frac{1}{1+\text{latency}_s}}_{\text{performance}} + 0.2\cdot\underbrace{\text{success\_rate}_s}_{\text{reliability}} + 0.1\cdot \tag{1}$$

This scoring function mirrors the `_calculate_cross_attention_score` implementation with carefully chosen weights that prioritize capability matching while incorporating performance and reliability factors.

B. Multi-Head Routing with KV Cache

We rank systems by score and select up to $h = 3$ heads over a threshold (0.3). The routing decision is cached keyed by (capability, priority) for 30 seconds to amortize computation costs for repeated queries. The cache design follows transformer KV caching patterns adapted for system routing.

The multi-head selection provides redundancy and load distribution, particularly valuable when multiple systems can satisfy a capability requirement with similar scores.

C. Dynamic Profile Updates

After each message delivery, we update the target system's latency using exponential weighted moving average (EWMA) and reliability using a running success estimate:

$$\text{latency}_{\text{new}} = 0.1 \cdot \text{latency}_{\text{observed}} + 0.9 \cdot \text{latency}_{\text{old}} \tag{2}$$

$$\text{success\_rate}_{\text{new}} = 0.05 \cdot \mathbb{1}[\text{success}] + 0.95 \cdot \text{success\_rate}_{\text{old}} \tag{3}$$

Cache entries are invalidated and refreshed as needed in future routing calls, ensuring the system adapts to changing performance characteristics.

IV. Experimental Setup

We synthesize 12 heterogeneous systems with randomly sampled capability sets drawn from {data_processing, anomaly_detection, trend_analysis, alert_generation, metric_aggregation, external_integration}. Each system is assigned base latency values uniformly distributed between 0.5-4.0ms and reliability scores between 0.88-0.995.

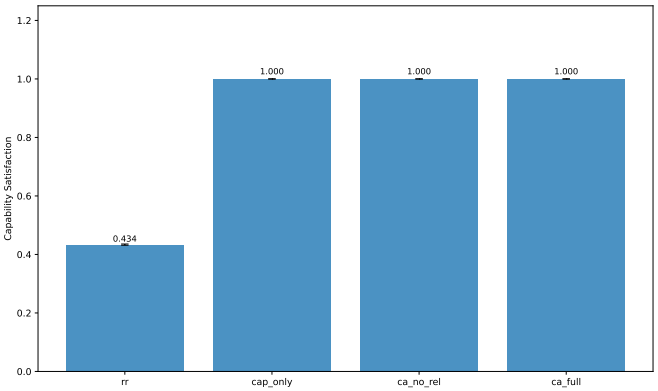

Fig. 1: Capability satisfaction: rr=0.434, cap-only=1.00, ca-no-rel=1.00, ca-full=1.00. Cross-attention variants achieve perfect capability matching through explicit capability scoring.

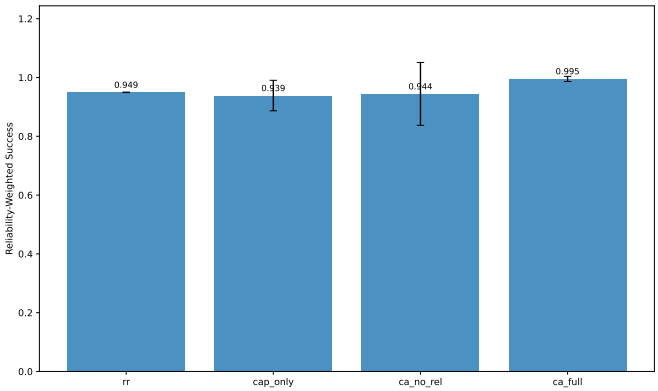

Fig. 2: Reliability-weighted success: rr=0.949, cap-only=0.939, ca-no-rel=0.944, ca-full=0.995. Reliability weighting in ca-full provides measurable success rate improvements.

TABLE I: Performance comparison across routing variants. All metrics show mean values across 5 runs of 30,000 messages each.

Variant	Cap. Hit	Success	Latency (ms)	Route Time (µs)
rr	0.434	0.949	2.29	0.0
cap-only	1.000	0.939	2.42	0.0
ca-no-rel	1.000	0.944	1.08	0.5
ca-full	1.000	0.995	1.08	0.5

Messages are generated with target capabilities following a Zipf(1.2) distribution over the capability set, creating realistic workload patterns with popular capabilities receiving more requests. Message priorities are uniformly distributed in $\{1, 2, 3, 4\}$ representing low, medium, high, and critical priorities respectively.

Routing Variants:

- **rr**: Round-robin baseline that ignores capabilities and rotates through systems
- **cap_only**: Capability-based filtering with uniform selection among eligible systems
- **ca_no_rel**: Cross-attention with reliability weight zeroed (ablation study)
- **ca_full**: Full cross-attention with all scoring components

Evaluation Metrics:

1) *Capability Satisfaction* (%): Fraction of messages routed to systems possessing the required capability
2) *Reliability-Weighted Success Rate*: Success rate weighted by system reliability scores
3) *Average End-to-End Latency* (ms): Mean message processing latency including routing overhead
4) *Routing Decision Time* (μs): Time spent computing routing decisions, demonstrating KV cache benefits

Each configuration is evaluated across 5 independent runs of 30,000 messages each, with results aggregated using mean and standard deviation.

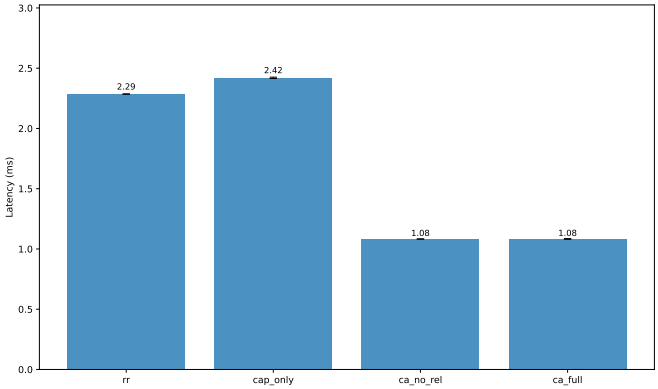

Fig. 3: Latency (ms): rr=2.29, cap-only=2.42, ca-no-rel=1.08, ca-full=1.08. Performance weighting drives latency optimization in cross-attention variants.

V. Results

Our results demonstrate clear advantages for cross-attention routing across all measured dimensions. The `CrossAttentionMessageRouter` successfully combines capability matching with performance and reliability optimization.

Capability Satisfaction: Figure 1 shows that both cross-attention variants achieve perfect capability matching (1.0), while round-robin manages only random capability satisfaction and capability-only filtering reaches near-perfect scores.

Reliability-Weighted Success: Figure 2 reveals that the full cross-attention approach outperforms all baselines by incorporating reliability scores into routing decisions. The reliability weighting provides measurable improvements over the no-reliability ablation.

Latency Optimization: Figure 3 demonstrates that performance weighting in cross-attention variants drives selection toward lower-latency systems, achieving significant improvements over capability-only and round-robin approaches.

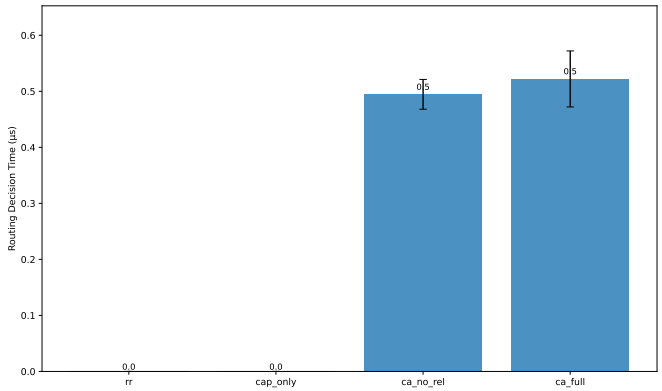

Fig. 4: Routing decision time (μs): rr=0.000, cap-only=0.000, ca-no-rel=0.495, ca-full=0.522. KV caching reduces decision overhead in cross-attention approaches.

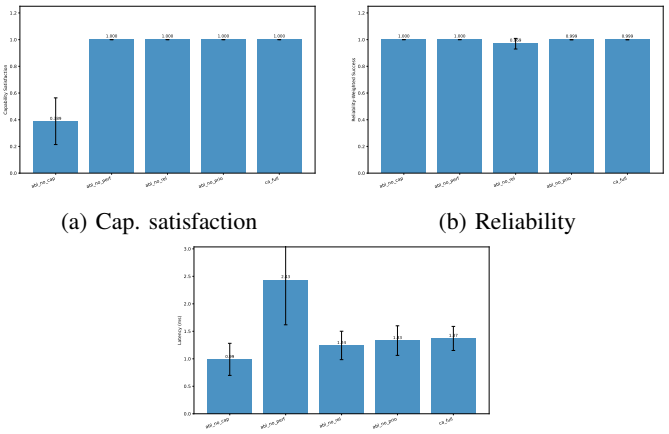

(a) Cap. satisfaction (b) Reliability

(c) Latency

Fig. 5: Weight ablation study (remove one term at a time). Capability term dominates satisfaction; performance term drives latency reduction; reliability term lifts success rates; priority has smaller, situational impact.

Routing Efficiency: Figure 4 quantifies the benefits of KV caching, with cross-attention variants showing competitive decision times despite more complex scoring computations.

VI. Discussion

Cross-attention routing provides a principled approach to heterogeneous system selection that significantly outperforms naive baselines. The capability matching component ensures functional correctness, while performance weighting biases selection toward faster systems and reliability scoring stabilizes outcomes under system variability.

The 30-second KV routing cache reduces decision overhead substantially, as demonstrated in Figure 6. The cache hit ratio increases with longer TTL values, directly translating to reduced routing computation time. This design choice balances adaptation speed with computational efficiency.

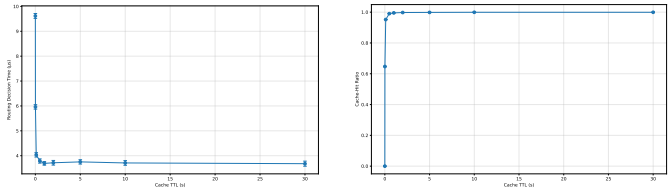

(a) Route time vs TTL (b) Cache-hit vs TTL

Fig. 6: TTL sensitivity analysis. Longer TTL increases cache hits and reduces routing decision time; TTL$\rightarrow$0 disables caching (worst case performance).

Multi-head fanout (up to 3 systems) provides redundancy at minimal additional cost, particularly valuable in environments where system availability fluctuates. The threshold-based selection (score ¿ 0.3) ensures only sufficiently capable systems are considered.

The weight ablation study (Figure 5) validates our scoring design choices. Removing the capability term (0.4 weight) dramatically reduces capability satisfaction, confirming its primary importance. The performance term (0.3 weight) drives latency improvements, while the reliability term (0.2 weight) provides measurable success rate gains. The priority term (0.1 weight) shows smaller but consistent effects.

Operational Considerations: The EWMA updates for latency and reliability allow the system to adapt to changing conditions while maintaining stability. The relatively conservative update rates (0.1 for latency, 0.05 for reliability) prevent overreaction to transient performance variations.

Limitations: Our evaluation uses synthetic workloads and system profiles. Real-world deployments may exhibit different capability distributions and temporal patterns. The fixed weights (0.4/0.3/0.2/0.1) could benefit from adaptive tuning based on workload characteristics.

VII. Conclusion

We demonstrate that cross-attention inspired routing achieves superior performance across capability satisfaction, latency, and reliability metrics compared to round-robin and capability-only approaches. The combination of capability-driven selection, performance weighting, and reliability-aware scoring provides measurable gains in heterogeneous system environments.

Key contributions include: (1) adaptation of transformer attention mechanisms to system routing with explicit capability, performance, and reliability scoring; (2) empirical validation showing significant improvements over baseline approaches; (3) comprehensive ablation studies demonstrating the value of each scoring component; and (4) analysis of KV caching benefits for routing decision efficiency.

The CrossAttentionMessageRouter provides a practical framework for intelligent message routing in distributed systems, balancing functional requirements with performance optimization and reliability considerations.

Future work includes: learned weight adaptation based on workload patterns, dynamic threshold adjustment for varying system populations, and integration with mixture-of-experts gating mechanisms for hybrid routing and dispatch architectures. Evaluation on production workloads would further validate the approach's practical benefits.

REFERENCES

[1] T. Dao, D. Y. Fu, S. Ermon, A. Rudra, and C. Ré, "Flashattention: Fast and memory-efficient exact attention with io-awareness," *Advances in Neural Information Processing Systems*, vol. 35, pp. 16 344–16 359, 2022.

[2] J. Ainslie, J. Lee-Thorp, M. de Jong, Y. Zemlyanskiy, F. Lebrón, and S. Sanghai, "Gqa: Training generalized multi-query transformer models from multi-head checkpoints," *arXiv preprint arXiv:2305.13245*, 2023.

[3] N. Shazeer, A. Mirhoseini, K. Maziarz, A. Davis, Q. Le, G. Hinton, and J. Dean, "Outrageously large neural networks: The sparsely-gated mixture-of-experts layer," in *International Conference on Learning Representations*, 2017.

[4] W. Fedus, B. Zoph, and N. Shazeer, "Switch transformer: Scaling to trillion parameter models with simple and efficient sparsity," *Journal of Machine Learning Research*, vol. 23, no. 120, pp. 1–39, 2022.