DFly2: Turbocharging Long-Context Speculative Decoding by Sharing the Target's Sparsity

Hong Liu, Rui Cen, Junhan Shi, Lin Niu, Runzhi Fan, Zhaojing Zhou, Haodi Lei, Jiebin Zhang, Saiyong Yang, Junqiang Zheng, Zhenghai Xue, Guanghua Yu [*], Jianchen Zhu

Tencent Inc.
[*]Corresponding author

https://github.com/Tencent/AngelSpec
lucayu@tencent.com, irisxliu@tencent.com

arXiv:2611.34776v1 [cs.CL] 25 Nov 2026

Abstract

To avoid the quadratic cost of dense attention over long contexts, models such as DeepSeek-V4 and Hy4 have turned natively sparse. Speculative decoding adds a further, lossless speedup with fewer target passes, and matters most for the long contexts these models serve, where decoding is slowest. Block-parallel drafters, however, still work in a paradigm built for dense targets, keeping a context cache of their own over the whole prompt or a sliding window. The full one takes a larger share of the KV per token the more the target compresses its own, and either costs more than memory, in prefill time, prefix-cache hits and decode concurrency. Drafters that read the target's KV cache avoid one of their own, but work on this line settles only which cache, not which positions, and still reads every position or a recent window. Neither rule is made for the query at hand, and computing a selection per query in the drafter would bring back a cost that grows with the context. Yet a natively sparse target already picks positions for each query, over a cache it keeps anyway. Both of the drafter's choices are therefore already made, and paid for, by the target: its cache and its selection. We measure how far one selection carries across a block: made a step early, it still covers close to nine tenths of what each later token selects on DS-V4.1-Flash. We therefore present **DFly2, a block-parallel drafter that shares the target's sparsity**. Each draft layer reads the KV its paired target layer cached (**KV sharing**), at the positions that layer's indexer selected, a fixed number from anywhere in the context (**Index sharing**). DFly2 runs no scoring pass and keeps no cache of its own, so none of those costs apply. We evaluate on three targets whose sparse attention and architecture differ, DS-V4.1-Flash, GLM-5.3-Flash and Hy4 flash, over general-purpose, long-context and agentic benchmarks at two temperatures. Across all of them, DFly2 accepts more tokens per step than DFlash and DSpark trained on the same recipe. Its margin over the stronger baseline widens from **3.9%–10.5%** on general workloads to **10.0%–31.5%** on long-context and agentic workloads of up to 128K. Serving DS-V4.1-Flash at high concurrency and matched per-user decode speed, it delivers **62%** and **81%** more throughput than that baseline on long-context and agentic sets at 64K.

1 Introduction

Agents and multi-turn chat have pushed serving to long contexts, where the cost of dense attention grows quadratically as the context lengthens. Recent long-context models such as DeepSeek-V4, GLM-5 and Hy4 have answered by turning natively sparse (Yuan et al., 2025; DeepSeek-AI, 2025, 2026; GLM-5 Team, 2026; Tencent Hy Team, 2026a). Each sparse layer attends only to a fixed budget of positions, which a trained indexer picks for the current query. They also store less, compressing the KV cache they keep for the context (DeepSeek-

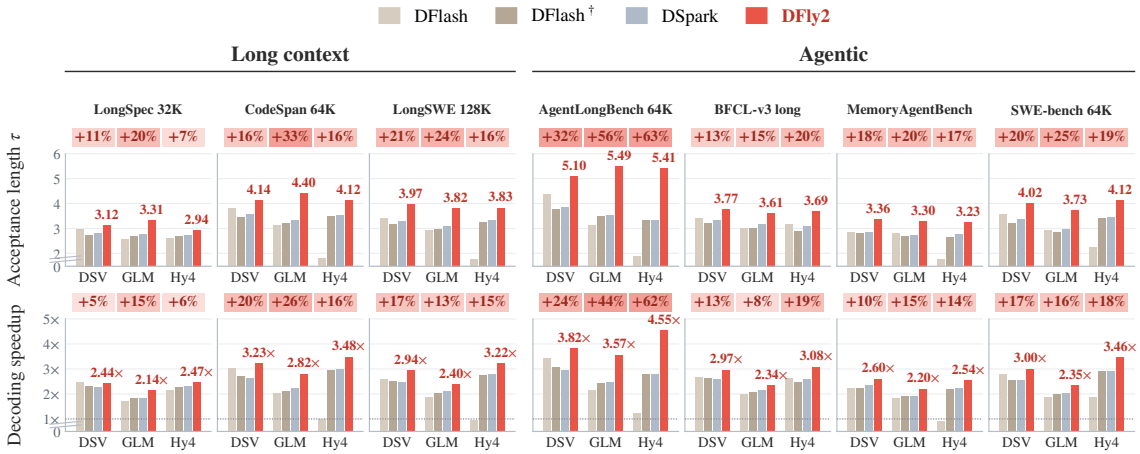

Figure 1: Accepted length τ and decoding speedup on DS-V4.1-Flash, GLM-5.3-Flash and Hy4 flash, greedy at a concurrency of one. Chips give DFly2's gain over the better of DFlash† and DSpark.

AI, 2025, 2026; GLM-5 Team, 2026; Tencent Hy Team, 2026a). At a million tokens, DeepSeek-V4-Flash needs 10% of the per-token FLOPs and 7% of the KV cache of DeepSeek-V3.2 (DeepSeek-AI, 2026).

Speculative decoding, a mainstream lossless accelerator, speeds up decoding further with fewer target passes, and serving needs it most in agentic and multi-turn workloads, whose long contexts slow decoding down (Leviathan et al., 2023; Chen et al., 2023; Sadhukhan et al., 2025; Cheng et al., 2026). Long-context speculative decoding thus increasingly runs on natively sparse targets: DeepSeek-V3.2, DeepSeek-V4, GLM-5 and Hy4 all ship with a drafter (DeepSeek-AI, 2026; Cheng et al., 2026; GLM-5 Team, 2026; Tencent Hy Team, 2026a). Sparsity has made each decode step of such a target cheaper, but drafters still work in a paradigm designed for dense targets. We focus on block-parallel drafters, which make two choices about the target's context: where it comes from, meaning which cache and which positions, and how it is injected (Chen et al., 2026). The mainstream answer takes the context from a cache of the drafter's own, over positions picked by a rule that ignores the current query (Chen et al., 2026; Zhang et al., 2026; Lin et al., 2026; Liu et al., 2026a; Cheng et al., 2026; Rheinboldt et al., 2026; Huang et al., 2026). That cache holds keys and values the drafter projects from the target's hidden states and injects into every draft layer. DFlash caches every position, and this cache weighs more the more the target saves (Section 3.1). It adds a fifth to the KV per token on Hy4 flash and three quarters on DeepSeek-V4.1-Flash (DS-V4.1-Flash). DSpark bounds that cost with a sliding window, which on DeepSeek-V4 spans the last 128 tokens of the context (Cheng et al., 2026).

Full or windowed, a cache of the drafter's own costs more than the memory it takes. The drafter has to run on the prefill instances under prefill-decode disaggregation and build the cache over the prompt, which delays the first token (Sections 3.1 and 5.2). A full cache also leaves the target fewer and shorter prefix-cache hits, and removing it lets DS-V4.1-Flash hold 1.73× as many requests at 32K tokens (Section 3.1).

Drafters that read the target's KV cache instead avoid a growing cache of their own, a line that goes back to GliDe (Du et al., 2024) and runs through LongSpec, H-Spec and LongSpark (Yang et al., 2025; Jiang et al., 2026; He et al., 2026). They settle which cache to read, but not which positions. Like DFlash and DSpark, these drafters still read every position or a recent window, by rules that ignore the current query because the dense targets they were developed on offered no selection to follow (Section 6). Attention over all positions also grows with the context, while a fixed number of positions costs the same at every length (Section 3.1). For drafters with a cache of their own, the better rule even depends on the target. On long-context and agentic workloads, a full cache beats a window on DS-V4.1-Flash, loses to one on GLM-5.3-Flash, and collapses on Hy4 flash once prompts pass the 32K the drafter was trained at (Section 5.2). A rule carried over from dense targets thus needs to be adapted to each target. Each of these rules is fixed when the drafter is designed, not made for the query at hand, and scoring the context instead would bring back a cost that grows with it. The cache a drafter reads can be the target's own, so could the selection it follows come from the target too?

A natively sparse target already makes that selection at every step: its trained indexers pick a fixed budget of positions for the current query, from a KV cache the target already keeps. This is the selection that drafting has wanted and given up: top-k attention "can achieve very high acceptance rate", yet "it is not a practical

draft option because of its prohibitively high KV selection cost" (Sadhukhan et al., 2025). Verification-guided methods recover it by extracting it from a dense target's attention (Zhao et al., 2026; Yue et al., 2026; Cha et al., 2025), a step a natively sparse target does not need. Our key insight is simple: a natively sparse target has already settled where a drafter's context should come from, with its cache and its selection, and paid for both. One selection also holds across a block: made one step early, for the token just before it, it serves every slot. On DS-V4.1-Flash it covers close to nine tenths of what each later token would select, against less than a fifth for a window of the same size (Section 3.3).

We present **DFly2, a block-parallel drafter that shares the target's sparsity**: the compact KV cache it maintains and the positions its sparse layers attend to. Each draft layer reads the KV that a paired target layer has cached (**KV sharing**), at the positions that layer's indexer selected, a fixed number drawn from anywhere in the context (**Index sharing**). DFly2 thus runs no scoring pass and keeps no cache of its own, so it needs no place on the prefill instances, inherits every prefix-cache hit of the target, and costs no decode concurrency (Section 4). The selection was made for the target's queries, yet the drafter is a far smaller model with queries of its own. Each draft layer starts from its paired layer's weights, including the query latent that layer shares with its indexer, so draft queries begin in the space that produced the selection. With the target's last-layer hidden state, which no cache holds, gated into the block, DFly2 reuses four things the target has already computed: its KV, its selection, its layer weights and that state. An intra-block correction then rereads the shared context slot by slot.

We instantiate DFly2 on three targets that differ in their variant of DeepSeek Sparse Attention (DSA) (DeepSeek-AI, 2025) and in architecture, taking each target's selection in the form it gives (Section 4). DS-V4.1-Flash keeps one cache shared across its layers as YOCO does (Sun et al., 2024b), GLM-5.3-Flash interleaves DSA layers with linear attention ones, and Hy4 flash compresses the index of its sparse layers in two stages (Section 2). On all three, DFly2 accepts more tokens per step than DFlash and DSpark trained on the same recipe, on every benchmark and temperature (Section 5.2). Its margin over the stronger baseline grows from 3.9%–10.5% on general workloads to 10.0%–31.5% on long-context and agentic ones. On agentic workloads, whose prompts mostly run longer, its average stays within 4% of its general average, while DSpark's falls by 11% to 21%. The target's selection reaches anywhere in the context yet stays a fixed size, so DFly2 inherits its length generalization. What DFly2 reads is not tuned to a target but taken from it, which is why the same design holds on all three architectures. Accepted length turns into speed, and with nothing added at prefill, DFly2 averages a higher end-to-end speedup than the baselines on every target and workload family (Section 5.3).

Our contributions are threefold:

- **A drafter can take the target's selection as it is.** A context of the drafter's own costs more than memory, reaching prefill, prefix caching and decode concurrency. Our insight is that the selection it needs is already computed, and paid for, by the target. We measure how far one selection carries: made one step early, it still covers close to nine tenths of what each later token in the block would select on DS-V4.1-Flash.
- **DFly2 shares the target's selection alongside its cache.** DFly2 takes both: each draft layer reads its paired target layer's KV cache, at the positions that layer's indexer selected, and starts from that layer's weights. DFly2 therefore keeps no cache, runs no scoring pass, and reads a fixed number of positions from anywhere in the context, so the costs above do not apply. We add two pieces: the last-layer state that no cache holds, and an intra-block correction that rereads the shared context.
- **Gains that grow with the context, on all three targets.** Against baselines trained on the same recipe, DFly2's margin in accepted length over the stronger one grows from 3.9%–10.5% on general workloads to 10.0%–31.5% on long-context and agentic ones. It generalizes past its training length: on agentic workloads its average accepted length stays within 4% of its general average. Serving at high concurrency on DS-V4.1-Flash, at matched per-user decode speed, it delivers 28% more throughput on short text and 62% and 81% more on a long-context and an agentic set at 64K.

2 Background

Speculative decoding. At each step a drafter proposes γ tokens and the target verifies them in one forward pass, accepting the longest prefix consistent with its distribution and appending one token of its own, the bonus token (Leviathan et al., 2023; Chen et al., 2023). With τ the mean accepted length, which counts the bonus

Table 1: How a drafter receives the target's context. Only reading the target's cache leaves the drafter without a context KV of its own.

Route	Context enters as	Drafter's own context KV
Input conditioning (MTP, EAGLE)	target hidden states at the input	every position
KV injection (DFlash, DSpark)	H_{ctx} in every layer's K/V	every position
Reading the target's cache (GliDe)	the target's stored K/V, read in place	none

token and so satisfies $1 \le \tau \le \gamma + 1$, the latency per generated token is

$$T_{\text{tok}} \;=\; \frac{T_{\text{draft}} + T_{\text{verify}}}{\tau}. \tag{1}$$

Drafting time in turn splits into a context-independent part T_0, spent in projections, feed-forward layers and the output head, and attention over the m context positions whose KV the drafter reads,

$$T_{\text{draft}} \;=\; T_0 + T_{\text{attn}}(m). \tag{2}$$

Speedup thus comes from raising τ while keeping T_{draft} small, and T_{draft} grows with the context only through m.

Block-parallel drafters. A block-parallel drafter fills a whole block in one forward pass (Chen et al., 2026). Its input is an anchor token followed by γ mask tokens, the draft slots, which attend to one another and to the context.[1] Whatever its architecture, a drafter is conditioned on the target's context, which reaches it in one of three ways (Table 1). Autoregressive drafters such as MTP (Gloeckle et al., 2024) and EAGLE (Li et al., 2025) take target hidden states at their input and keep a KV cache of their own over the whole context. The other two routes put the context into the keys and values of every draft layer i,

$$K_i = \big[\, X_i \;;\; W_i^K H_{\text{blk}} \,\big], \qquad V_i = \big[\, Y_i \;;\; W_i^V H_{\text{blk}} \,\big], \tag{3}$$

where H_{blk} holds the block's hidden states and $[\,\cdot\,;\,\cdot\,]$ stacks along the sequence. KV injection sets $(X_i, Y_i) = (W_i^K H_{\text{ctx}}, W_i^V H_{\text{ctx}})$, with H_{ctx} fused and projected from the hidden states of several target layers (Chen et al., 2026; Cheng et al., 2026). Reading the target's cache instead builds them from the latents the target has already stored for a paired layer (Du et al., 2024). The first two routes thus leave the drafter with KV of its own for every context position, and the third adds none.

Natively sparse targets. The three targets we study are natively sparse: each is trained with a variant of DeepSeek Sparse Attention (DSA) (DeepSeek-AI, 2025), which recent long-context models have widely adopted (DeepSeek-AI, 2026; GLM-5 Team, 2026). In each sparse attention layer, a lightweight indexer gives every preceding token s an index score $I_{t,s}$ for the query token $\mathbf{h}_t$, and the layer attends only to the key-value latents $\mathbf{c}_s$ that score highest to produce the output $\mathbf{u}_t$,

$$\mathbf{u}_t = \text{Attn}\big(\mathbf{h}_t,\ \{\mathbf{c}_s \mid s \in \mathcal{S}_t\}\big), \qquad \mathcal{S}_t = \big\{\, s \mid I_{t,s} \in \text{Top-k}(I_{t,:}) \,\big\}. \tag{4}$$

The selection $\mathcal{S}_t$ is made afresh for every query and has a fixed size, while the scoring pass still runs over the whole context (Fu et al., 2026). These models also shrink what they store: most layers keep only a sliding window or a recurrent state, and the three differ in the remaining layers. DS-V4.1-Flash (DeepSeek-AI, 2026) stores a compressed latent in four of its 40 layers and shares it across the others as YOCO does (Sun et al., 2024b), GLM-5.3-Flash (GLM-5 Team, 2026) interleaves 11 DSA layers with 34 linear attention layers, and Hy4 flash (Tencent Hy Team, 2026b) compresses the index of its sparse layers in two stages. The KV that grows with the context is therefore small, and both this compact layout and the selection $\mathcal{S}_t$ are part of the target itself.

[1] We call positions rows, as in the target's attention matrix. The anchor, the previous round's bonus token, sits at row A, and the target generated it from its query at row $A-1$. Draft slot k predicts the token at row $A+k$ and so stands in for the target's query at row $A+k-1$.

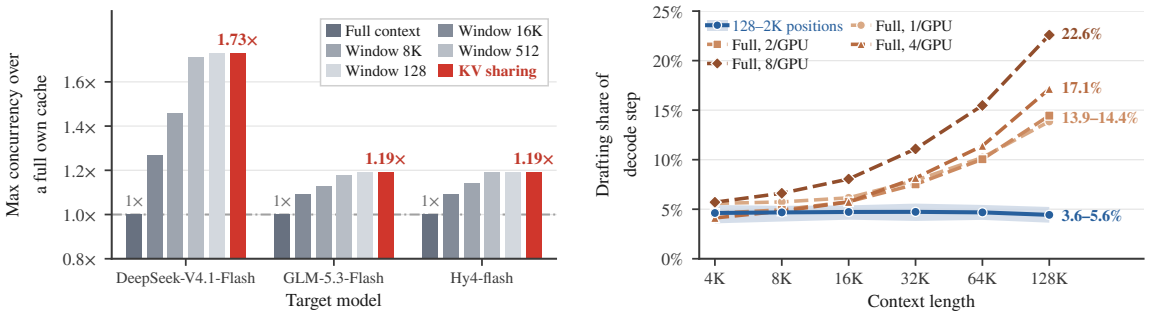

(a) Decode concurrency at 32K **(b) Drafting cost**

Figure 2: What a drafter's own cache costs at decode. (a) Concurrency at 32K tokens, with a full own cache as 1×, estimated from model configurations including fixed per-request state. (b) DS-V4.1-Flash, timed per step by the server, with the band measured as a window.

3 Observations

A drafter makes two choices about its context, whether to keep its own KV cache and which positions to read, and both point to the target. Its own cache costs it at prefill, in prefix caching and at decode, and sharing the target's KV avoids all three (Section 3.1). At long contexts it should read only a bounded number of positions, and the target already selects a fixed number of them for its own attention (Section 3.2). Made one step before the block, this selection still covers most of what later tokens would select (Section 3.3).

3.1 Should the drafter keep its own KV cache?

Owning a cache of its own costs a drafter in three places, and attention over the context adds a fourth cost at every step.

Prefill. At prefill, a cache of its own adds compute but mainly constrains where the drafter runs. Under prefill-decode disaggregation the drafter must also run on the prefill instances, where it holds its weights, projects an injection at every prompt position and transfers the resulting cache to the decode instance. KV sharing adds nothing at prefill, because the latents the drafter reads are ones the target has already written.

Prefix caching. A cache of its own also leaves the target fewer and shorter prefix-cache hits. The drafter's own cache is projected from the target's hidden states, which prefix caching does not keep. It must therefore either be kept in the prefix cache or be rebuilt by recomputing the target. If kept, it takes over two fifths of each cached prefix's memory on DS-V4.1-Flash (Table 2), so the pool holds fewer prefixes for multi-turn and agentic workloads to reuse. Rebuilding it shortens the target's own hit, since the target must recompute the hidden states the hit would have skipped. After each hit, vLLM (Kwon et al., 2023) recomputes the last 128 tokens to restore the target's sliding windows on DS-V4.1-Flash, so a drafter window avoids both only if it is no longer than those 128 tokens. Under KV sharing the drafter has no cache of its own to keep or rebuild, so every target hit is also a drafter hit.

Table 2: A drafter's own cache costs most where the target has saved most. Per-token KV of all caches that grow with the context, from model configurations (bf16), for a three-layer drafter.

Target	Target KV	Drafter's own cache	Increase
DS-V4.1-Flash	4.4 KiB	3.4 KiB	**1.76**×
GLM-5.3-Flash	11.7 KiB	3.0 KiB	1.26×
Hy4 flash	17.2 KiB	3.4 KiB	1.20×

Table 3: KV sharing removes the drafter's own cache without limiting its reach. DS-V4.1-Flash, $w = 128$ or 512, concurrency estimated as in Figure 2 a. †Except at $w = 128$.

	Drafter's own cache		
	Full	**Window w**	**KV sharing**
Drafter needed at prefill	yes, for all positions	yes, for the last w	**no**
Target's prefix-cache hits	fewer and shorter	can be cut short†	**unaffected**
Decode concurrency at 32K	1.00×	1.71–1.73×	**1.73**×
Reach at 32K	**whole context**	last w (0.4–1.6%)	**whole context**

Decode. A drafter that keeps its own cache for the whole context, as DFlash does, lowers decode concurrency. Every in-flight request's KV stays resident during decode, so a pool of M bytes holds at most $M/(bL + f)$ requests of length L. Here b is the per-token KV that grows with the context, including the drafter's own cache if it keeps one, and f the fixed per-request state. The concurrency lost is largest on the targets that have saved the most, because their layer design shrinks the target's KV but not the drafter's own cache. Even in the target's compact latent layout, three draft layers add three quarters to b on DS-V4.1-Flash (Table 2). Removing that addition, as KV sharing does, lets the same pool hold 1.73× as many requests at 32K tokens (Figure 2 a). A sliding-window drafter (Cheng et al., 2026) trades reach for concurrency: the closer it comes to that ratio, the less of the context its window covers (Figure 2 a). KV sharing attains that ratio without a window, so the drafter can read anywhere in the context.

Drafting. The last cost comes from attention over the context rather than from owning a cache. By Equation (2), drafting time grows with the context only through $T_{\text{attn}}(m)$, whose compute and memory traffic scale with the number m of positions read. A fixed budget of positions, as a window or a top-k selection, costs 4% to 6% of the decode step at every context length and load. The full cache costs as much at 4K tokens but climbs to 14% to 23% at 128K, faster under heavier load (Figure 2 b). This cost should follow the number of positions, wherever they lie and whichever cache holds them. A drafter that shares the target's KV but reads all of it then pays that rising cost too. At long contexts a drafter should therefore read a bounded number of positions, and the most useful ones cost as little as the most recent.

A sliding-window drafter shrinks its own cache but still has to run on the prefill side (Table 3). It recovers most of the concurrency and keeps prefix caching intact only by giving up reach. Trained with a 512- or 128-token window, DFlash loses 3% to 12% of its full-context accepted length on long contexts, the most at the 128 tokens that prefix caching requires. Sharing the target's KV removes the drafter's own cache and, by design, all three costs of owning it, while keeping the whole context within reach. With the read bounded as well, the drafter's cost should no longer grow with the context, and the main question left is which positions to read without losing that reach.

3.2 Which positions should the drafter read?

A bounded read can pick its positions by position, which yields a window of the most recent ones (Cheng et al., 2026) and costs reach (Section 3.1). It can instead pick by query, reaching the whole context at the price of a scoring pass over all positions. That pass grows with the context like a full read, giving back what the bound saved (Sadhukhan et al., 2025).

A natively sparse target already pays for this scoring pass. At every step its indexers select a fixed budget of positions for its sparse attention layers, 512 to 2,048 in released models.[2] A drafter that shares the target's KV can share the selection as well, reading by query at the cost of a window and with no scoring of its own (Figure 2 b). The drafter is trained to predict what the target would generate, so we expect it to need the same context the target attends to.

NSA (Yuan et al., 2025) lends indirect support: a selection trained with the model outperforms one computed afterwards, and the target's indexers were trained with the target. Following the target's selection still rests on

[2] 512 and 1,024 on DeepSeek-V4-Flash and V4-Pro (DeepSeek-AI, 2026), 2,048 on DeepSeek-V3.2 (DeepSeek-AI, 2025) and on GLM-5.3-Flash.

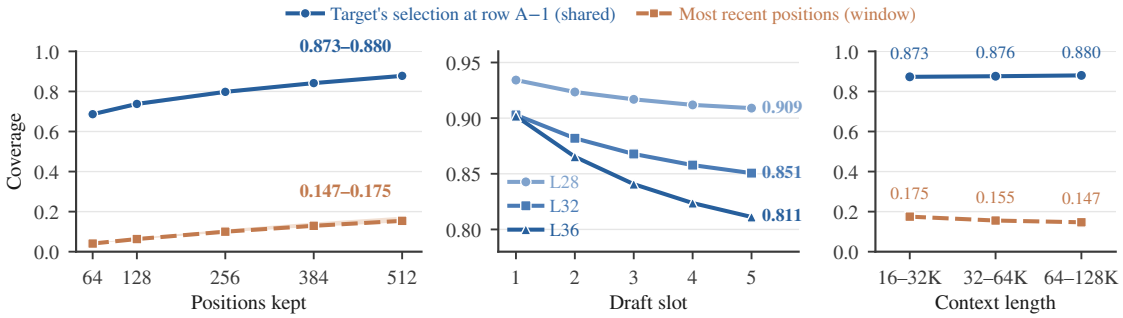

(a) By budget **(b) Along the block, 64K–128K** **(c) By context length**

Figure 3: Top-k mass coverage on DS-V4.1-Flash. (a) Mean over slots and layers. Bands span context lengths. (b) Budget 512, own scale. The window stays at 0.13–0.17. (c) Budget 512, averaged as in (a).

two assumptions. The drafter is a different model, so learning to use the selection is a matter of training that Section 5 tests end to end. The selection, made one step before the block, must still cover what the rows in the block would select, which Section 3.3 measures on the target alone, with no drafter involved.

3.3 Does the target's selection hold across the block?

Question. The second assumption is that a selection made for one query serves the queries that follow. The target selects one query at a time, so a drafter starting a block has only the selection made for the token just before the anchor. The drafter applies that selection to every token in the block. Each later token, however, has a query of its own, which depends on tokens the target has not yet processed and may favor other positions.

Measurement. In the notation of Section 2, the shared selection is $\mathcal{S}_{A-1}$, and draft slot k stands in for row $A+k-1$. We run the target on the actual continuation to obtain each row's own top-512, $\mathcal{S}_t$ of Equation (4). With p_t the index-score distribution of row t, the top-k mass coverage

$$\mathrm{Cov}(t) \;=\; \sum_{s\in\mathcal{S}_{A-1}} p_t(s) \;\Big/\; \sum_{s\in\mathcal{S}_t} p_t(s) \tag{5}$$

is the share of what row t would select that the shared selection already holds. Its denominator is the most any 512 positions can hold, so $\mathrm{Cov}(t) \le 1$. We measure it on DS-V4.1-Flash over 158,638 blocks of long-context evaluation data.

Findings. The selection largely holds. Averaged over slots and layers, the coverage of Equation (5) reaches close to nine tenths, while a window of the same size recovers less than a fifth at the same cost (Figure 3 a). Along the block it declines only gradually, and the deepest layer's last slot still keeps about four fifths (Figure 3 b), well above what the window reaches at any slot. Its margin over the window widens as the context grows (Figure 3 c), because the target spreads its picks over the whole context while the window covers less and less of it. A selection the target has already paid for therefore serves the whole block, not only the row it was made for.

Mechanism. The gradual decline points to the query's content rather than its position. Rotating the query at row $A-1$ to each later row leaves its top-512 unchanged in all 72 slot, layer and context cells we checked. The energy of an index query sits almost entirely in its low-frequency RoPE components, which a shift of a few positions barely moves. That leaves the content, which changes one token at a time, so the target's focus drifts rather than switches. The decline should therefore follow how far into the block a row sits rather than how long the context is.

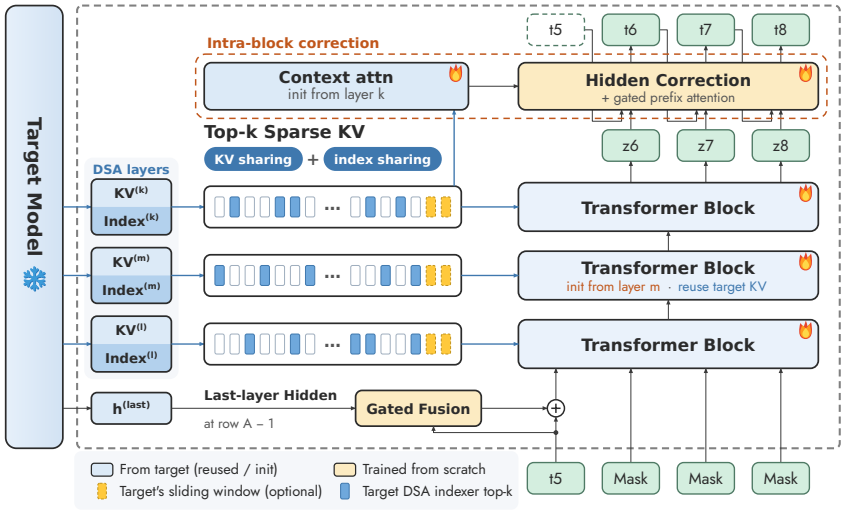

Figure 4: Overview of DFly2. Each of its three draft layers is paired with one of the target's DSA layers and reads the KV that layer cached, at the positions its indexer selected.

3.4 Summary

Both choices behind a drafter's context lead back to the target. Owning a cache of its own costs a drafter at prefill, in prefix caching and in decode concurrency, three costs that sharing the target's KV avoids by design without giving up reach (Section 3.1). Long contexts also call for a bounded read, and a natively sparse target has already paid for a query-aware selection once judged too costly for drafting (Sadhukhan et al., 2025) (Section 3.2). Made one step early, that selection largely holds across the block and covers far more than a window of the same size (Section 3.3). Section 4 turns both into mechanisms, *KV sharing* to read the target's KV cache and *Index sharing* to read the positions the target itself selected. The drafter then keeps no cache and runs no indexer of its own, and can still read anywhere in the context.

4 Method

A drafter must decide where its context comes from and which positions to read (Section 3). DFly2 leaves both decisions to the target (Figure 4). Its three layers fill one block in a single pass, each paired with one target layer. A draft layer reads the KV its paired layer has already cached, at the positions that layer selected.

KV sharing. The context a draft layer attends to is the target's (Figure 5a). Before the block, the KV injection of DFlash and DSpark projects every position into keys and values the drafter stores (Equation (3)) (Chen et al., 2026; Cheng et al., 2026). DFly2 instead reads the latents the paired layer has already cached,

$$(X_i,\, Y_i) \;=\; \big(\underbrace{W_i^{UK}}_{\text{trained}}\,\underbrace{\mathbf{C}_{\mathcal{R}_i}^{L_i}}_{\text{shared}},\; W_i^{UV}\,\mathbf{C}_{\mathcal{R}_i}^{L_i}\big), \qquad \mathbf{C}_{\mathcal{R}_i}^{L_i} \;=\; \{\mathbf{c}_s^{L_i}\}_{s\in\mathcal{R}_i}, \tag{6}$$

where $\mathbf{c}_s^{L_i}$ is the latent that layer L_i holds at position s and $\mathcal{R}_i$ is the set of positions draft layer i reads, itself taken from the target (Equation (8)). Inside the block the drafter still compresses its γ slots with a down-projection W_i^D and expands them with the same up-projections, so a draft layer's keys and values read

$$K_i \;=\; W_i^{UK}\big[\underbrace{\mathbf{C}_{\mathcal{R}_i}^{L_i}}_{\text{context}};\; \underbrace{W_i^D H_{\text{blk}}}_{\text{the block}}\big], \qquad V_i \;=\; W_i^{UV}\big[\mathbf{C}_{\mathcal{R}_i}^{L_i};\; W_i^D H_{\text{blk}}\big], \tag{7}$$

with H_{blk} the hidden states of the block and $W_i^{UK}W_i^D$ the factored form of W_i^K in Equation (3). Both terms are latents, so one attention covers the context and the block at once. What the drafter shares at run time is the target's cache, not its weights. The target supplies $\mathbf{C}_{\mathcal{R}_i}^{L_i}$, while every projection in Equations (6) and (7)

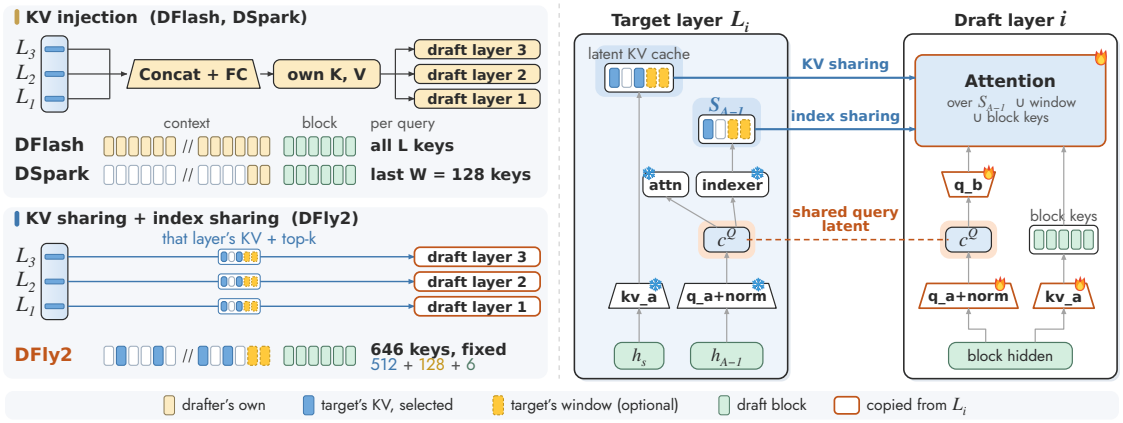

(a) Where each drafter's context KV comes from **(b) What DFly2 shares within one paired layer**

Figure 5: (a) DFlash and DSpark build their own KV, while DFly2 reads the target's KV at its top-k positions. (b) Only the latent KV and the selection pass at run time, so the drafter needs no indexer.

belongs to the drafter and is trained (Figure 5b). Only the target's DSA layers keep reusable latents, so the L_i are drawn from them in depth order, one per draft layer. These layers do not select the same positions, so the drafter sees several views of the context rather than one view repeated (Table 4 lists L and Section 5.4 ablates the choice).

The drafter stores nothing of its own per context position. In Equations (6) and (7) the context comes from the target's cache, and the rest is projection matrices of fixed size. The three costs of Section 3.1 therefore do not apply: the drafter needs no place on the prefill instances, leaves the target's prefix-cache hits unaffected, and costs no decode concurrency. A short window shrinks these costs without removing them.

Index sharing. A draft layer reads where its paired layer reads, and $\mathcal{R}_i$ in Equation (6) is that set:

$$\mathcal{R}_i \;=\; \mathcal{S}_{A-1}^{(L_i)} \cup \mathcal{N}, \qquad m \;=\; |\mathcal{R}_i|. \tag{8}$$

Here $\mathcal{S}_{A-1}^{(L_i)}$ is the top-k the indexer of layer L_i scored at row $A-1$ (Equation (4)), and $\mathcal{N}$ the target's near window before the anchor. What the second term adds follows from the target's structure. One such set serves all γ slots of a layer, and a selection made one step early largely holds across the block (Section 3.3).

The drafter neither scores the context nor reads more of it as the context grows. The verify pass that produced the anchor already made the selection, so a draft layer runs no scoring pass of its own and carries no indexer (Figure 5b). Both terms are fixed in size by the target, so the number of positions a draft step reads, m, does not grow with the context (Figure 5a). A draft step's attention over the context is therefore bounded, and so is the drafting time T_{draft} (Equation (2)).

The two terms differ across the three targets, and a draft layer takes them as each target gives them, with nothing to tune (Table 4). The first term is a top-k over single positions on DS-V4.1-Flash and over groups of adjacent positions on GLM-5.3-Flash, so the same k can mean several times as many positions. Where the target stores its near window apart from the latents the top-k picks from, as DS-V4.1-Flash does, the second term adds what the top-k cannot pick. On GLM-5.3-Flash and Hy4 flash the near window lies within those latents, so the top-k can already pick it (Section 5.4 ablates the term). Section 5.5 asks whether reading beyond $\mathcal{R}_i$ would help.

Initializing from the target. The paired layer also gives a draft layer its starting weights, so one pairing settles what the draft layer reads, where it reads and what it starts from. A draft layer's **attention** has the form of the paired layer's and starts from that layer's weights,

$$O_i \;=\; W_i^{O}\,\mathrm{Attn}\big(\underbrace{W_i^{UQ}\,\mathrm{Norm}(W_i^{DQ}H_{\text{blk}})}_{Q_i},\; K_i,\; V_i\big), \qquad W_i^{\star}\big|_{\text{init}} \;=\; W_{L_i}^{\star}, \tag{9}$$

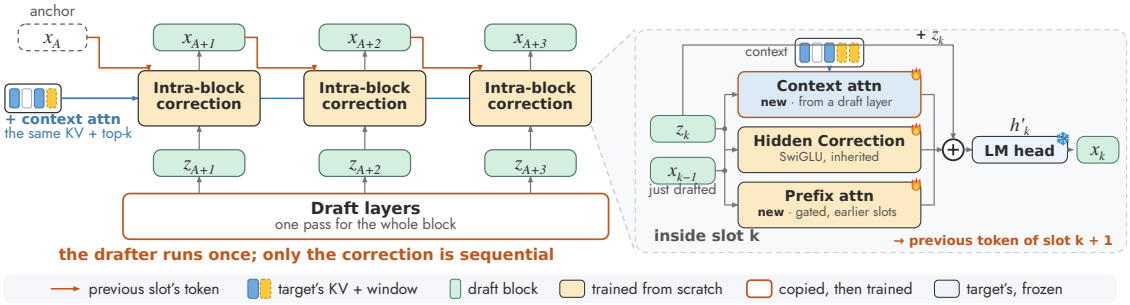

Figure 6: Intra-block correction across the block (left) and inside slot k (right). Slot k is corrected from its hidden state and the token drafted at slot $k-1$.

with K_i and V_i from Equation (7) and $\star$ running over DQ, UQ, D, UK, UV and O. The norms are copied with the matrices, and all of them are trained from there. Orthrus likewise starts its trainable projections from the target's (Nguyen et al., 2026). On a natively sparse target $W^{DQ}_{L_i}$ also feeds the indexer queries, so the copy inherits the target's *shared query latent* even without the indexer (Figure 5b). A draft query therefore starts in the query space that produced the top-k it reuses, on all three targets.

The **feed-forward** side also starts from the paired layer, but a drafter is small by design, so a draft layer draws only a fraction of that layer's routed experts,

$$\theta_i^{\text{exp}}\big|_{\text{init}} \;=\; \theta_{L_i}^{\text{exp}}\big[\,\mathcal{E}_i\,\big], \qquad \mathcal{E}_i \;\sim\; \text{Unif}\big\{\,\mathcal{B}\subseteq\mathcal{E}^{L_i}\;:\;|\mathcal{B}|=E\,\big\}, \tag{10}$$

where θ^{exp} is a layer's routed experts, $\mathcal{E}^{L_i}$ their ids at layer L_i, and E how many a draft layer holds. Only the choice of experts is random, and their weights are the target's. Each drawn expert keeps its router row and correction bias, and the shared expert and the post-attention norm are copied whole. Section 5.4 measures each half of the initialization alone and the two together.

Last-layer hidden injection. The target's last-layer hidden state goes only to its LM head and MTP layer, so no cache holds it and reading the target's KV does not reach it. Yet it is the state the target produced the anchor from, and MTP heads draft from it (Gloeckle et al., 2024). DFly2 fuses it with the anchor's embedding, as the target's MTP layer does, and adds the result at the anchor slot through a gate that starts at zero (Figure 4). In a disaggregated deployment this adds one vector per request to the prefill-to-decode transfer.

Intra-block correction. The correction DFly2 runs inside the block returns to the context. Block-parallel drafters correct slot k from its hidden state and the token drafted at slot $k-1$. TreeFlash's SwiGLU over that pair is the *hidden correction* we inherit unchanged (Rheinboldt et al., 2026). Such an update can only rework what the slot already holds: after *the value in row 3 is* it cannot recover the number, and that is where corrections elsewhere stop (Section 6). Reading the context would normally mean storing or projecting something per position, and here it costs neither, since the context is the target's and already in its cache. Alongside the inherited update, a *context attention* reads one draft layer's context and a *gated prefix attention* the block up to slot k (Figure 6),

$$\mathbf{h}_k \;\leftarrow\; \mathbf{h}_k \;+\; \underbrace{\text{SwiGLU}(\mathbf{z}_k)}_{\text{inherited}} \;+\; W_c^O\Big[\,\underbrace{\text{Attn}\big(\mathbf{q}_k,\,\mathbf{C}^{L_i}_{\mathcal{R}_i}\big)}_{\text{context}} \;+\; \underbrace{\mathbf{g}_k \odot \text{Attn}\big(\mathbf{q}_k,\,\mathbf{h}_{0:k}\big)}_{\text{the block}}\,\Big], \tag{11}$$

where i is the draft layer it reads, $\mathbf{z}_k = [\,\text{Norm}(\mathbf{h}_k)\,;\,\text{Norm}(\text{Emb}(y_{k-1}))\,]$ is that same pair, and the query $\mathbf{q}_k$ and the per-head gate $\mathbf{g}_k$ are projected from it. The correction's own output projection W_c^O starts at zero with $\mathbf{g}_k$, so it begins as the inherited one alone. It runs slot by slot, the one sequential part of a draft step (Section 5.4 weighs its cost against what it buys).

Training. Only the drafter's parameters are trained, and Section 5.1 gives the data and the schedule. The objective is single-stage: each slot is matched to the target's distribution there, and the slots are weighted by

Table 4: What each target settles. Positions per step is the top-k, the near window and the block's six rows. A dash marks a window the top-k already covers.

Target	Paired layers L	Top-k	Near window	Positions per step
DS-V4.1-Flash	$[28, 32, 36]$	512	128	646
GLM-5.3-Flash	$[35, 39, 43]$	512 pools of four	–	2,054
Hy4 flash	$[33, 37, 41]$	2,048	–	2,054

the accepted length that depends on them,

$$\mathcal{L} \;=\; \frac{\sum_{k=1}^{\gamma} w_k\, \ell_k}{\sum_{k=1}^{\gamma} w_k}, \qquad \ell_k \;=\; \lambda_k\, \mathrm{KL}\big(p_k \,\|\, \hat{p}_k\big) \;+\; (1-\lambda_k)\, \mathrm{TV}\big(p_k, \hat{p}_k\big), \tag{12}$$

where p_k is the target's distribution at slot k, $\hat{p}_k$ the drafter's, and ℓ_k the LK distillation loss (Samarin et al., 2026), whose λ_k follows the overlap between the two. The D-PACE weights w_k (Wu et al., 2026) are the accepted length still expected from slot k on, taken from the drafter's probabilities at the earlier slots and carrying no gradient. Serving takes the selection from the verify pass, and training, which has none, recomputes it with a frozen copy of each paired layer's indexer. A draft layer therefore sees the same latents and the same positions in training as at serving.

5 Experiments

5.1 Experimental setup

Models and baselines. We evaluate on three natively sparse targets, DS-V4.1-Flash (DeepSeek-AI, 2026), GLM-5.3-Flash (GLM-5 Team, 2026) and Hy4 flash (Tencent Hy Team, 2026b). Their sparse attention differs in each case (Section 2), so each sets its own paired layers and top-k, the near window where it keeps one, and the corresponding number of positions per draft step (Table 4). Every drafter has three layers of 128 routed experts, 6B to 14B parameters of which 0.4B to 1.1B are active per token, and reuses the target's embedding and output head. The baselines cover both ways of owning a cache (Section 3.1): DFlash (Chen et al., 2026) keeps one over the whole context, and DSpark (Cheng et al., 2026) bounds its own to a sliding window. DFlash† is DFlash held to that same window, 128 tokens and 512 on Hy4 flash, which separates the window from DSpark's other changes. We train all three ourselves with the same recipe, data and number of steps as DFly2. They have the same three-layer MoE stack and come within 2% of its size. Unless noted otherwise, every number we report is measured on NVIDIA H20 GPUs.

Benchmarks. The three workload families, all taken from earlier work, differ in how much context a request carries. The general-purpose family is the twelve evaluation sets that ship with DeepSpec, the training repository released with DSpark (Cheng et al., 2026). The main tables report seven of them: GSM8K (Cobbe et al., 2021), MATH-500 (Lightman et al., 2024), AIME25 (MAA, 2025), MBPP (Austin et al., 2021), HumanEval (Chen et al., 2021), LiveCodeBench (Jain et al., 2025) (LCB) and MT-Bench (Zheng et al., 2023). The long-context family follows LongSpark (He et al., 2026): continuation tasks from LongSpec's corpus (Yang et al., 2025) at 32K, LongSpark's own CodeSpan at 64K, and LongSWE-Bench (Rando et al., 2025) at 128K, LSWE 128K in the tables. The agentic family follows DLite (Xu et al., 2026): AgentLongBench (Fang et al., 2026) (ALB), MemoryAgentBench (Hu et al., 2026) (MemAgent), the long-context multi-turn category of BFCL-v3 (Patil et al., 2025), and LongSWE-Bench issues at 64K (SWE 64K). Each family runs under four sampling settings, greedy ($T = 0$) and stochastic ($T = 1$, or $T = 0.9$ on Hy4 flash) crossed with 2K chat and 8K thinking. The main tables report the 2K chat setting, thinking disabled and at most 2048 generated tokens. The ablation tables average each family over its sets and all four settings, twelve family-setting cells per drafter.

Training data. We train on DFly-SFT, a general-purpose instruction-tuning corpus of 989,861 conversations and 7.42B tokens, all of it from public datasets. The conversations fall into ten capability domains (Table 5), labeled by their source where it gives a label and by an LLM annotator otherwise. Samples are capped at

Table 5: DFly-SFT data mix. Mix is the share of samples. Token counts and lengths use the GLM-5.3-Flash tokenizer, and Think is the fraction of samples whose assistant turns carry an explicit reasoning trace.

Data type	Mix	Samples	Tokens	Median len.	Think
Math (advanced)	19.2%	190,000	2.90B	13,982	98.6%
Code (competitive)	11.1%	110,000	1.96B	18,175	94.7%
Code (engineering)	9.6%	95,479	0.56B	4,064	46.9%
Knowledge & STEM	15.2%	150,000	0.48B	2,075	51.3%
Instruction following	7.1%	70,000	0.47B	5,751	85.7%
Code (function-level)	11.1%	110,000	0.41B	686	37.9%
Writing	10.1%	100,000	0.24B	1,425	40.1%
Text transformation	6.1%	60,000	0.19B	1,972	37.4%
Logical reasoning	2.5%	24,382	0.11B	1,461	65.0%
Math (word problems)	8.1%	80,000	0.09B	477	20.2%
Total	100.0%	989,861	7.42B	3,621	61.6%

32,768 tokens, and each domain is filled with its longest samples first, which makes the mixture long-form and reasoning-heavy. An explicit reasoning trace appears in 61.6% of samples and 34.9% run past 8K tokens, so a drafter meets long contexts during training. The corpus is deduplicated globally and decontaminated against the twelve DeepSpec evaluation sets (Cheng et al., 2026) with an n-gram overlap test.

Implementation details. The same recipe trains every drafter, on all three targets and in every ablation: three epochs at a global batch of 256, sequence length 32K, with Muon (Jordan et al., 2024), on the objective of Equation (12). All speculative methods share a block of $\gamma = 5$ and run on vLLM (Kwon et al., 2023), where DFly2 registers no KV group and reads the pages its paired layers already hold. Accepted length is measured on a single tp8 engine. Each accepted length is the mean of the last three runs. Speedups come from each target's own deployment at a concurrency of one, prefill and decode on separate tp8 instances for DS-V4.1-Flash and GLM-5.3-Flash and one tp8 instance for Hy4 flash. The concurrency sweep of Figure 7 runs DS-V4.1-Flash greedily, every instance at DP8 $\times$ TP1 with EP8. Short text pools the twelve general-purpose sets on one prefill and one decode instance, at one to 128 concurrent requests. CodeSpan 64K and SWE 64K run at one to 32 with eight prefill instances per decode instance, since at that length a single prefill instance cannot keep the decode batch fed.

5.2 Experimental results

General-purpose workloads. DFly2 accepts more than both baselines in all eighteen cells, and on the general-purpose family it is 3.9% to 10.5% above DSpark, the stronger baseline there (Table 6). On short text, how much a drafter keeps in its own cache makes little difference, as holding DFlash to a 128-token window, 512 on Hy4 flash, moves its accepted length by at most 2%. DFly2 learns in training to draft from the positions the target's sparse attention selects, and trained on its near window alone instead, it accepts less on this family for all three targets (Table 10). Training on the target's sparse selection thus helps the drafter generalize.

Long-context and agentic workloads. DFly2 holds its accepted length as prompts outgrow the 32K every drafter was trained at (Table 7). On the agentic family, whose prompts mostly run longer, its average stays within 4% of its general-purpose average on every target and temperature, while DSpark's falls 11% to 21%. Its margin over the stronger baseline widens to 10.0% to 31.5% on the long-context and agentic families. The baselines show the two limits of a drafter keeping its own context. A window cannot see past itself, and a full cache sees everything but risks breaking where training stops: on Hy4 flash, DFlash falls to as few as 1.08 tokens a step once prompts pass 32K. DFly2 instead reads the target's KV at the positions the target's sparse indexer has already selected. That selection can reach anywhere in the context, bounded by neither a near window nor the training length, yet it stays a fixed top-k. A draft step on a long prompt thus reads as many positions as it did in training, so what DFly2 learned there carries over, and it inherits the target's length generalization.

Table 6: Mean accepted length τ on general-purpose benchmarks, each cell at temperature 0/1 (0.9 on Hy4 flash). Bold and underline mark the best and second best per target and temperature.

		Math			Code			Chat	Overall
Target	Drafter	GSM8K	MATH	AIME25	MBPP	HumanEval	LCB	MT-Bench	Avg.
DS-V4.1-Flash	DFlash	3.95/3.64	3.93/3.42	3.54/2.97	3.82/3.63	4.25/3.99	3.34/2.80	2.70/2.51	3.65/3.28
	DFlash†	3.94/3.59	3.91/3.40	3.51/2.94	3.81/3.58	4.08/3.82	3.26/2.77	2.68/2.53	3.60/3.23
	DSpark	4.24/4.00	4.18/3.74	3.69/3.24	4.00/3.83	4.27/4.07	3.44/3.04	2.82/2.67	3.81/3.51
	DFly2	**4.62/4.45**	**4.43/4.10**	**4.01/3.60**	**4.28/4.17**	**4.70/4.56**	**3.76/3.38**	**3.03/2.91**	**4.12/3.88**
GLM-5.3-Flash	DFlash	3.93/3.65	3.73/3.29	3.25/2.77	3.76/3.38	3.84/3.44	2.91/2.64	2.58/2.34	3.43/3.07
	DFlash†	4.00/3.73	3.83/3.38	3.36/2.92	3.79/3.43	3.74/3.40	2.99/2.73	2.66/2.35	3.48/3.13
	DSpark	4.23/3.97	4.06/3.61	3.58/3.17	4.02/3.66	3.91/3.62	3.16/2.92	2.80/2.50	3.68/3.35
	DFly2	**4.53/4.37**	**4.31/4.00**	**3.88/3.51**	**4.24/4.01**	**4.36/4.03**	**3.41/3.24**	**2.96/2.72**	**3.96/3.69**
Hy4 flash	DFlash	3.86/3.77	4.11/3.95	3.80/3.69	3.81/3.68	3.02/3.04	3.59/3.51	2.90/2.71	3.58/3.48
	DFlash†	3.87/3.78	4.13/3.98	3.80/3.67	3.83/3.69	3.00/3.01	3.61/3.52	2.91/2.74	3.59/3.48
	DSpark	4.16/4.09	4.42/4.29	4.05/3.98	4.02/3.92	3.17/3.22	3.79/3.72	3.07/2.92	3.81/3.73
	DFly2	**4.40/4.43**	**4.55/4.52**	**4.22/4.20**	**4.16/4.11**	**3.34/3.42**	**3.92/3.91**	**3.16/3.03**	**3.96/3.95**

Table 7: Mean accepted length τ on long-context and agentic benchmarks, formatted as in Table 6. Long and Agen. average the three long-context and four agentic sets.

		Long context			Agentic				Group average	
Target	Drafter	LongSpec 32K	CodeSpan 64K	LSWE 128K	ALB	MemAgent	BFCL-v3	SWE 64K	Long	Agen.
DS-V4.1-Flash	DFlash	2.98/2.55	3.81/3.13	3.42/2.96	4.38/4.13	2.84/2.60	3.40/3.11	3.58/3.14	3.40/2.88	3.55/3.25
	DFlash†	2.74/2.40	3.43/2.86	3.15/2.77	3.78/3.62	2.80/2.52	3.21/2.98	3.19/2.80	3.11/2.68	3.24/2.98
	DSpark	2.82/2.48	3.58/2.92	3.29/2.98	3.87/3.59	2.85/2.65	3.33/3.09	3.35/2.99	3.23/2.79	3.35/3.08
	DFly2	**3.12/2.81**	**4.14/3.59**	**3.97/3.55**	**5.10/4.72**	**3.36/3.13**	**3.77/3.55**	**4.02/3.60**	**3.74/3.32**	**4.06/3.75**
GLM-5.3-Flash	DFlash	2.55/2.17	3.12/2.48	2.92/2.64	3.12/1.92	2.79/2.50	2.99/2.75	2.91/2.60	2.86/2.43	2.95/2.44
	DFlash†	2.69/2.23	3.22/2.60	2.97/2.65	3.50/3.45	2.68/2.48	3.00/2.73	2.86/2.61	2.96/2.50	3.01/2.82
	DSpark	2.75/2.37	3.32/2.77	3.08/2.79	3.53/3.73	2.74/2.58	3.15/2.91	2.98/2.74	3.05/2.64	3.10/2.99
	DFly2	**3.31/2.92**	**4.40/3.63**	**3.82/3.52**	**5.49/4.77**	**3.30/3.12**	**3.61/3.35**	**3.73/3.48**	**3.85/3.36**	**4.03/3.68**
Hy4 flash	DFlash	2.59/2.41	1.17/1.10	1.12/1.09	1.50/1.36	1.08/1.06	3.18/3.20	2.26/1.88	1.63/1.53	2.00/1.88
	DFlash†	2.68/2.50	3.47/3.28	3.23/3.06	3.32/3.06	2.63/2.47	2.90/2.92	3.40/2.94	3.13/2.95	3.06/2.85
	DSpark	2.74/2.56	3.54/3.35	3.31/3.11	3.32/3.10	2.77/2.56	3.08/2.98	3.45/3.14	3.20/3.01	3.16/2.95
	DFly2	**2.94/2.75**	**4.12/4.03**	**3.83/3.60**	**5.41/5.22**	**3.23/2.95**	**3.69/3.67**	**4.12/3.70**	**3.63/3.46**	**4.11/3.88**

5.3 Serving efficiency

Throughput. DFly2 has the highest average speedup on every target and family, in decoding and end to end alike (Tables 8 and 9). Its lead in accepted length carries into decoding speed (Equation (1)) and, like it, is widest on the long-context and agentic families. Over the stronger baseline, DFly2 decodes 2% to 6% faster on the general-purpose family, 6% to 18% faster on the long-context one and 12% to 30% faster on the agentic one. End to end, both DFlash variants on DS-V4.1-Flash pay for building their own cache over the prompt with up to twice the time to the first token (Section 3.1), and end the agentic sets slower than not speculating. DFly2 builds no cache of its own at prefill and thus leaves prefill speed unaffected. On long prompts its advantage therefore has two sources, faster decoding from its higher accepted length and no added overhead at prefill.

High-concurrency throughput. We compare throughput and per-user decode speed across concurrency levels on short text and on a long-context and an agentic set, CodeSpan 64K and SWE 64K (Figure 7). DFly2's curve lies outside every baseline's at every level, and the gap does not close as the batch grows. At the same per-user speed as DSpark, DFly2 delivers 28% higher throughput on short text and 62% and 81% higher on the two 64K sets, the largest gains again on long prompts. At the same throughput, each user decodes 15% to 24% faster. DFly2's advantage at a single request thus largely carries over to high concurrency.

Table 8: Speedup over the target alone on general-purpose benchmarks, greedy, at a concurrency of one. Each cell gives decoding / end-to-end speedup in × over the target alone, the latter including prefill. Bold and underline mark the best and second best per target and speedup.

		Math			Code			Chat	Overall
Target	Drafter	GSM8K	MATH	AIME25	MBPP	HumanEval	LCB	MT-Bench	Avg.
DS-V4.1-Flash	DFlash	3.23/2.87	3.26/3.14	2.86/2.73	3.11/3.04	3.41/3.29	2.67/2.54	2.21/**2.16**	2.96/2.82
	DFlash†	3.21/2.83	3.19/3.06	2.75/2.71	3.09/3.02	3.28/3.17	2.62/2.49	2.19/2.14	2.91/2.77
	DSpark	3.30/2.91	3.24/3.11	2.80/2.76	3.09/3.02	3.31/3.18	2.63/2.52	2.15/2.11	2.93/2.80
	DFly2	**3.56/3.09**	**3.42/3.27**	**3.03/2.98**	**3.28/3.20**	**3.57/3.36**	**2.87/2.74**	**2.27**/2.15	**3.14/2.97**
GLM-5.3-Flash	DFlash	2.69/2.31	2.54/2.42	2.20/2.17	2.57/2.44	2.59/2.50	1.98/1.89	1.79/1.77	2.34/2.21
	DFlash†	2.77/2.39	2.62/2.50	2.26/2.22	2.63/2.49	2.56/2.47	2.04/1.98	1.83/1.80	2.39/2.26
	DSpark	2.85/**2.42**	2.67/2.55	2.32/2.28	**2.71/2.57**	2.62/2.54	2.10/2.04	1.85/1.82	2.45/2.32
	DFly2	**2.90**/2.40	**2.71/2.56**	**2.43/2.39**	**2.71**/2.56	**2.74/2.64**	**2.16/2.10**	**1.88/1.85**	**2.50/2.36**
Hy4 flash	DFlash	3.23/2.98	3.47/3.35	3.21/3.16	3.21/3.09	2.55/2.50	3.03/2.84	2.44/2.39	3.02/2.90
	DFlash†	3.25/3.00	3.49/3.38	3.22/3.18	3.24/3.13	2.53/2.48	3.04/2.87	2.44/2.38	3.03/2.92
	DSpark	3.47/3.16	3.71/3.57	3.42/3.37	3.38/3.26	2.67/2.60	3.17/2.97	2.58/2.51	3.20/3.06
	DFly2	**3.70/3.34**	**3.85/3.69**	**3.56/3.51**	**3.51/3.38**	**2.81/2.75**	**3.30/3.07**	**2.66/2.59**	**3.34/3.19**

Table 9: Speedup on long-context and agentic benchmarks, measured and formatted as in Table 8. DFlash† reads a window at decode but still builds a full cache at prefill.

		Long context			Agentic				Group average	
Target	Drafter	LongSpec 32K	CodeSpan 64K	LSWE 128K	ALB	MemAgent	BFCL-v3	SWE 64K	Long	Agen.
DS-V4.1-Flash	DFlash	**2.46**/1.75	3.05/1.67	2.59/1.16	3.46/0.69	2.22/0.63	2.67/0.78	2.79/1.49	2.70/1.52	2.78/0.90
	DFlash†	2.32/1.70	2.70/1.55	2.52/1.13	3.09/0.83	2.25/0.69	2.63/0.84	2.55/1.47	2.51/1.46	2.63/0.96
	DSpark	2.28/1.99	2.63/2.06	2.49/1.68	2.95/1.21	2.37/**1.27**	2.58/**1.37**	2.54/1.89	2.47/1.91	2.61/1.44
	DFly2	2.44/**2.10**	**3.23/2.37**	**2.94/1.86**	**3.82/1.25**	**2.60**/1.24	**2.97**/1.36	**3.00/2.13**	**2.87/2.11**	**3.10/1.50**
GLM-5.3-Flash	DFlash	1.70/1.47	2.03/1.46	1.89/1.28	2.15/1.09	1.83/**1.01**	2.02/1.16	1.88/1.36	1.87/1.40	1.97/1.15
	DFlash†	1.86/1.56	2.12/1.48	2.04/1.28	2.45/1.07	1.92/0.97	2.08/1.09	1.99/1.50	2.00/1.44	2.11/1.16
	DSpark	1.84/1.53	2.24/1.58	2.12/1.38	2.48/1.08	1.90/0.97	2.17/1.13	2.02/1.49	2.07/1.49	2.14/1.17
	DFly2	**2.14/1.75**	**2.82/1.79**	**2.40/1.39**	**3.57/1.10**	**2.20**/0.98	**2.34/1.17**	**2.35/1.62**	**2.45/1.64**	**2.61/1.22**
Hy4 flash	DFlash	2.15/1.93	0.96/0.96	0.90/1.00	1.23/1.06	0.87/0.97	2.64/1.87	1.87/1.46	1.34/1.30	1.65/1.34
	DFlash†	2.28/2.03	2.96/2.26	2.75/1.60	2.81/1.26	2.18/1.29	2.46/1.82	2.90/1.80	2.66/1.96	2.59/1.54
	DSpark	2.32/2.06	3.00/2.28	2.79/1.60	2.80/1.26	2.23/1.27	2.59/1.85	2.92/1.81	2.70/1.98	2.63/1.55
	DFly2	**2.47/2.17**	**3.48/2.51**	**3.22/1.68**	**4.55/1.35**	**2.54/1.31**	**3.08/2.04**	**3.46/1.96**	**3.06/2.12**	**3.41/1.67**

5.4 Ablation study

What sharing buys. Of the two things the drafter shares, accepted length comes from the selection, whose scoring pass over every position the target has already run (Table 10 and section 3.2). Removing its top-k leaves only the near window, the largest loss here and two to six times as costly on the long-context family as on the general-purpose one. A window chooses by position, and the longer the context, the more it misses. The near window itself is worth about 3% on DS-V4.1-Flash and nothing on the other two. On DS-V4.1-Flash the indexer scores one set of latents shared across layers, while each layer also caches its own for the most recent 128 positions, which the top-k cannot reach (Section 4). A drafter with its own cache over that same near window accepts at most 3% more on GLM-5.3-Flash and Hy4 flash. Sharing the cache therefore affects accepted length only slightly while removing the three costs of Section 3.1. Because the drafter reads the selection as each target makes it, how much each term adds follows the target's structure, and one design serves all three with nothing to tune in what it reads (Section 4).

Which layers to pair with. The layers the drafter pairs with decide how many views of the context its three layers read. Moving the pairing from layers 28, 32 and 36 to the adjacent layers 37 to 39 costs 4.6% and lowers all twelve cells (Table 11). On DS-V4.1-Flash layers 28, 32 and 36 each run an indexer, while layers

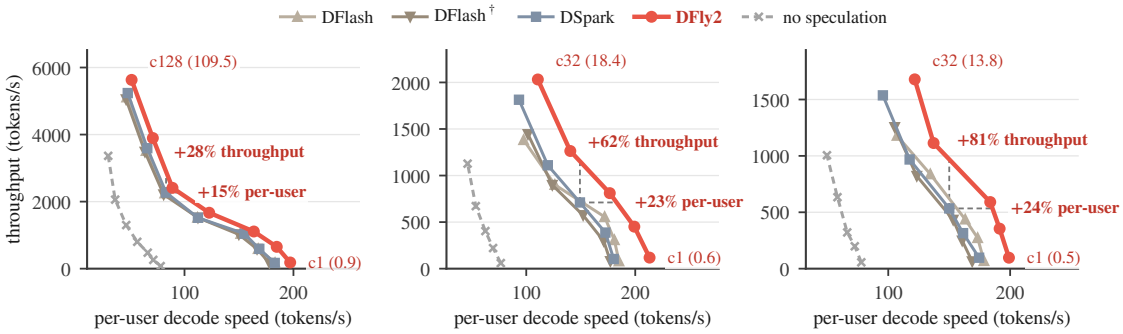

(a) Short text, 1P1D **(b) CodeSpan 64K, 8P1D** **(c) SWE 64K, 8P1D**

Figure 7: Throughput against per-user decode speed on DS-V4.1-Flash, greedy, one point per concurrency level. Labels give the concurrency and, in parentheses, the requests decoding. Dashed guides compare DFly2 with DSpark at equal per-user speed and at equal throughput. Throughput counts the whole deployment.

Table 10: Mean accepted length τ with one part of what DFly2 reads removed, averaged over each family's sets and four sampling settings. Bold marks the best per column.

	DS-V4.1-Flash			GLM-5.3-Flash			Hy4 flash		
What the drafter reads	Gen.	Long	Agen.	Gen.	Long	Agen.	Gen.	Long	Agen.
Full (DFly2)	**3.57**	**3.62**	**3.75**	3.56	3.67	**3.91**	**3.36**	**3.46**	3.76
w/o near window $\mathcal{N}$	3.45	3.50	3.63	**3.57**	**3.68**	3.86	**3.36**	3.45	**3.77**
w/o the target's top-k	3.14	2.76	2.95	3.42	3.10	3.34	3.31	3.12	3.13
w/o sharing (own cache)	3.46	3.34	3.37	3.46	3.18	3.43	3.32	3.14	3.17

37 to 39 reuse the 512 latents layer 36 picked, so the draft layers read one view three times. On a target that shares one cache across layers, the layers to pair with are thus those with an indexer of their own. Each of them contributes its own top-k, so the drafter gathers a richer picture of where the target attends.

Where the weights start. Both halves of a draft layer start from its paired layer, the attention copied from it and the feed-forward side drawn from its experts (Equations (9) and (10)), and each half helps about equally. Copying the attention alone, with the feed-forward side random, costs 1.6%, and the reverse costs 1.1% (Table 11). The copied weights come from the same layer whose latents the draft layer reads. One pairing thus gives a draft layer, at no cost, both the context it reads and weights already matched to it, aligning the draft layer with the target from the first step of training.

What the drafter adds. Both parts the drafter adds on top of what it shares raise accepted length. Removing the two attention reads of the intra-block correction (Equation (11)) costs 2.8%, and removing the last-layer hidden injection costs 1.6% (Table 11). The two reads take under 1% of a decode step on a tp8 engine. The correction lets each slot look back at the context and at the slots before it, and the injection supplies the last-layer state from which the target produced the anchor. Feeding the injection to every slot instead performs on par, so one entry point suffices. Sharing hands the drafter the target's cache, and these two parts add what no cache holds, the block in progress and the state behind the anchor.

5.5 Analysis

What YOCO opens. DS-V4.1-Flash follows YOCO (Sun et al., 2024b) in keeping one set of latents that all layers share, and the drafter reads that same cache. Every indexer of the target therefore selects over what the drafter reads, which opens a design space in how the drafter picks its top-k. We explore three dimensions, earlier rows, more layers and a larger budget per row, and score each read by its top-k mass recall (Figure 8). All three start from DFly2's own read, the target's top-512 at the row before the anchor, which on average

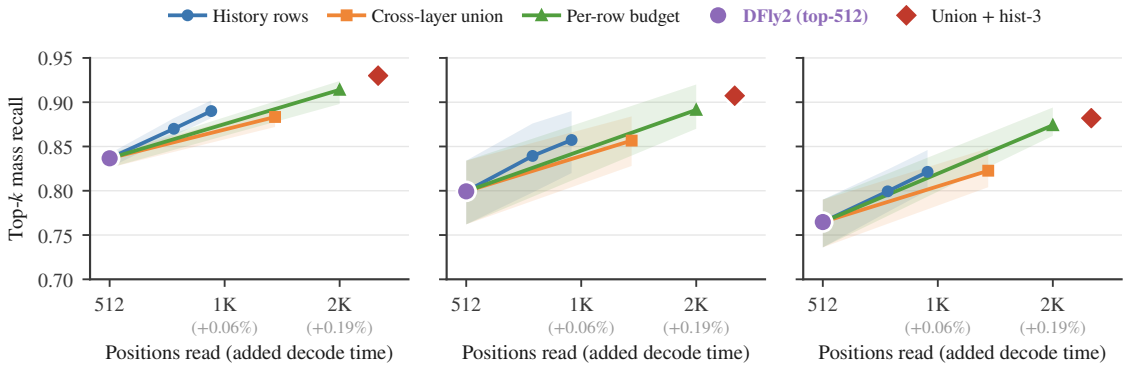

(a) Paired layer 28 **(b) Paired layer 32** **(c) Paired layer 36**

Figure 8: Top-k mass recall on DS-V4.1-Flash, the index-score share of each slot's own top-512 that a read includes. Lines average slots 1–5 and three context lengths, with bands spanning the lengths.

Table 11: Mean accepted length τ on DS-V4.1-Flash, one change to DFly2 at a time, over each family's sets and four sampling settings.

Dimension	Change	Gen.	Long	Agen.
–	**DFly2**	3.57	3.62	3.75
Pairing	Adjacent 37–39	3.36	3.47	3.61
Initialization	Attention only	3.51	3.56	3.69
	FFN only	3.51	3.58	3.72
Correction	w/o attn. reads	3.47	3.51	3.65
Injection	w/o injection	3.51	3.56	3.68
	Broadcast	**3.58**	**3.64**	**3.77**

Table 12: As in Table 11, with DFly2 reading beyond the target's top-k. Candidates reads every first-stage candidate block, up to 16K positions.

Axis widened	Change	Gen.	Long	Agen.
–	**DFly2**	3.57	3.62	3.75
History rows	hist-2	3.53	3.63	3.76
	hist-3	3.57	3.67	3.79
Per-row budget	top-1024	3.57	3.68	3.79
Cross-layer	Union	3.57	3.68	3.81
Both	Union+hist-3	3.57	3.71	**3.84**
Selection stage	Candidates	**3.58**	**3.72**	3.82

already recalls 0.80 of each slot's own top-512 by index score. Adding two earlier rows reaches 0.86 at under 1K positions, the union of the target's five indexers 0.85 at 1.4K, and a 2K budget per row 0.89. Earlier rows thus gain the most recall per added position, about twice the other two, and combining them with the union reaches 0.91 at 2.6K positions. Each read stays cheap, as 2K positions add 0.19% to a decode step, against 14% to 23% for the whole context at 128K (Section 3.1).

Capped at 2K positions, the YOCO union with history matches the accepted length of reading every first-stage candidate block, up to eight times as many (Table 12). On the general-purpose sets, contexts during drafting average under 2K positions, so the reads differ little in what they see and accepted length rises by at most 0.3%, comparable to the variation between runs. On the long-context and agentic families, the union raises accepted length over DFly2's own read by 2.6% and 2.4%, against 2.7% and 1.9% for the candidate read. Up to 16K positions, the target's top-k thus holds nearly all that the drafter can use, and the dimensions YOCO opens recover the rest at little cost.

6 Related Work

Block-parallel drafters. Block-parallel drafters build a context KV of their own that spans the prompt and grows with it. DFlash (Chen et al., 2026) fuses hidden states from several target layers, injects them into every draft layer as extra keys and values, and drafts a whole block in one forward pass. Its successors refine the conditioning signal (Zhang et al., 2026; Lin et al., 2026) or model the dependency between slots (Cheng et al., 2026; Rheinboldt et al., 2026; Huang et al., 2026), and DFly (Liu et al., 2026a) does both. All of them keep this cache, and their intra-block corrections do not return to the context. DLite (Xu et al., 2026) avoids the cost by keeping no cache at all and conditioning on hidden states from many target layers at the latest position.

DFly2 also builds no context KV of its own, but it still reads the context through the target's cache, and so does its intra-block correction.

Reading the target's KV cache. Reading the target's KV cache instead of building one goes back to GliDe with CAPE (Du et al., 2024), but drafters on this line choose positions by rules that ignore the current query. GliDe, LongSpec (Yang et al., 2025), KVShot (Liu et al., 2026b) and Orthrus (Nguyen et al., 2026) read the whole cache, H-Spec (Jiang et al., 2026) reads a fixed window, and LongSpark (He et al., 2026) reads a recent window plus a summary pooled by fixed queries. All of them were developed on dense targets, which attend to every position and so offer no selection to follow. DFly2 also reads the target's cache, but at the positions a natively sparse target attends to.

Query-aware selection. Query-aware selection gives good drafts, but prior work either finds it too costly for drafting or extracts it from a dense target's attention. TriForce (Sun et al., 2024a) drafts from a cache retrieved by the current query, whereas MagicDec (Sadhukhan et al., 2025) finds per-query search too costly at large batch sizes and drafts from a static sparse cache. Verification-guided methods read the selection off the target's attention during verification. Either the target itself then drafts (Zhao et al., 2026; Yue et al., 2026), or a separate drafter keeps a cache of its own and reads only the chunks the target ranks highest (Cha et al., 2025). In every case, the drafting side has to compute or extract the selection itself.

Speculative decoding on natively sparse targets. Natively sparse targets such as DeepSeek-V3.2, DeepSeek-V4 and GLM-5 compute a query-aware selection as part of their own attention (Yuan et al., 2025; DeepSeek-AI, 2025, 2026; GLM-5 Team, 2026), and within a model one selection can already serve several layers (Bai et al., 2026). Systems work on these targets speeds up the target's side of speculative decoding, through sparse verification kernels (Wang et al., 2026), indexer prediction (Fu et al., 2026) and tree verification over compressed attention (Liu & Li, 2026). The drafters paired with these targets, however, still bring a context of their own: the MTP modules shipped with these models keep their own KV cache (GLM-5 Team, 2026; Xu et al., 2026), and DSpark serves DeepSeek-V4 with its own 128-token window (Cheng et al., 2026). GLM-5.2 goes as far as reusing a selection across draft steps, but the selection is its MTP module's own.[3] To our knowledge, no prior drafter reuses the target's own selection. DFly2 reuses both the target's cache and its selection, so the drafter keeps no cache and computes no selection of its own.

7 Conclusion

Long contexts have turned targets natively sparse, and speculative decoding increasingly runs on them. Their drafters still bring a context of their own, a cache that costs more than the memory it takes, read at positions a rule fixed before the query arrives. A natively sparse target picks positions for each query already, over a cache it keeps anyway, and has paid for both. DFly2 shares both, reading each paired target layer's cached KV (*KV sharing*) at the positions that layer's indexer selected (*Index sharing*). It stores nothing per context position, scores nothing, and reads the same number of positions at any prompt length.

On three targets whose sparse attention and architecture differ, DFly2 accepts more tokens per step than DFlash and DSpark trained on the same recipe. Its margin over the stronger one widens from 3.9%–10.5% on general workloads to 10.0%–31.5% on long-context and agentic workloads of up to 128K. At high concurrency on DS-V4.1-Flash, at matched per-user decode speed, that becomes 62% and 81% more throughput at 64K. The gain grows with the context because what DFly2 reads does not, the target's selection being a fixed size that can still fall anywhere. Long contexts are where serving needs speculative decoding most, and they are where following the target's own choices pays most.

[3] Enabled by `index_share_for_mtp_iteration` in the released GLM-5.2 configuration (https://huggingface.co/zai-org/GLM-5.2). The sharing mechanism follows IndexCache (Bai et al., 2026).

References

Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021.

Yushi Bai, Qian Dong, Ting Jiang, Xin Lv, Zhengxiao Du, Aohan Zeng, Jie Tang, and Juanzi Li. IndexCache: Accelerating sparse attention via cross-layer index reuse. *arXiv preprint arXiv:2603.12201*, 2026.

Jungyoub Cha, Hyunjong Kim, and Sungzoon Cho. SpecExtend: A drop-in enhancement for speculative decoding of long sequences. *arXiv preprint arXiv:2505.20776*, 2025.

Charlie Chen, Sebastian Borgeaud, Geoffrey Irving, Jean-Baptiste Lespiau, Laurent Sifre, and John Jumper. Accelerating large language model decoding with speculative sampling. *arXiv preprint arXiv:2302.01318*, 2023.

Jian Chen, Yesheng Liang, and Zhijian Liu. DFlash: Block diffusion for flash speculative decoding. *arXiv preprint arXiv:2602.06036*, 2026.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.

Xin Cheng, Xingkai Yu, Chenze Shao, Jiashi Li, Yunfan Xiong, et al. DSpark: Confidence-scheduled speculative decoding with semi-autoregressive generation. *arXiv preprint arXiv:2607.05147*, 2026.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

DeepSeek-AI. DeepSeek-V3.2: Pushing the frontier of open large language models. *arXiv preprint arXiv:2512.02556*, 2025.

DeepSeek-AI. DeepSeek-V4: Towards highly efficient million-token context intelligence. *arXiv preprint arXiv:2606.19348*, 2026.

Cunxiao Du, Jing Jiang, Yuanchen Xu, Jiawei Wu, Sicheng Yu, Yongqi Li, Shenggui Li, Kai Xu, Liqiang Nie, Zhaopeng Tu, and Yang You. GLIDE with a CAPE: A low-hassle method to accelerate speculative decoding. In *International Conference on Machine Learning (ICML)*, 2024.

Shicheng Fang, Yuxin Wang, XiaoRan Liu, Jiahao Lu, Chuanyuan Tan, Xinchi Chen, Yining Zheng, Xuanjing Huang, and Xipeng Qiu. AgentLongBench: A controllable long benchmark for long-contexts agents via environment rollouts. *arXiv preprint arXiv:2601.20730*, 2026.

Yao Fu, Jiahan Chang, Yueying Li, Ritchie Zhao, Bryce Long, Mohammad Mahdi Kamani, Samkit Jain, Rahul Raman, Tara Safavi, et al. Spin: Shadow predictive indexer for sparse attention. *arXiv preprint arXiv:2610.09025*, 2026.

GLM-5 Team. GLM-5: From vibe coding to agentic engineering. *arXiv preprint arXiv:2602.15763*, 2026.

Fabian Gloeckle, Badr Youbi Idrissi, Baptiste Rozière, David Lopez-Paz, and Gabriel Synnaeve. Better & faster large language models via multi-token prediction. In *International Conference on Machine Learning (ICML)*, 2024.

Hao-Yuan He, Peng-Fei Liu, Chen Chen, Si Shen, Quan He, Runzhou Liu, et al. Efficient speculative decoding with a fixed-cost parallel drafter. *arXiv preprint arXiv:2609.37029*, 2026.

Yuanzhe Hu, Yu Wang, and Julian McAuley. Evaluating memory in LLM agents via incremental multi-turn interactions. In *International Conference on Learning Representations (ICLR)*, 2026. arXiv:2507.05257.

Jianuo Huang, Yaojie Zhang, Qituan Zhang, Hao Lin, Hanlin Xu, et al. Domino: Decoupling causal modeling from autoregressive drafting in speculative decoding. *arXiv preprint arXiv:2605.29707*, 2026.

Naman Jain, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. LiveCodeBench: Holistic and contamination free evaluation of large language models for code. In *International Conference on Learning Representations (ICLR)*, 2025.

Weifan Jiang, Krishna Teja Chitty-Venkata, Megan Flynn, Reed Meyerson, Zhenting Qi, Tianyu Wu, Eldar Kurtic, Minlan Yu, and Alexandre Marques. H-Spec: Parallel speculative decoding without a drafter-side KV cache. *arXiv preprint arXiv:2609.24197*, 2026.

Keller Jordan, Yuchen Jin, Vlado Boza, Jiacheng You, Franz Cesista, Laker Newhouse, and Jeremy Bernstein. Muon: An optimizer for hidden layers in neural networks. https://kellerjordan.github.io/posts/muon/, 2024.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with PagedAttention. In *Proceedings of the 29th Symposium on Operating Systems Principles (SOSP)*, 2023.

Yaniv Leviathan, Matan Kalman, and Yossi Matias. Fast inference from transformers via speculative decoding. In *International Conference on Machine Learning (ICML)*, 2023.

Yuhui Li, Fangyun Wei, Chao Zhang, and Hongyang Zhang. EAGLE-3: Scaling up inference acceleration of large language models via training-time test. *arXiv preprint arXiv:2503.01840*, 2025.

Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *International Conference on Learning Representations (ICLR)*, 2024.

Luxi Lin, Shuang Peng, Rui Ma, Junhao Hua, Shuwei Fan, Zhengda Qin, et al. HyperDFlash: Hyper-connection-aligned block speculative decoding with gated residual reduction. *arXiv preprint arXiv:2606.26744*, 2026.

Changxu Liu and Zhaogeng Li. Adapting tree-structured speculative decoding to DeepSeek-V4 for efficient inference. *arXiv preprint arXiv:2609.24698*, 2026.

Hong Liu, Rui Cen, Junhan Shi, Guangshuo Qin, Jiebin Zhang, Tianyu Liu, et al. AngelSpec: Towards real-world high performance inference with speculative decoding. *arXiv preprint arXiv:2607.25852*, 2026a.

Tianyu Liu, Yuhao Shen, Xinyi Hu, Baolin Zhang, Hengxin Zhang, Jun Dai, et al. When hidden states drift: Can KV caches rescue long-range speculative decoding? *arXiv preprint arXiv:2604.26412*, 2026b.

MAA. American invitational mathematics examination (AIME). https://maa.org/math-competitions/american-invitational-mathematics-examination-aime, 2025.

Chien Van Nguyen, Chaitra Hegde, et al. Orthrus: Memory-efficient parallel token generation via dual-view diffusion. *arXiv preprint arXiv:2605.12825*, 2026.

Shishir G. Patil, Huanzhi Mao, Fanjia Yan, Charlie Cheng-Jie Ji, Vishnu Suresh, Ion Stoica, and Joseph E. Gonzalez. The Berkeley function calling leaderboard (BFCL): From tool use to agentic evaluation of large language models. In *International Conference on Machine Learning (ICML)*, 2025.

Stefano Rando, Luca Romani, Alessio Sampieri, Luca Franco, John Yang, Yuta Kyuragi, Fabio Galasso, and Tatsunori Hashimoto. LongCodeBench: Evaluating coding LLMs at 1M context windows. In *Conference on Language Modeling (COLM)*, 2025.

Peer Rheinboldt, Frédéric Berdoz, and Roger Wattenhofer. TreeFlash: Parallel AR-approximation for faster speculative decoding. *arXiv preprint arXiv:2606.03819*, 2026.

Ranajoy Sadhukhan, Jian Chen, Zhuoming Chen, Vashisth Tiwari, Ruihang Lai, Jinyuan Shi, Ian En-Hsu Yen, Avner May, Tianqi Chen, and Beidi Chen. MagicDec: Breaking the latency-throughput tradeoff for long context generation with speculative decoding. In *International Conference on Learning Representations (ICLR)*, 2025.

Alexander Samarin, Sergei Krutikov, Anton Shevtsov, Sergei Skvortsov, Filipp Fisin, and Alexander Golubev. LK losses: Direct acceptance rate optimization for speculative decoding. *arXiv preprint arXiv:2602.23881*, 2026.

Hanshi Sun, Zhuoming Chen, Xinyu Yang, Yuandong Tian, and Beidi Chen. TriForce: Lossless acceleration of long sequence generation with hierarchical speculative decoding. In *Conference on Language Modeling (COLM)*, 2024a.

Yutao Sun, Li Dong, Yi Zhu, Shaohan Huang, Wenhui Wang, Shuming Ma, Quanlu Zhang, Jianyong Wang, and Furu Wei. You only cache once: Decoder-decoder architectures for language models. *arXiv preprint arXiv:2405.05254*, 2024b.

Tencent Hy Team. Hy4. https://hy.tencent.ai/research/hy4, 2026a.

Tencent Hy Team. Hy4 flash. https://hy.tencent.ai/research/hy4-flash, 2026b.

Zhibin Wang, Ziyu Zhong, Nuo Shen, Yuhang Zhou, Rong Gu, Sheng Zhong, et al. SSV: Sparse speculative verification for efficient LLM inference. *arXiv preprint arXiv:2605.19893*, 2026. Version 1 appeared as SpecSA.

Tianyu Wu, Yu Yao, Zhenting Qi, Han Zheng, Zhuohan Wang, Haoran Ma, Lawrence Liao, Himabindu Lakkaraju, Ju Li, and Yilun Du. D-PACE: Dynamic position-aware cross-entropy for parallel speculative drafting. *arXiv preprint arXiv:2605.18810*, 2026.

Jiaming Xu, Hanzhen Wang, Tianlang Zhao, Jiewen Xiao, Chengze Yuan, et al. DLite: History-free speculative decoding for agentic AI. *arXiv preprint arXiv:2610.23787*, 2026.

Penghui Yang, Cunxiao Du, Fengzhuo Zhang, Haonan Wang, Tianyu Pang, Chao Du, and Bo An. LongSpec: Long-context lossless speculative decoding with efficient drafting and verification. *arXiv preprint arXiv:2502.17421*, 2025.

Jingyang Yuan, Huazuo Gao, Damai Dai, Junyu Luo, Liang Zhao, Zhengyan Zhang, Zhenda Xie, Y. X. Wei, Lean Wang, et al. Native sparse attention: Hardware-aligned and natively trainable sparse attention. *arXiv preprint arXiv:2502.11089*, 2025.

Yikang Yue, Yuqi Xue, and Jian Huang. Vegas: Self-speculative decoding with verification-guided sparse attention. In *International Conference on Machine Learning (ICML)*, 2026. arXiv:2602.07223; version 1 appeared as SpecAttn.

Jiebin Zhang, Zhenghan Yu, Song Liu, Eugene J. Yu, Zheng Li, et al. DFlare: Scaling up draft capacity for block diffusion speculative decoding. *arXiv preprint arXiv:2606.02091*, 2026.

Yilong Zhao, Jiaming Tang, Kan Zhu, Zihao Ye, Chi-Chih Chang, Chaofan Lin, et al. Accelerating large-scale reasoning model inference: Self-speculative decoding with sparse attention. In *Proceedings of Machine Learning and Systems (MLSys)*, 2026.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, et al. Judging LLM-as-a-judge with MT-Bench and Chatbot Arena. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.