arXiv:2611.34051v1 [cs.CL] 25 Nov 2026

Stable Gated DeltaNet-2: Understanding and Mitigating Recurrent State Instability

Xichong Zhang[1,2*] Xixuan Hao[3*] Zhichen Liu[4*] Le Zhou[5*]
Zeyu Zhang[1] Boyu Chen[1] Mingfei Lu[1] Zehao Lin[1] Bo Tang[1]
Yonggang Zhang[6†] Linyun Liu[1†] Feiyu Xiong[1] Jianlin Su[7] Zhiyu Li[1]

[1]MemTensor [2]USTC [3]HKUST(GZ) [4]SUSTech [5]SJTU [6]HKUST [7]kexue.fm

[*]Equal contribution. [†]Corresponding authors.

Abstract

Linear attention offers an efficient solution for long-sequence modeling via linear complexity and fixed-size recurrent states. Gated DeltaNet-2 (GDN-2) introduces independent channel-wise control over erasure and writing, enhancing state-update expressiveness and in-context retrieval, yet it suffers from training failures during long-context training. We trace these failures to the asymmetry between reading and modification directions during erasure, which yields non-normal erase operators whose directional gains accumulate across tokens, triggering exponential state growth even without new writes. To resolve this, we propose **Stable GDN-2**, employing symmetric erasure to ensure non-expansive state propagation and decay-regulated residual writing to preserve writing flexibility. Theoretically, we establish an $O(\sqrt{T})$ upper bound on recurrent state growth while retaining chunkwise parallel computation. Experiments on 1.6B parameter models show that **Stable GDN-2** achieves stable training alongside improved long-context retrieval. Compared to GDN-2, it boosts the 128K average RULER score by 6.80 points in the recurrent setting and average accuracy across nine downstream tasks by 0.62 points in the hybrid setting.

Contact:
Xichong Zhang <yyu18@mail.ustc.edu.cn>,
Xixuan Hao <xhao390@connect.hkust-gz.edu.cn>,
Zhichen Liu <liuzc2024@mail.sustech.edu.cn>,
Le Zhou <zhoule116@sjtu.edu.cn>,
Yonggang Zhang <zhangyg@ust.hk>,
Linyun Liu <liuly@memtensor.cn>,
Jianlin Su <bojone@spaces.ac.cn>,
Zhiyu Li <lizy@memtensor.cn>.

1. Introduction

"A small leak will sink a great ship."

– Benjamin Franklin

The growing demand for long-context understanding and extended generation (Sun et al., 2025; Zhang et al., 2026) has made efficient sequence modeling an increasingly important challenge for large language models. To address the computational and memory bottlenecks of standard softmax attention (Llama, 2023; Lu et al., 2026), linear attention (Katharopoulos et al., 2020; Sun et al., 2023) has emerged as an efficient alternative that compresses historical information into a fixed-size recurrent state, thereby enabling linear-time sequence processing and constant memory during autoregressive decoding. Building on this approach, recent advances in gating mechanisms and delta-rule-based updates (Yang et al., 2024a; 2025; 2024b) have substantially enhanced the expressiveness of recurrent memory, enabling more fine-grained control over information retention and modification. Gated DeltaNet-2 (GDN-2) (Hatamizadeh et al., 2026) further advances this line of work by replacing the shared scalar update gate in Gated DeltaNet (GDN) (Yang et al., 2025) and Kimi Delta Attention (KDA) (Team et al., 2025) with independent channel-wise erase and write gates, supporting more expressive recurrent state updates and improved in-context retrieval. Specifically, the erase gate reweights the key coordinates used to read content for erasure, while the modification is applied along the original key, allowing more flexible selection of content for erasure.

However, we argue that **the asymmetry between the reading and modification directions during erasure in GDN-2 constitutes an inherent source of *instability* in its recurrent dynamics**, potentially triggering catastrophic state explosion and ultimately causing training failures at long context. Our analysis shows that the asymmetric coupling between the erase gate and the key can introduce shear, preserving some state components while inducing additional

components along other directions, as illustrated in Figure 1(a). The resulting non-normal erasure operator can amplify certain state directions despite all its eigenvalues having magnitude at most one. When amplified output directions of earlier transitions align with expanding input directions of subsequent transitions, these gains can accumulate across the sequence, allowing the recurrent state to continue growing to the point of numerical overflow, even after new writes cease. As Franklin famously remarked, "A small leak will sink a great ship."

To thoroughly investigate the underlying mechanisms of this instability, we examine recurrent state dynamics and their numerical and representational consequences in practical training. At a context length of 64K, we observe numerical failures across both pretraining and post-training settings, affecting recurrent as well as hybrid GDN-2 architectures. More importantly, numerical failures represent only the most visible manifestation of a broader issue: even in runs that complete training with finite losses, individual entries of the recurrent state can exceed 10^{36}, indicating substantial internal state amplification. Replay analysis shows that amplification from asymmetric erasure can concentrate state energy in a dominant direction and sharply reduce effective capacity before numerical failure, with little warning in the training loss. These observations highlight a fundamental limitation of loss-based monitoring and motivate a closer investigation of the underlying recurrent state dynamics that govern memory propagation over long sequences.

These findings motivate **Stable GDN-2**, which controls recurrent state amplification through symmetric erasure while retaining channel-wise gating. Distributing the erase gate symmetrically across the two key factors aligns the reading and modification directions, ensuring non-expansive propagation under the standard key normalization and gate constraints, as illustrated in Figure 1(b). To complement the resulting gated write, we introduce a residual writing path whose amplitude is regulated by decay. This design supports additional writing in channels where the erase gate is small while controlling the energy introduced by the residual. The update retains an input-dependent affine structure compatible with chunk-wise parallel computation.

We evaluate **Stable GDN-2** through 64K-context pretraining in recurrent and hybrid architectures, comparing language modeling, downstream performance, and long-context retrieval against GDN-2 and existing stabilization methods. The results show that controlling recurrent state amplification while retaining channelwise erase gating supports improved retrieval and downstream performance, alongside competitive language modeling quality.

In summary, our contributions lie in the following aspects:

- *Understanding recurrent state instability in GDN-2.* We characterize recurrent state instability in GDN-2 theoretically and empirically, showing that amplification can persist even without further writes. To the best of our knowledge, this is the first work to empirically identify training instability in GDN-2 and trace its origin to asymmetric erasure.

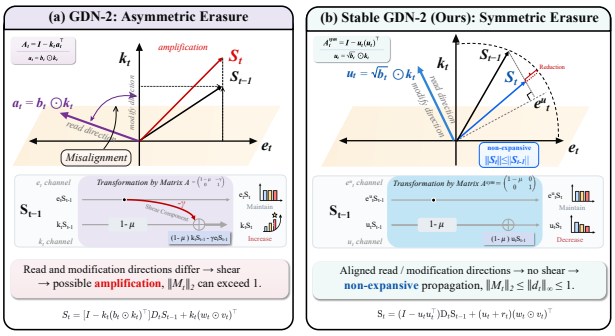

Figure 1. Geometric comparison between GDN-2 and **Stable GDN-2**.

- *A principled stabilization mechanism for GDN-2.* We propose **Stable GDN-2**, which employs symmetric erasure to guarantee non-expansive state propagation while retaining channel-wise gating. We further introduce a decay-regulated residual writing path to preserve writing flexibility without compromising the stability guarantee. The resulting formulation supports chunkwise parallel computation.

- *Comprehensive empirical evaluation.* We evaluate **Stable GDN-2** against GDN-2 and existing stabilization methods using recurrent and hybrid models pretrained on 64K-token sequences. Compared with GDN-2, **Stable GDN-2** improves the average RULER score at 128K by 6.80 points in the recurrent setting and average accuracy across nine downstream tasks by 0.62 percentage points in the hybrid setting, while maintaining competitive language modeling performance.

2. Preliminary

We consider a single attention head and omit head and layer indices. At position t, the query and key are $\boldsymbol{q}_t, \boldsymbol{k}_t \in \mathbb{R}^{d_k}$, the value is $\boldsymbol{v}_t \in \mathbb{R}^{d_v}$, and the recurrent state is $\mathbf{S}_t \in \mathbb{R}^{d_k \times d_v}$. All vectors are column vectors. We use $\odot$ for elementwise multiplication and $\mathrm{Diag}(\boldsymbol{x})$ for the diagonal matrix with diagonal $\boldsymbol{x}$.

2.1. Linear Attention and the Gated Delta Rule

Linear attention accumulates historical key–value associations in a fixed-size matrix state and retrieves them using the current query (Katharopoulos et al., 2020). Following recent delta-rule models (Yang et al., 2024b; 2025), we adopt the

linear-kernel formulation without an attention normalizer:

$$\mathbf{S}_t = \mathbf{S}_{t-1} + \boldsymbol{k}_t \boldsymbol{v}_t^\top, \quad \boldsymbol{o}_t = \mathbf{S}_t^\top \boldsymbol{q}_t. \tag{1}$$

Each outer product adds a new association to the memory, whose dimensions are independent of sequence length. This additive update does not explicitly revise the content already associated with the incoming key.

The delta rule corrects an existing association using the residual between the incoming value and the value retrieved from memory (Schlag et al., 2021; Yang et al., 2024b). For a scalar update gate $\beta_t \in [0, 1]$, DeltaNet applies

$$\begin{aligned} \mathbf{S}_t &= \mathbf{S}_{t-1} + \beta_t \boldsymbol{k}_t \big(\boldsymbol{v}_t - \mathbf{S}_{t-1}^\top \boldsymbol{k}_t\big)^\top \\ &= \big(\mathbf{I} - \beta_t \boldsymbol{k}_t \boldsymbol{k}_t^\top\big)\mathbf{S}_{t-1} + \beta_t \boldsymbol{k}_t \boldsymbol{v}_t^\top. \end{aligned} \tag{2}$$

Here, $\mathbf{S}_{t-1}^\top \boldsymbol{k}_t$ is the value predicted by the current memory. We consider ℓ_2-normalized keys, $\|\boldsymbol{k}_t\|_2 = 1$, throughout the delta-rule recurrences below. The expanded form separates the removal of retrieved content from the insertion of the new value. Both operations are controlled by the same scalar gate β_t. Gated DeltaNet introduces a forget gate that decays the existing memory before applying the delta-rule correction (Yang et al., 2025). The resulting update takes the form

$$\mathbf{S}_t = \big(\mathbf{I} - \beta_t \boldsymbol{k}_t \boldsymbol{k}_t^\top\big)\mathbf{D}_t \mathbf{S}_{t-1} + \beta_t \boldsymbol{k}_t \boldsymbol{v}_t^\top. \tag{3}$$

Gated DeltaNet uses a scalar decay, $\mathbf{D}_t = d_t \mathbf{I}$ with $d_t \in (0, 1]$. Kimi Delta Attention instead uses $\mathbf{D}_t = \mathrm{Diag}(\mathbf{d}_t)$, where $\mathbf{d}_t \in (0, 1]^{d_k}$ contains a separate decay factor for each key channel (Team et al., 2025). Setting $\mathbf{D}_t = \mathbf{I}$ recovers DeltaNet. Although KDA provides finer control over forgetting, its erase and write operations remain tied to the same scalar update gate.

2.2. Gated DeltaNet-2

GDN-2 replaces the shared scalar update gate with separate channel-wise gates for erasing and writing (Hatamizadeh et al., 2026). Let $\boldsymbol{b}_t \in (0, 1)^{d_k}$ and $\boldsymbol{w}_t \in (0, 1)^{d_v}$ denote the erase and write gates, respectively.

Retaining KDA's diagonal decay, GDN-2 updates the state as

$$\mathbf{S}_t = \big[\mathbf{I} - \boldsymbol{k}_t(\boldsymbol{b}_t \odot \boldsymbol{k}_t)^\top\big]\mathbf{D}_t \mathbf{S}_{t-1} + \boldsymbol{k}_t(\boldsymbol{w}_t \odot \boldsymbol{v}_t)^\top. \tag{4}$$

The erase gate weights the key coordinates used to retrieve old content for removal, while the write gate scales the value coordinates committed to memory. These gates are input-dependent and parameterized independently.

GDN-2 allows the erase operation to read and modify the state along different directions. In GDN and KDA, both the retrieval of old content and the subsequent modification use the key $\boldsymbol{k}_t$. In GDN-2, retrieval uses $\boldsymbol{b}_t \odot \boldsymbol{k}_t$, while the modification is still applied along $\boldsymbol{k}_t$. These vectors need not be parallel, so the erase matrix $\mathbf{I} - \boldsymbol{k}_t(\boldsymbol{b}_t \odot \boldsymbol{k}_t)^\top$ need not be symmetric. This structural change arises from the channel-wise erase gate; the independent write gate determines the content injected into the state. The recurrence reduces to KDA when $\boldsymbol{b}_t = \beta_t \mathbf{1}_{d_k}$ and $\boldsymbol{w}_t = \beta_t \mathbf{1}_{d_v}$, and further reduces to Gated DeltaNet when $\mathbf{D}_t = d_t \mathbf{I}$.

GDN-2 evaluates the same recurrence through chunkwise parallel computation during training (Hatamizadeh et al., 2026). For chunks of length C, let $\mathbf{S}_{[n]} = \mathbf{S}_{nC}$ denote the state entering chunk n, which contains positions $nC + 1, \ldots, (n + 1)C$; its exit state is $\mathbf{S}_{[n+1]} = \mathbf{S}_{(n+1)C}$. Token interactions within each chunk are computed through parallel matrix operations, while the boundary states carry information between successive chunks.

3. Understanding State Instability in GDN-2

In GDN-2, the asymmetric coupling between the channel-wise erase gate and the key yields a generally non-normal erase operator, breaking the non-expansiveness guaranteed by symmetric erasure in GDN and KDA. Across time-varying transitions, directional gains from this non-normality can accumulate when amplified state directions align with subsequent expanding directions, allowing stored memory to grow even after new writes cease. In the following, we theoretically formalize this mechanism (Section 3.1) and show that it manifests in long-context training as both sudden numerical failures and extreme state excursions hidden beneath normal loss trajectories (Section 3.2).

3.1. Theoretical Analysis

GDN-2 employs channel-wise erase gates to offer fine-grained control over recurrent states. Erasure reads the decayed recurrent state $\mathbf{D}_t \mathbf{S}_{t-1}$ using the gated key $\mathbf{b}_t \odot \mathbf{k}_t$ and applies the resulting modification along the original key $\mathbf{k}_t$. The stability question is whether this asymmetric modification can amplify the recurrent state even when every gate entry lies in $(0, 1)$.

The nontrivial action of erasure is confined to the subspace spanned by the key and its gated counterpart. Omitting the time index, let $\mathbf{a} = \mathbf{b} \odot \mathbf{k}$ and $\mathbf{A} = \mathbf{I} - \mathbf{k}\mathbf{a}^\top$. For $\|\mathbf{k}\|_2 = 1$, choose $\mathbf{e}_\parallel = \mathbf{k}$ and decompose

$$\begin{aligned} &\mathbf{a} = \mu \mathbf{e}_\parallel + \gamma \mathbf{e}_\perp, \\ &\mu = \mathbf{k}^\top \mathbf{a}, \qquad \gamma = \|\mathbf{a} - \mu \mathbf{k}\|_2, \end{aligned} \tag{5}$$

where $\mathbf{e}_\perp$ is a unit vector orthogonal to $\mathbf{e}_\parallel$; when $\gamma = 0$, any such vector may be chosen. Expressing $\mathbf{A}\mathbf{e}_\parallel = (1 - \mu)\mathbf{e}_\parallel$ and $\mathbf{A}\mathbf{e}_\perp = \mathbf{e}_\perp - \gamma \mathbf{e}_\parallel$ in this orthonormal basis gives the

two-dimensional representation

$$\mathbf{T} = \begin{pmatrix} 1-\mu & -\gamma \\ 0 & 1 \end{pmatrix}, \tag{6}$$

with the identity acting on the orthogonal complement. Here, μ controls attenuation along the key, while γ measures the shear. It means that an orthogonal component is retained while contributing an additional component along the key. This coupling depends jointly on the erase gate and the key.

Nonzero shear γ creates directions in which erasure amplifies the recurrent state. The following theorem gives the exact maximum gain of a single erasure step.

Theorem 3.1 (Single-step amplification in GDN-2)**.** *Let* $\mathbf{A} = \mathbf{I} - \mathbf{k}(\mathbf{b} \odot \mathbf{k})^\top$*, where* $d_k \geq 2$*,* $\|\mathbf{k}\|_2 = 1$*, and* $\mathbf{b} \in [0,1]^{d_k}$*. With* μ *and* γ *defined in Eq. (5), its operator spectral norm is*

$$\begin{aligned} \|\mathbf{A}\|_2 &= \sigma_{\max}(\mathbf{A}) \\ &= \frac{\sqrt{\mu^2+\gamma^2} + \sqrt{(2-\mu)^2+\gamma^2}}{2}. \end{aligned} \tag{7}$$

In particular, $\|\mathbf{A}\|_2 = 1$ *if and only if* $\gamma = 0$*, equivalently,* $\mathbf{a}$ *and* $\mathbf{k}$ *are collinear. Otherwise,* $\|\mathbf{A}\|_2 > 1$*.*

The proof is given in Appendix A.1. The spectral norm $\|\mathbf{A}\|_2$ measures the maximum energy amplification factor of the recurrent state $\|\mathbf{AS}\|_F / \|\mathbf{S}\|_F$. Accordingly, Theorem 3.1 shows that erasure alone can increase the energy of recurrent state, defined as $E(\mathbf{S}) = \|\mathbf{S}\|_F^2$, even without any new write. For GDN and KDA, scalar erase gating gives $\mathbf{a} = \beta\mathbf{k}$ and $\gamma = 0$, recovering non-expansive symmetric erasure. As a simple example of the geometry in Figure 1, consider a single-column recurrent state $\mathbf{S}_t = \mathbf{e}_\perp$. When $\gamma > 0$, erasure gives

$$\begin{aligned} \mathbf{AS}_t &= \mathbf{e}_\perp - \gamma \mathbf{e}_\parallel, \\ E(\mathbf{AS}_t) &= 1 + \gamma^2 > E(\mathbf{S}_t) = 1. \end{aligned} \tag{8}$$

This amplification can persist in the full transition $\mathbf{M} = \mathbf{AD}$. In particular, an admissible scalar decay $\mathbf{D} = d\mathbf{I}$ with $1/\|\mathbf{A}\|_2 < d < 1$ gives $\|\mathbf{M}\|_2 = d\|\mathbf{A}\|_2 > 1$.

The analysis of a single update shows that a full transition $\mathbf{M}_t = \mathbf{A}_t\mathbf{D}_t$ can increase recurrent state energy even without a new write. Along a token sequence, the recurrent state evolves through successive transitions that depend on the input. Their composition determines whether this amplification can persist. Suppose no further content is written after token s. The recurrent state at token $t > s$ then satisfies

$$\mathbf{S}_t = \mathbf{\Phi}_{t:s+1}\mathbf{S}_s, \tag{9}$$

where $\mathbf{\Phi}_{t:s+1} = \mathbf{M}_t \cdots \mathbf{M}_{s+1}$. When an amplified output direction of $\mathbf{M}_j$ aligns with an expanding input direction of $\mathbf{M}_{j+1}$, the transition at the next token can amplify it further. Repeated alignment across tokens can therefore allow gains to accumulate, posing a risk of sustained recurrent state energy growth even after writing has ceased.

Such alignment can be realized by an admissible sequence of GDN-2 updates, yielding unbounded recurrent state growth after a single finite write.

Corollary 3.2 (Unbounded growth without further writes)**.** *There exist normalized keys, together with erase, write, and decay gates whose entries lie in* $(0,1)$*, such that the GDN-2 recurrence, initialized at* $\mathbf{S}_0 = \mathbf{0}$*, satisfies*

$$\lim_{t\to\infty} \|\mathbf{S}_t\|_F = \infty, \tag{10}$$

after a finite nonzero write at the first token, even with no writes at subsequent tokens.

Appendix A.2 provides an explicit example that alternates two admissible transitions with $\mathbf{M}_b = \mathbf{M}_a^\top$ and $\lambda = \sigma_{\max}(\mathbf{M}_a)^2 > 1$. Starting from an initial write along a leading right singular direction of $\mathbf{M}_a$, each pair returns the recurrent state to the same direction and multiplies its norm by λ, yielding exponential growth without further writes.

In summary, the asymmetric coupling between the erase gate and the key in GDN-2 can yield a non-normal erase operator with expanding directions, breaking the non-expansiveness of symmetric erasure in GDN and KDA. When amplified output directions align with expanding input directions at subsequent tokens, the gains can accumulate, allowing sustained growth in recurrent state energy without further writes and creating a risk of numerical overflow in finite precision arithmetic. The analysis establishes the existence of such trajectories, while the occurrence and extent of amplification in practical training depend on how the transitions induced by the token sequence and learned parameters interact with the recurrent state.

3.2. Empirical Characterization

To examine how the amplification mechanism identified above manifests in training, we study GDN-2 under pretraining, post-training, and post-training after pretraining, using Pure Linear and Hybrid architectures with four combinations of data and initialization seeds. Pretraining uses 64K sequences to examine accumulated gains over long uninterrupted recurrent state propagation. We record a numerical failure when the training loss becomes NaN (not a number) and prevents further training. Across the two architectures, Table 1 reports failures in 3 of 8 pretraining runs and 4 of 8 post-training runs; all five post-training stages launched after successful pretraining complete. Detailed evaluation protocols and failure accounting are provided in Appendix B.1.

Replaying one failed pretraining run traces its loss NaN to

Table 1. Training failures across regimes and architectures. The designated trajectories and additional repeats are detailed in Appendix B.1.

Architecture		Pretraining	Post-training	Pretraining → Post-training
GDN-2	Recurrent Model	1/4	1/4	0/3
	Hybrid Model	2/4	3/4	0/2

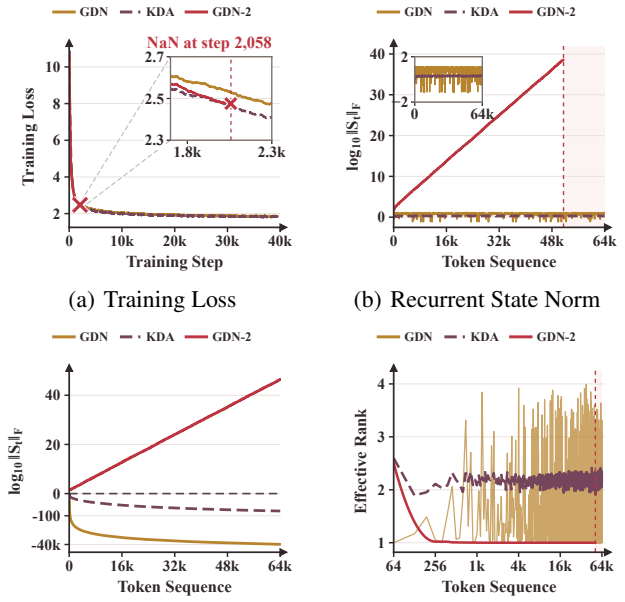

(a) Training Loss (b) Recurrent State Norm

(c) Propagation without Writes (d) Recurrent State Capacity

Figure 2. The trace of a training failure.

recurrent state amplification that persists without further writes. Figure 2(a) shows GDN-2 failing at step 2,058, alongside the training trajectories of GDN and KDA. On the triggering sequence, Figure 2(b) shows sustained growth in the GDN-2 recurrent state norm, with overflow observed at token 52,736 and nonfinite values subsequently reaching the logits and loss. GDN and KDA do not exhibit the same growth in the corresponding replay. To isolate propagation, Figure 2(c) retains the first 64 token updates and sets subsequent write terms to zero, while holding the extracted keys, gates, and decay factors fixed. The recurrent state then follows Eq. (9) with $s = 64$. GDN and KDA decay after writes cease, whereas GDN-2 continues to grow. This trajectory exhibits the mechanism established in Corollary 3.2 on a real token sequence. The non-normal erase operator A_t induced by the asymmetric coupling between the erase gate and the key permits expanding directions; when amplified output directions align with expanding input directions at subsequent tokens, gains can accumulate and drive recurrent state growth even without further writes.

This amplification is accompanied by a loss of effective representational dimensions before numerical overflow. Following Nazari & Rusch (2026), we measure the effective rank using $(\mathbf{S}_t) = \|\mathbf{S}_t\|_F^2/\|\mathbf{S}_t\|_2^2$ to characterize the effective capacity of the recurrent state. In Figure 2(d), the capacity of recurrent state using GDN-2 rapidly approaches one and remains near this value. The concentration of energy into a single dominant direction reflects degraded utilization of the recurrent state's representational dimensions.

Extreme recurrent state magnitudes also occur during runs that eventually complete training. Table 4 in Appendix B.3 shows that the Pure Linear model with Seed 2 completes pretraining even though the largest recorded absolute recurrent state entry during training reaches 3.157×10^{36}. Such values remain within the finite representable range of FP32, so substantial amplification need not immediately produce loss NaN. These observations connect asymmetric erasure and accumulated transition gains to recurrent state growth in actual training, with consequences ranging from numerical overflow to extreme values along otherwise successful training trajectories. They motivate directly controlling the gain of recurrent state propagation.

4. Stable GDN-2

The analysis above shows that the asymmetric coupling between the channel-wise erase gate and the key in GDN-2 can amplify the recurrent state. *Can we retain the fine-grained control of erase while restoring non-expansive recurrent state propagation?* To this end, we propose **Stable GDN-2**, which restores non-expansive propagation through symmetric erasure and introduces a residual write, regulated by decay, to supplement writing suppressed by symmetric gating.

4.1. Symmetric Erasure and Residual Writing

The channel-wise erase gate in GDN-2 provides fine-grained control over erasure, but its asymmetric coupling with the key in $\mathbf{A}_t$ (Eq. (4)) introduces a risk of recurrent state amplification. To retain channel-wise erase gating while aligning the reading and modification directions, we split the erase gate into two square-root factors and apply one to each key factor. Defining $\mathbf{u}_t = \sqrt{\mathbf{b}_t} \odot \mathbf{k}_t$ gives the symmetric erase operator

$$\begin{aligned} \mathbf{A}_t^{\mathrm{sym}} &= \mathbf{I} - (\sqrt{\mathbf{b}_t} \odot \mathbf{k}_t)(\sqrt{\mathbf{b}_t} \odot \mathbf{k}_t)^\top \\ &= \mathbf{I} - \mathbf{u}_t(\mathbf{u}_t)^\top. \end{aligned} \tag{11}$$

This preserves the diagonal attenuation of the original GDN-2 operator, $(\mathbf{A}_t^{\mathrm{sym}})_{ii} = (\mathbf{A}_t)_{ii} = 1 - b_{t,i}k_{t,i}^2$.

Using $\mathbf{u}_t$ for reading, modification, and writing, we can rewrite the delta update in this symmetric formulation as

$$\mathbf{S}_t = \mathbf{D}_t\mathbf{S}_{t-1} + \mathbf{u}_t\left[\mathbf{z}_t - (\mathbf{D}_t\mathbf{S}_{t-1})^\top\mathbf{u}_t\right]^\top, \tag{12}$$

where $\mathbf{z}_t = \mathbf{w}_t \odot \mathbf{v}_t$ is the gated value. This retains the delta update structure and ensures non-expansive recurrent state

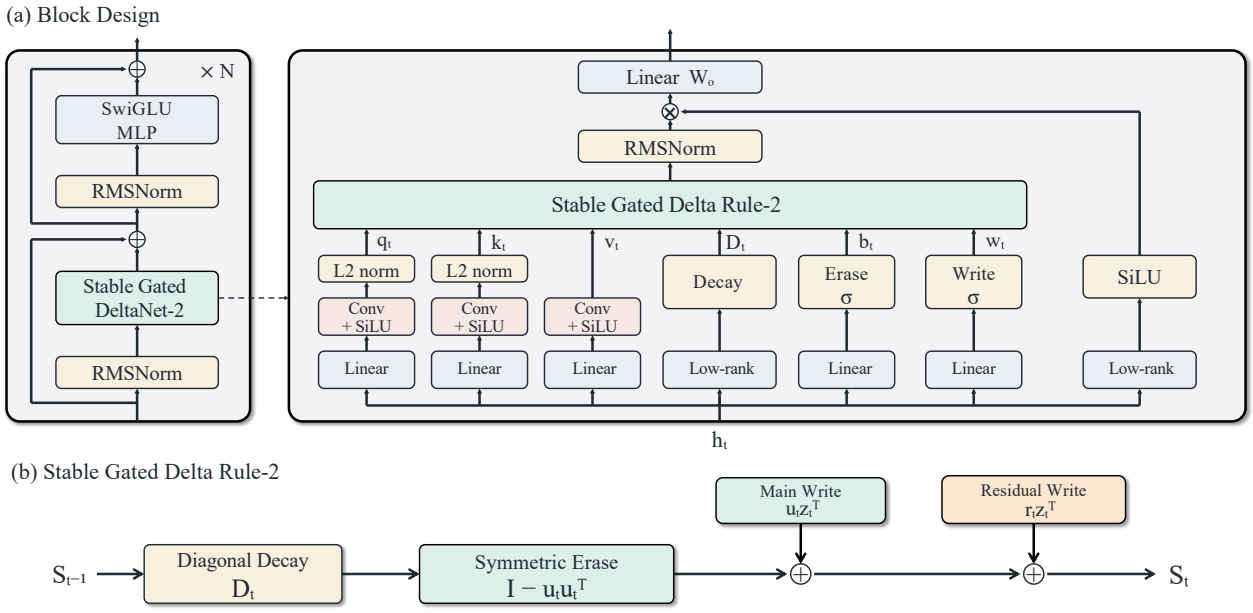

Figure 3. **Stable GDN-2** architecture and recurrent state update.

propagation. However, as $b_{t,i}$ approaches zero, both erasure and the base write in channel i vanish. This tightly couples erasing stored content with injecting new information, restricting the ability to write while suppressing erasure.

To supplement the suppressed writes, we introduce a residual along $\mathbf{k}_t - \mathbf{u}_t$ and regulate its amplitude by decay. Decay retains a fraction $d_{t,i}^2$ of the recurrent state energy in channel i. We therefore scale the residual by the square root of the complementary fraction,

$$\mathbf{r}_t = \sqrt{1 - \mathbf{d}_t^2} \odot (\mathbf{k}_t - \mathbf{u}_t). \tag{13}$$

As $d_{t,i}$ approaches one, the residual write vanishes in that channel; stronger decay permits a larger additional write. The complete update becomes

$$\mathbf{S}_t = \mathbf{A}_t^{\mathrm{sym}} \mathbf{D}_t \mathbf{S}_{t-1} + (\mathbf{u}_t + \mathbf{r}_t)\mathbf{z}_t^\top. \tag{14}$$

Beyond non-expansive propagation, this design yields a tighter bound on recurrent state energy under continued writing. Section 4.2 establishes that its growth is controlled by cumulative write energy, $\sum_{t=1}^{T} \|\mathbf{z}_t\|_2^2$.

Figure 3 shows the architecture and the two write contributions. We retain the gate parameterization of GDN-2 (Hatamizadeh et al., 2026), using separate sigmoid projections for the erase $\mathbf{b}_t$ and write gates $\mathbf{w}_t$ and the same exponential decay parameterization with rates constrained by softplus. Queries and keys follow the original ℓ_2 normalization; $\mathbf{u}_t$ and $\mathbf{r}_t$ retain their gate-dependent magnitudes without further normalization.

4.2. Stability Guarantees

Symmetric erasure removes the expanding directions identified in Theorem 3.1. For the transition $\mathbf{M}_t = \mathbf{A}_t^{\mathrm{sym}} \mathbf{D}_t$ in Eq. (14), this prevents recurrent state amplification in the absence of new writes.

Proposition 4.1 (Non-expansive transitions without writes)**.** *The symmetric erase operator in **Stable GDN-2** satisfies* $\|\mathbf{A}_t^{\mathrm{sym}}\|_2 \le 1$*, and the corresponding transition obeys*

$$\|\mathbf{M}_t\|_2 \le \|\mathbf{d}_t\|_\infty \le 1. \tag{15}$$

When the write term at token t is zero, the recurrent state therefore satisfies $\|\mathbf{S}_t\|_F \le \|\mathbf{d}_t\|_\infty \|\mathbf{S}_{t-1}\|_F \le \|\mathbf{S}_{t-1}\|_F$.

Proof. Key normalization and the erase gate parameterization give $\|\mathbf{u}_t\|_2^2 = \sum_i b_{t,i} k_{t,i}^2 \le \|\mathbf{k}_t\|_2^2 \le 1$. Hence $0 \preceq \mathbf{A}_t^{\mathrm{sym}} \preceq \mathbf{I}$. As in the analysis of symmetric erasure in GDN and KDA in Section 3.1, the erase operator has spectral norm at most one. By submultiplicativity, $\|\mathbf{M}_t\|_2 \le \|\mathbf{A}_t^{\mathrm{sym}}\|_2 \|\mathbf{D}_t\|_2 \le \|\mathbf{d}_t\|_\infty \le 1$. With no new write, $\mathbf{S}_t = \mathbf{M}_t \mathbf{S}_{t-1}$, which gives the stated Frobenius norm bound. □

Consequently, if writes cease after position s, then $\|\mathbf{S}_t\|_F \le \|\mathbf{\Phi}_{t:s+1}\|_2 \|\mathbf{S}_s\|_F \le \|\mathbf{S}_s\|_F$ for every $t > s$, using the propagation notation from Eq. (9). With continued writing, the recurrent state can grow through the accumulation of new content, with each contribution propagated with gain at most one. Expanding Eq. (14) gives $\mathbf{S}_T =$

$\mathbf{\Phi}_{T:1}\mathbf{S}_0 + \sum_{t=1}^{T} \mathbf{\Phi}_{T:t+1}(\mathbf{u}_t + \mathbf{r}_t)\mathbf{z}_t^\top$, with empty products equal to $\mathbf{I}$. Together with $|(\mathbf{u}_t + \mathbf{r}_t)_i| \le |k_{t,i}|$, this yields

$$\begin{aligned}\|\mathbf{S}_T\|_F &\le \|\mathbf{S}_0\|_F + \sum_{t=1}^{T} \|(\mathbf{u}_t + \mathbf{r}_t)\mathbf{z}_t^\top\|_F \\ &\le \|\mathbf{S}_0\|_F + \sum_{t=1}^{T} \|\mathbf{k}_t\mathbf{z}_t^\top\|_F.\end{aligned} \tag{16}$$

Appendix C.1 gives the coordinatewise argument. This bound has the same additive form as recurrent state accumulation in the Linear Transformer of Katharopoulos et al. (2020). For $\mathbf{S}_0 = \mathbf{0}$ and $\|\mathbf{v}_t\|_2 \le V$, it gives $\|\mathbf{S}_T\|_F \le TV$, since $\|\mathbf{z}_t\|_2 = \|\mathbf{w}_t \odot \mathbf{v}_t\|_2 \le \|\mathbf{v}_t\|_2$.

Decay regulation tightens this bound from $O(T)$ to $O(\sqrt{T})$ by jointly controlling propagation and writing.

Proposition 4.2 (Recurrent state norm bound)**.** *Consider the **Stable GDN-2** recurrence in Eq. (14), initialized at $\mathbf{S}_0 = \mathbf{0}$. If $\|\mathbf{v}_t\|_2 \le V$ for all tokens, with a constant $V > 0$, then for every $T \ge 1$,*

$$\|\mathbf{S}_T\|_F \le \sqrt{\frac{3T}{2}}\, V. \tag{17}$$

The proof in Appendix C.2 jointly bounds propagation and writing, without requiring decay factors to remain uniformly below one. Symmetric erasure eliminates the exponential amplification through transition products identified in Section 3.1, while decay regulation further provides an $O(\sqrt{T})$ bound on the recurrent state norm under bounded values and zero initialization. Removing decay regulation preserves non-expansive transitions but need not admit an $O(\sqrt{T})$ bound with a constant independent of the gate values, as shown in Appendix C.3.

4.3. Chunkwise Parallel Training

Now we show that **Stable GDN-2** can keep the same chunkwise structure as GDN and KDA. Its transition remains Diagonal-Plus-Low-Rank (DPLR), while the residual path contributes an additional input-dependent write. Defining the modification vector $\boldsymbol{\delta}_t = \mathbf{z}_t - (\mathbf{D}_t\mathbf{S}_{t-1})^\top\mathbf{u}_t$, then Eq. (14) becomes

$$\mathbf{S}_t = \mathbf{D}_t\mathbf{S}_{t-1} + \mathbf{u}_t\boldsymbol{\delta}_t^\top + \mathbf{r}_t\mathbf{z}_t^\top. \tag{18}$$

Each token introduces a single state-dependent correction $\boldsymbol{\delta}_t$, while the residual write $\mathbf{r}_t\mathbf{z}_t^\top$ is fully determined by the input. The residual contributions therefore enter the correction equations as known terms on the right-hand side. Since each correction depends only on earlier corrections, these equations form a single unit lower triangular system over the token positions within a chunk.

Consider a chunk of C tokens, with local indices $i = 1, \dots, C$ and initial state $\mathbf{S}_{\text{in}}$. Following the decay factorization used in GDN-2, write $\mathbf{d}_i = \exp(\mathbf{g}_i)$ and define the cumulative decay vector $\boldsymbol{\gamma}_i = \exp\left(\sum_{j=1}^{i} \mathbf{g}_j\right)$, with $\boldsymbol{\gamma}_0 = \mathbf{1}$, where exponentials are applied elementwise. Let $\Delta, Z \in \mathbb{R}^{C\times d_v}$ contain rows $\boldsymbol{\delta}_i^\top$ and $\mathbf{z}_i^\top$, respectively, and define the decay-weighted factors

$$\begin{aligned}E_{i,:} &= (\boldsymbol{\gamma}_i \odot \mathbf{u}_i)^\top, \\ \bar{U}_{i,:} &= (\mathbf{u}_i \oslash \boldsymbol{\gamma}_i)^\top, \\ \bar{R}_{i,:} &= (\mathbf{r}_i \oslash \boldsymbol{\gamma}_i)^\top.\end{aligned} \tag{19}$$

where $\oslash$ denotes element-wise division. The causal interactions are

$$\begin{aligned}L &= \operatorname{tril}\left(E(\bar{U})^\top, -1\right), \\ J &= \operatorname{tril}\left(E\bar{R}^\top, -1\right).\end{aligned} \tag{20}$$

Here, the strictly lower triangular masks restrict interactions to earlier positions. Thus, $L\Delta$ and JZ capture the effects of past delta corrections and residual writes, respectively, giving

$$(I_C + L)\Delta = (I_C - J)Z - E\mathbf{S}_{\text{in}}. \tag{21}$$

The DPLR transition determines the coefficient matrix $I_C + L$, while the residual path adds $-JZ$ to the right-hand side. Since JZ depends only on inputs within the chunk, it can be precomputed independently of the incoming state and incorporated into the same triangular solve.

To separate the incoming-state dependence, we define the erase-side and write-side WY auxiliaries $\mathcal{W}$ and $\mathcal{U}$ through

$$(I_C + L)\begin{bmatrix}\mathcal{W} & \mathcal{U}\end{bmatrix} = \begin{bmatrix}E & (I_C - J)Z\end{bmatrix}. \tag{22}$$

Both auxiliaries depend only on chunk inputs and can therefore be computed independently across chunks. Once the incoming state is available, the corrections are recovered as $\Delta = \mathcal{U} - \mathcal{W}\mathbf{S}_{\text{in}}$.

The final state and all token outputs within the chunk are computed by aggregating contributions from the incoming state and both update paths. Define the factors propagated to the chunk boundary as $U_{\text{tail}} = \bar{U}\operatorname{Diag}(\boldsymbol{\gamma}_C)$ and $R_{\text{tail}} = \bar{R}\operatorname{Diag}(\boldsymbol{\gamma}_C)$. For output computation, let $(Q_\gamma)_{i,:} = (\boldsymbol{\gamma}_i \odot \mathbf{q}_i)^\top$ and define

$$\begin{aligned}P_u &= \operatorname{tril}\left(Q_\gamma(\bar{U})^\top\right), \\ P_r &= \operatorname{tril}\left(Q_\gamma\bar{R}^\top\right).\end{aligned} \tag{23}$$

These masks include the diagonal because each output reads the state after the current update. The outgoing state and output block are then

$$\begin{aligned}\mathbf{S}_{\text{out}} &= \operatorname{Diag}(\boldsymbol{\gamma}_C)\mathbf{S}_{\text{in}} + (U_{\text{tail}})^\top\Delta + R_{\text{tail}}^\top Z, \\ O &= Q_\gamma\mathbf{S}_{\text{in}} + P_u\Delta + P_r Z,\end{aligned} \tag{24}$$

where row i of O is $\mathbf{o}_i^\top$. Together, these expressions complete the forward computation within each chunk, yielding all token outputs and the state passed to the next chunk.

Together, Eqs. (22) and (24) organize the complete **Stable GDN-2** update into the standard chunkwise parallel framework. The auxiliary matrices can be prepared independently for each chunk, followed by recurrent state propagation across chunks and parallel computation of the output blocks. Within this schedule, residual writing contributes additional matrix products to the correction, state, and output computations while retaining the same $C \times C$ triangular system size. **Stable GDN-2** can therefore reuse the chunkwise training structure of delta-rule models, with linear complexity in sequence length for fixed chunk size and head dimensions.

5. Experiments

5.1. Experimental Setup

5.1.1. Datasets

In this work, we train models through 64K-context pretraining on Prolong dataset (Gao et al., 2025). We merge the 64K and 512K ProLong data, retokenize the combined corpus with the TinyLlama (Zhang et al., 2024) tokenizer, and repack it into uniform sequences of 64K tokens. The resulting training stream contains approximately 82.78B tokens and uses next-token prediction over all non-padding tokens. Recurrent states are reset at each training-sample boundary and propagate continuously across documents within each packed sequence.

5.1.2. Baselines

We compare **Stable GDN-2** with four baselines under the same training protocol. GDN-2 (Hatamizadeh et al., 2026) provides the direct reference for assessing the effect of stabilization, while KDA (Team et al., 2025) represents a related delta-rule model with channel-wise forgetting. D-Clip (Su, 2026) caps each decay factor at $1/\|\mathbf{A}_t\|_2$, making the full transition non-expansive through stronger decay while preserving the original erase gate. β-Guard (Hatamizadeh, 2026) instead shrinks channel-wise deviations of the erase gate from its key-weighted mean when needed to satisfy a sufficient condition for non-expansive transitions, preserving both this mean and the original decay. Both methods retain the original GDN-2 write term. In comparison, **Stable GDN-2** combines symmetric erasure with decay-regulated residual writing, as described in Section 4.1.

5.1.3. Evaluation and Implementation

Evaluation. Following (Hatamizadeh et al., 2026), we assess zero-shot common-sense reasoning on nine tasks: PIQA (Bisk et al., 2020), HellaSwag (Zellers et al., 2019), WinoGrande (Sakaguchi et al., 2020), ARC-C and ARC-E (Clark et al., 2018), OpenBookQA (Mihaylov et al., 2018), SocialIQA (Sap et al., 2019), BoolQ (Clark et al., 2019), and LAMBADA (Paperno et al., 2016). We use normalized accuracy for PIQA, HellaSwag, ARC-C, ARC-E, and OpenBookQA, and accuracy for the remaining tasks. For language modeling, we report WikiText (Merity et al., 2016) perplexity at 8K, 16K, 32K, and 64K windows, together with LAMBADA perplexity. We evaluate long-context retrieval capabilities on 11 task types from RULER (Hsieh et al., 2024). Specifically, S1, S2, S3, and MK1 correspond to `niah_single_1`, `niah_single_2`, `niah_single_3`, and `niah_multikey_1`, respectively. The other tasks include Multi-Query (MQ), Multi-Value (MV), Common-Word Extraction (CWE), Frequent-Word Extraction (FWE), Question Answering on HotpotQA (QA-H) and SQuAD (QA-S), and Variable Tracking (VT). Evaluation uses context lengths of 32K, 64K, and 128K.

Implementation. We consider two distinct 20-layer architectures with different attention configurations. The first is a recurrent model consisting exclusively of linear-attention layers, while the second adopts a hybrid design that repeats three linear-attention layers followed by one full-attention layer five times. Full attention covers the entire causal context. Both architectures use a hidden size of 2,304 and an FFN intermediate size of 6,208. Each linear-attention layer has 16 heads with 128-dimensional keys and values. GDN-2 and **Stable GDN-2** have equal parameter counts within each architecture, approximately 1.60B for the recurrent model and 1.55B for the hybrid model.

All main-comparison models are randomly initialized and trained for one data epoch, comprising 39,473 optimizer steps with a global batch of 32 sequences, or approximately 2.10M token positions per step. We use AdamW with $\beta_1 = 0.9$, $\beta_2 = 0.95$, weight decay 0.1, and gradient clipping at 1.0. The learning rate increases linearly over 500 warmup steps and then remains constant at 3×10^{-4}. Training uses BF16 mixed precision and full activation recomputation on 16 NVIDIA H800 GPUs per run. Main performance comparisons use initialization seed 42 and data-order seed 2026, and evaluate the final checkpoint at step 39,473.

5.2. Main Results

To demonstrate the effectiveness of **Stable GDN-2**, we compare it against four representative baselines, including GDN-2, KDA, D-Clip, and β-Guard, across both Recurrent and Hybrid architectures. We evaluate their downstream task performance, language modeling quality, and long-context retrieval capabilities. The results are reported in Table 2 and 3.

Downstream Performance and Language Modeling. Table 2 reports the downstream task accuracy and language

Table 2. Performance comparison on zero-shot common-sense reasoning and language modeling. Bold and underlined values indicate the best and second-best results within each setting, respectively.

Method	Accuracy (%) ↑										WikiText PPL ↓				LAMBADA PPL ↓
	Average	PIQA	Hella Swag	Wino Grande	ARC-C	ARC-E	Open BookQA	Social IQA	BoolQ	LAMBADA	8K	16K	32K	64K	
Recurrent Setting															
GDN-2	**50.39**	71.27	52.21	56.35	30.29	54.42	35.20	40.12	58.04	**55.56**	7.849	7.838	7.825	7.804	**7.856**
D-Clip	49.10	70.67	51.65	55.96	28.41	53.11	34.60	**41.76**	51.99	53.72	7.889	7.877	7.873	7.871	8.691
β-Guard	48.89	70.29	51.46	**57.46**	29.69	55.05	32.80	41.15	50.89	51.25	7.858	7.835	7.820	7.806	9.543
KDA	49.02	**71.60**	**52.95**	56.99	**31.31**	53.70	**37.00**	40.53	44.13	53.00	**7.697**	**7.695**	**7.688**	**7.677**	9.250
Ours	**50.39**	71.44	51.57	54.93	30.20	**58.71**	34.80	39.87	**59.69**	52.28	7.801	7.788	7.776	7.759	9.333
Hybrid Setting															
GDN-2	49.91	70.95	52.05	**56.91**	28.67	53.75	34.80	41.56	56.91	53.64	7.546	**7.507**	**7.495**	**7.493**	9.001
D-Clip	50.29	**72.09**	52.50	55.96	**31.31**	**56.27**	**35.00**	41.10	54.07	54.30	**7.542**	7.515	7.518	7.519	8.277
β-Guard	49.68	71.00	51.76	54.46	30.29	52.31	32.60	41.45	57.83	55.42	7.601	7.572	7.562	7.561	**8.153**
KDA	50.36	71.65	**52.59**	56.43	30.63	55.64	34.00	41.20	55.38	**55.75**	7.544	7.510	7.504	7.504	8.399
Ours	**50.53**	70.73	52.29	54.06	29.78	56.23	**35.00**	**41.66**	**59.42**	55.64	7.594	7.564	7.557	7.558	8.586

modeling perplexity of different methods under the Recurrent and Hybrid settings. Under the Recurrent setting, our method achieves an average downstream accuracy of 50.39%, matching the original GDN-2 while outperforming D-Clip, β-Guard, and KDA. On WikiText, our method consistently achieves lower perplexity than GDN-2, D-Clip, and β-Guard across all evaluated context windows from 8K to 64K, with PPLs of 7.801, 7.788, 7.776, and 7.759, respectively. KDA, however, achieves lower perplexity, indicating that different stabilization designs exhibit distinct trade-offs in language modeling performance.

Under the Hybrid setting, our method achieves the highest average downstream accuracy of 50.53%, improving over GDN-2 by 0.62 percentage points and slightly outperforming KDA (50.36%) and D-Clip (50.29%). On WikiText, our method exhibits slightly higher perplexity than GDN-2, D-Clip, and KDA, while remaining comparable to β-Guard, with relatively small differences across context windows.

Long-Context Retrieval. Table 3 evaluates the ability of different methods to retrieve and utilize historical information at context lengths of 32K, 64K, and 128K. In the Recurrent setting, our method achieves the highest average RULER scores across all three context lengths, reaching 17.35, 16.34, and 11.62, respectively. Notably, at 128K, our method achieves 11.62, outperforming GDN-2 (4.82), D-Clip (7.02), β-Guard (10.82), and KDA (9.96). A similar trend is observed on NIAH, where our method achieves the best results at 32K and 64K, with scores of 29.90 and 29.10, respectively. At 128K, however, β-Guard and KDA slightly outperform our method.

For the Hybrid setting, our method achieves the highest average RULER score at 32K, reaching 31.72 and outperforming D-Clip (30.37) and the other baselines. At 64K, our method achieves 18.69, slightly below β-Guard (19.13), while outperforming D-Clip, GDN-2, and KDA. At 128K, β-Guard and D-Clip achieve slightly higher scores of 10.52 and 8.85, respectively, compared with 8.78 for our method. Furthermore, our method exhibits a particularly pronounced advantage on Variable Tracking, achieving scores of 31.88 and 17.80 at 32K and 64K, respectively, substantially outperforming all other baselines. These results demonstrate that **Stable GDN-2** maintains strong historical information retrieval capabilities across different context lengths and architectural settings, with particularly competitive long-context retrieval performance under the Recurrent setting.

Overall Analysis. Taken together, **Stable GDN-2** achieves a favorable balance between recurrent state stability and modeling capability. Compared with stabilization methods such as D-Clip and β-Guard, which constrain state propagation or erasure strength, our method retains channel-wise gating and flexible information writing while achieving competitive performance across multiple language modeling and long-context retrieval benchmarks. In particular, under the Recurrent setting, our method consistently achieves the highest average RULER scores across all three evaluated context lengths. These results suggest that effective control of recurrent state amplification can be achieved without sacrificing the model's ability to utilize long-term historical information through a carefully designed combination of state erasure and information writing mechanisms.

Ablation Studies. To investigate the effectiveness of different components of **Stable GDN-2**, we conduct ablation studies on different components. Figure 4 presents ablation results to examine the contributions of residual writing

Table 3. RULER benchmark results (%) at 32K, 64K, and 128K context lengths. NIAH denotes the average score across four needle-in-a-haystack tasks (S1, S2, S3, and MK1). The remaining metrics cover multi-query (MQ), multi-value (MV), common-word extraction (CWE), frequent-word extraction (FWE), question answering on HotpotQA (QA-H) and SQuAD (QA-S), and variable tracking (VT).

Recurrent Setting

Model	Seq Len	Avg.	NIAH	MQ	MV	CWE	FWE	QA-H	QA-S	VT
GDN-2	32K	15.85	28.65	4.40	10.00	0.20	19.73	**12.60**	7.95	**4.84**
	64K	12.95	24.45	2.25	4.30	0.20	13.33	**11.60**	**6.95**	6.04
	128K	4.82	2.35	**2.40**	2.55	0.32	12.87	11.80	6.15	7.48
D-Clip	32K	17.34	27.80	5.15	11.40	5.40	**35.40**	**12.60**	6.93	2.68
	64K	11.94	15.75	3.35	**5.45**	1.60	**34.13**	10.60	6.15	**7.08**
	128K	7.02	6.45	2.25	**3.00**	0.50	**29.53**	10.60	5.42	0.12
β-Guard	32K	16.29	29.60	4.05	**11.90**	6.20	16.33	11.60	7.40	3.28
	64K	15.27	28.25	2.15	3.05	2.04	25.80	11.40	6.55	4.00
	128K	10.82	**17.95**	1.70	2.40	0.32	18.20	**12.00**	5.93	6.68
KDA	32K	16.51	27.75	2.55	5.70	**9.62**	31.33	11.40	7.60	2.40
	64K	14.55	27.05	2.00	2.85	**3.14**	24.47	8.60	6.87	3.88
	128K	9.96	17.85	2.30	2.35	**1.24**	11.53	8.80	6.07	5.84
Ours	32K	**17.35**	**29.90**	**7.15**	8.80	2.70	30.13	10.40	**8.80**	3.28
	64K	**16.34**	**29.10**	**3.50**	5.10	0.48	32.33	9.60	6.13	6.24
	128K	**11.62**	17.30	**2.40**	2.55	0.40	27.20	10.80	**6.42**	**8.88**

Hybrid Setting

Model	Seq Len	Avg.	NIAH	MQ	MV	CWE	FWE	QA-H	QA-S	VT
GDN-2	32K	25.39	50.80	18.70	23.20	0.36	4.53	19.00	**10.27**	0.00
	64K	15.50	27.20	10.15	17.70	0.08	14.33	12.40	7.07	0.00
	128K	8.57	11.85	2.00	6.50	0.02	20.27	11.80	6.33	0.00
D-Clip	32K	30.37	**62.00**	**28.10**	**25.50**	2.94	4.60	17.20	6.73	1.00
	64K	18.65	**33.30**	**18.25**	**20.20**	0.56	11.93	12.80	5.37	2.84
	128K	8.85	**13.70**	**10.40**	**12.10**	0.16	4.07	10.80	4.93	0.08
β-Guard	32K	27.40	51.45	16.70	19.85	0.06	**25.87**	21.00	8.87	3.24
	64K	**19.13**	30.90	12.95	18.60	0.04	**32.33**	14.00	7.00	1.92
	128K	**10.52**	13.30	6.50	10.50	0.04	25.27	13.00	**7.22**	0.00
KDA	32K	24.31	42.55	16.85	22.35	1.92	22.47	**21.40**	6.93	5.28
	64K	14.49	21.95	9.80	13.45	0.30	24.47	**14.80**	**7.27**	1.48
	128K	8.20	8.00	2.65	4.20	0.08	**31.73**	**14.20**	5.22	0.08
Ours	32K	**31.72**	59.05	22.60	22.55	**4.34**	12.67	12.20	6.47	**31.88**
	64K	18.69	29.90	17.10	18.90	**1.30**	14.53	10.80	5.57	**17.80**
	128K	8.78	11.75	8.10	10.80	**0.72**	13.53	10.00	3.67	**2.80**

(RW) and decay regulation (DR). *Ours w/o RW* removes residual writing by setting $\mathbf{r}_t = \mathbf{0}$ in Eq. (14), yielding $\mathbf{S}_t = \left[\mathbf{I} - \mathbf{u}_t(\mathbf{u}_t)^\top\right]\mathbf{D}_t\mathbf{S}_{t-1} + \mathbf{u}_t\mathbf{z}_t^\top$. *Ours w/o DR* removes the factor $\sqrt{1-\mathbf{d}_t^2}$ from the residual write, giving $\mathbf{r}_t = \mathbf{k}_t - \mathbf{u}_t$. The total write key then becomes $\mathbf{k}_t$, resulting in $\mathbf{S}_t = \left[\mathbf{I} - \mathbf{u}_t(\mathbf{u}_t)^\top\right]\mathbf{D}_t\mathbf{S}_{t-1} + \mathbf{k}_t\mathbf{z}_t^\top$. Figure 4(a) reports the average zero-shot accuracy across the nine benchmarks in Table 2, while Figure 4(b) reports WikiText-2 token perplexity, both with an 8K context. Figure 4(c) reports the mean score over four RULER retrieval tasks (`NIAH single-1/2/3` and `multikey-1`) at 32K, 64K, and 128K context lengths, averaging all 12 task–length combinations equally.

Removing either component consistently degrades performance across the evaluated metrics. Specifically, removing RW reduces the short-task average accuracy from 50.39% to 49.43% and the RULER score from 25.43 to 22.22, accompanied by an increase in WikiText perplexity from 7.8004 to 7.8104. This demonstrates the importance of residual writing in recovering writing capacity suppressed by symmetric erasure. Removing DR leads to a more pronounced degradation in long-context retrieval, with the RULER score dropping to 13.53, while also reducing short-task accuracy to 49.18%. These results highlight the importance of regulating residual writes to control their cumulative effects and preserve long-term memory capabilities. Overall, the two components play complementary roles, with residual writing restoring flexible memory updates and decay regulation controlling the additional energy introduced by this pathway.

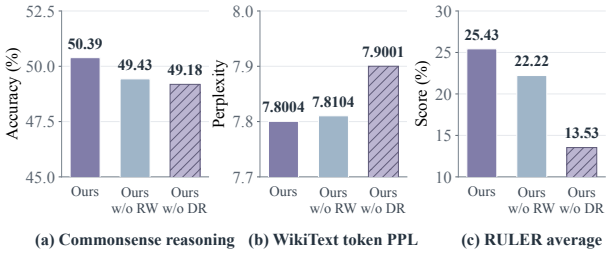

Figure 4. Results for ablation studies.

6. Related Work

6.1. Linear Attention and Delta-rule Models

Recent advances in linear attention have improved sequence modeling through more expressive gating mechanisms and recurrent state updates (Yang et al., 2024a; 2025; Team et al., 2025; Sun et al., 2026; Hatamizadeh et al., 2026). The basic formulation uses kernel feature maps to summarize the prefix in a recurrent state of key–value outer products, allowing queries to retrieve information at a cost independent of context length (Katharopoulos et al., 2020). Interpreting this recurrence as a fast-weight programmer (Schlag et al., 2021) motivates replacing additive writes with delta-rule updates that selectively revise existing associations (Irie et al., 2021), while chunkwise parallel algorithms make these updates practical for training at scale (Yang et al., 2024b). A parallel thread governs how long the state retains its contents, through fixed exponential decay in RetNet (Sun et al., 2023), input-dependent per-channel decay in gated linear attention (Yang et al., 2024a), depth-wise gate floors and state expansion in HGRN and HGRN2 (Qin et al., 2023; 2024), and selective discretizations in the Mamba line (Gu & Dao, 2024; Dao & Gu, 2024; Lahoti et al., 2026). Gated DeltaNet applied the delta correction to a decayed state (Yang et al., 2025), and Kimi Delta Attention refined the decay to a diagonal (Team et al., 2025). GDN-2 (Hatamizadeh et al., 2026) then split the shared scalar gate into independent channel-wise erase and write gates, an asymmetry EDA approaches from the write-address side (Li et al., 2026), while DeltaProduct (Siems et al., 2025) and RWKV-7 (Peng et al., 2025) pursue expressive transitions by other means. In our work, we identify recurrent state amplification as a fundamental source of instability in GDN-2 and propose **Stable GDN-2** to ensure non-expansive memory propagation through symmetric erasure.

6.2. Memory Capacity and Stabilization of Recurrent Models

A fixed-size state bounds how much history a recurrent model can carry. Attention reads as associative memory (Bietti et al., 2023); capacity rather than parameter count separates attention from gated convolution, and the resulting recall–throughput trade-off can be quantified (Arora et al., 2024a;b); no fixed latent state reproduces arbitrary inputs verbatim (Jelassi et al., 2024). Designs respond by widening the state (Qin et al., 2024), letting it grow logarithmically (Guo et al., 2026), or externalizing it (Behrouz et al., 2025). Whether a state of a given size propagates without numerical degradation is a separate question with older roots: long-range learning is tied to products of per-step Jacobians (Bengio et al., 1994), gradients are rescaled by their norm (Pascanu et al., 2013), and transitions may be confined to unitary matrices (Arjovsky et al., 2016). In transformers the analogue targets logits rather than states, via query–key normalization (Wortsman et al., 2024) or direct bounds on how far logits may move in one step (Anson & Aitchison, 2025). Each of these constrains a single transition; accumulation is left implicit, which suffices for normal transitions but not for the non-symmetric erase matrix of GDN-2. In our work, we show that non-expansive state propagation of GDN-2 can be guaranteed through symmetric erasure, preventing cumulative amplification while preserving flexible memory updates.

7. Conclusion

In this work, we identify asymmetric erasure as a fundamental source of recurrent state amplification in GDN-2, showing theoretically how non-normal state transitions can accumulate across tokens and lead to numerical instability. To address this issue, we propose **Stable GDN-2**, which combines symmetric erasure with decay-regulated residual writing to ensure non-expansive state propagation while preserving flexible memory updates and chunkwise parallel computation. Extensive experiments demonstrate that our method stabilizes long-context training and delivers strong long-context retrieval and downstream task performance, while maintaining competitive language modeling quality. These findings highlight the importance of controlling recurrent state dynamics and demonstrate that numerical stability can be substantially improved without sacrificing the expressive capacity of recurrent memory.

References

Anson, B. and Aitchison, L. Controlling changes to attention logits, 2025.

Arjovsky, M., Shah, A., and Bengio, Y. Unitary evolution recurrent neural networks. In Balcan, M. F. and Weinberger, K. Q. (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 1120–1128, New York, New York, USA, 20–22 Jun 2016. PMLR.

Arora, S., Eyuboglu, S., Timalsina, A., Johnson, I., Poli, M., Zou, J. Y., Rudra, A., and Re, C. Zoology: Measuring and improving recall in efficient language models. In Kim,

B., Yue, Y., Chaudhuri, S., Fragkiadaki, K., Khan, M., and Sun, Y. (eds.), *International Conference on Learning Representations*, volume 2024, pp. 15664–15730, 2024a.

Arora, S., Eyuboglu, S., Zhang, M., Timalsina, A., Alberti, S., Zou, J., Rudra, A., and Re, C. Simple linear attention language models balance the recall-throughput tradeoff. In Salakhutdinov, R., Kolter, Z., Heller, K., Weller, A., Oliver, N., Scarlett, J., and Berkenkamp, F. (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 1763–1840. PMLR, 21–27 Jul 2024b.

Behrouz, A., Zhong, P., and Mirrokni, V. Titans: Learning to memorize at test time. In Belgrave, D., Zhang, C., Lin, H., Pascanu, R., Koniusz, P., Ghassemi, M., and Chen, N. (eds.), *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pp. 113506–113543. Curran Associates, Inc., 2025. doi: 10.52202/085713-3786.

Bengio, Y., Simard, P., and Frasconi, P. Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks*, 5(2):157–166, March 1994. ISSN 1941-0093. doi: 10.1109/72.279181.

Bietti, A., Cabannes, V., Bouchacourt, D., Jegou, H., and Bottou, L. Birth of a transformer: A memory viewpoint. In Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., and Levine, S. (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 1560–1588. Curran Associates, Inc., 2023. doi: 10.52202/075280-0077.

Bisk, Y., Zellers, R., Le bras, R., Gao, J., and Choi, Y. Piqa: Reasoning about physical commonsense in natural language. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):7432–7439, Apr. 2020. doi: 10.1609/aaai.v34i05.6239.

Clark, C., Lee, K., Chang, M.-W., Kwiatkowski, T., Collins, M., and Toutanova, K. BoolQ: Exploring the surprising difficulty of natural yes/no questions. In Burstein, J., Doran, C., and Solorio, T. (eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 2924–2936, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1300.

Clark, P., Cowhey, I., Etzioni, O., Khot, T., Sabharwal, A., Schoenick, C., and Tafjord, O. Think you have solved question answering? try arc, the ai2 reasoning challenge, 2018.

Dao, T. and Gu, A. Transformers are SSMs: Generalized models and efficient algorithms through structured state space duality. In Salakhutdinov, R., Kolter, Z., Heller, K., Weller, A., Oliver, N., Scarlett, J., and Berkenkamp, F. (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 10041–10071. PMLR, 21–27 Jul 2024.

Gao, T., Wettig, A., Yen, H., and Chen, D. How to train long-context language models (effectively). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 7376–7399, 2025.

Gu, A. and Dao, T. Mamba: Linear-time sequence modeling with selective state spaces, 2024.

Guo, H., Yang, S., Goel, T., Xing, E. P., Dao, T., and Kim, Y. Log-linear attention, 2026.

Hatamizadeh, A. Update chunk_gdn2.py. NVlabs/GatedDeltaNet-2, GitHub. https://github.com/NVlabs/GatedDeltaNet-2/commit/de8476de0fecae08fc16f0bf3154900244c61c73, October 2026. Commit de8476d, October 11, 2026.

Hatamizadeh, A., Choi, Y., and Kautz, J. Gated deltanet-2: Decoupling erase and write in linear attention, 2026.

Hsieh, C.-P., Sun, S., Kriman, S., Acharya, S., Rekesh, D., Jia, F., Zhang, Y., and Ginsburg, B. Ruler: What's the real context size of your long-context language models?, 2024.

Irie, K., Schlag, I., Csordás, R., and Schmidhuber, J. Going beyond linear transformers with recurrent fast weight programmers. *Advances in neural information processing systems*, 34:7703–7717, 2021.

Jelassi, S., Brandfonbrener, D., Kakade, S. M., and Malach, E. Repeat after me: Transformers are better than state space models at copying. In Salakhutdinov, R., Kolter, Z., Heller, K., Weller, A., Oliver, N., Scarlett, J., and Berkenkamp, F. (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 21502–21521. PMLR, 21–27 Jul 2024.

Katharopoulos, A., Vyas, A., Pappas, N., and Fleuret, F. Transformers are rnns: Fast autoregressive transformers with linear attention. In *International conference on machine learning*, pp. 5156–5165. PMLR, 2020.

Lahoti, A. S., Li, K., Chen, B., Wang, C., Bick, A., Kolter, Z., Dao, T., and Gu, A. Mamba-3: Improved sequence modeling using state space principles. In Vondrick, C., Hariharan, B., Raffel, C., Pinto, L., Yang, D., and Faust,

A. (eds.), *International Conference on Learning Representations*, volume 2026, pp. 86173–86202, 2026.

Li, X., Zhang, C., Luo, H., Lin, X., Wang, Z., Qiu, Z., Mao, Y., Chen, L., Yuan, M., Sun, M., Jiang, H., Zhang, S., Men, R., Hu, W., Cheng, G., Zheng, B., Liu, D., and Zhou, J. Erase-then-delta attention: Decoupling erase and write addresses in delta-rule linear attention, 2026.

Llama, H. T. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2407.21783*, 2023.

Lu, E., Jiang, Z., Liu, J., Du, Y., Jiang, T., Hong, C., Liu, S., He, W., Yuan, E., Wang, Y., et al. Moba: Mixture of block attention for long-context llms. *Advances in Neural Information Processing Systems*, 38:17790–17815, 2026.

Merity, S., Xiong, C., Bradbury, J., and Socher, R. Pointer sentinel mixture models, 2016.

Mihaylov, T., Clark, P., Khot, T., and Sabharwal, A. Can a suit of armor conduct electricity? a new dataset for open book question answering. In Riloff, E., Chiang, D., Hockenmaier, J., and Tsujii, J. (eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 2381–2391, Brussels, Belgium, October-November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1260.

Nazari, P. and Rusch, T. K. On state reduction in linear attention. *arXiv preprint arXiv:2602.04852*, 2026. URL https://arxiv.org/abs/2602.04852.

Paperno, D., Kruszewski, G., Lazaridou, A., Pham, N. Q., Bernardi, R., Pezzelle, S., Baroni, M., Boleda, G., and Fernández, R. The LAMBADA dataset: Word prediction requiring a broad discourse context. In Erk, K. and Smith, N. A. (eds.), *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1525–1534, Berlin, Germany, August 2016. Association for Computational Linguistics. doi: 10.18653/v1/P16-1144.

Pascanu, R., Mikolov, T., and Bengio, Y. On the difficulty of training recurrent neural networks. In Dasgupta, S. and McAllester, D. (eds.), *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pp. 1310–1318, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR.

Peng, B., Zhang, R., Goldstein, D., Alcaide, E., Du, X., Hou, H., Lin, J., Liu, J., Lu, J., Merrill, W., Song, G., Tan, K., Utpala, S., Wilce, N., Wind, J. S., Wu, T., Wuttke, D., and Zhou-Zheng, C. Rwkv-7 "goose" with expressive dynamic state evolution, 2025.

Qin, Z., Yang, S., and Zhong, Y. Hierarchically gated recurrent neural network for sequence modeling. In Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., and Levine, S. (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 33202–33221. Curran Associates, Inc., 2023. doi: 10.52202/075280-1442.

Qin, Z., Yang, S., Sun, W., Shen, X., Li, D., Sun, W., and Zhong, Y. HGRN2: Gated linear RNNs with state expansion. In *First Conference on Language Modeling*, 2024.

Sakaguchi, K., Le Bras, R., Bhagavatula, C., and Choi, Y. Winogrande: An adversarial winograd schema challenge at scale. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8732–8740, Apr. 2020. doi: 10.1609/aaai.v34i05.6399.

Sap, M., Rashkin, H., Chen, D., Le Bras, R., and Choi, Y. Social IQa: Commonsense reasoning about social interactions. In Inui, K., Jiang, J., Ng, V., and Wan, X. (eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 4463–4473, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1454.

Schlag, I., Irie, K., and Schmidhuber, J. Linear transformers are secretly fast weight programmers. In *International conference on machine learning*, pp. 9355–9366. PMLR, 2021.

Siems, J., Carstensen, T., Zela, A., Hutter, F., Pontil, M., and Grazzi, R. Deltaproduct: Improving state-tracking in linear rnns via householder products. In Belgrave, D., Zhang, C., Lin, H., Pascanu, R., Koniusz, P., Ghassemi, M., and Chen, N. (eds.), *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pp. 153738–153782. Curran Associates, Inc., 2025. doi: 10.52202/085713-5141.

Su, J. How to Make GDN2 More Stable? Scientific Spaces. https://spaces.ac.cn/archives/11945, October 2026. Blog post.

Sun, P., Hu, Y., Tan, J., Wang, X., Zhang, J., Lu, Y., Sun, Y., Xie, Y., and Cai, X. Fg^2-gdn: Enhancing long-context gated delta networks with doubly fine-grained control, 2026.

Sun, Y., Dong, L., Huang, S., Ma, S., Xia, Y., Xue, J., Wang, J., and Wei, F. Retentive network: A successor to transformer for large language models. *arXiv preprint arXiv:2307.08621*, 2023.

Sun, Y., Li, Z., Zhang, Y., Pan, T., Dong, B., Guo, Y., and Wang, J. Efficient attention mechanisms for large language models: A survey. *arXiv preprint arXiv:2507.19595*, 2025.

Team, K., Zhang, Y., Lin, Z., Yao, X., Hu, J., Meng, F., Liu, C., Men, X., Yang, S., Li, Z., Li, W., Lu, E., Liu, W., Chen, Y., Xu, W., Yu, L., Wang, Y., Fan, Y., Zhong, L., Yuan, E., Zhang, D., Zhang, Y., Liu, T. Y., Wang, H., Fang, S., He, W., Liu, S., Li, Y., Su, J., Qiu, J., Pang, B., Yan, J., Jiang, Z., Huang, W., Yin, B., You, J., Wei, C., Wang, Z., Hong, C., Chen, Y., Chen, G., Wang, Y., Zheng, H., Wang, F., Liu, Y., Dong, M., Zhang, Z., Pan, S., Wu, W., Wu, Y., Guan, L., Tao, J., Fu, G., Xu, X., Wang, Y., Lai, G., Wu, Y., Zhou, X., Yang, Z., and Du, Y. Kimi linear: An expressive, efficient attention architecture, 2025.

Wortsman, M., Liu, P., Xiao, L., Everett, K., Alemi, A., Adlam, B., Co-Reyes, J., Gur, I., Kumar, A., Novak, R., Pennington, J., Sohl-Dickstein, J., Xu, K., Lee, J., Gilmer, J., and Kornblith, S. Small-scale proxies for large-scale transformer training instabilities. In Kim, B., Yue, Y., Chaudhuri, S., Fragkiadaki, K., Khan, M., and Sun, Y. (eds.), *International Conference on Learning Representations*, volume 2024, pp. 49844–49869, 2024.

Xu, P., Ping, W., Wu, X., Xu, C., Liu, Z., Shoeybi, M., and Catanzaro, B. Chatqa 2: Bridging the gap to proprietary llms in long context and rag capabilities. In *International Conference on Learning Representations*, volume 2025, pp. 67244–67258, 2025.

Yang, S., Wang, B., Shen, Y., Panda, R., and Kim, Y. Gated linear attention transformers with hardware-efficient training. In *Forty-first International Conference on Machine Learning*, 2024a.

Yang, S., Wang, B., Zhang, Y., Shen, Y., and Kim, Y. Parallelizing linear transformers with the delta rule over sequence length. *Advances in neural information processing systems*, 37:115491–115522, 2024b.

Yang, S., Kautz, J., and Hatamizadeh, A. Gated delta networks: Improving mamba2 with delta rule. In Yue, Y., Garg, A., Peng, N., Sha, F., and Yu, R. (eds.), *International Conference on Learning Representations*, volume 2025, pp. 29687–29707, 2025.

Zellers, R., Holtzman, A., Bisk, Y., Farhadi, A., and Choi, Y. HellaSwag: Can a machine really finish your sentence? In Korhonen, A., Traum, D., and Màrquez, L. (eds.), *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 4791–4800, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1472.

Zhang, P., Zeng, G., Wang, T., and Lu, W. Tinyllama: An open-source small language model. *arXiv preprint arXiv:2401.02385*, 2024.

Zhang, Y., Sun, W., Ding, W., Sun, S., Xie, R., Wang, H., Chen, J., Liu, J., Wang, S., Zhang, Y., et al. A survey of linear attention: Algorithm, theory, application, and infrastructure. *Transactions on Machine Learning Research*, 2026.

Appendix

A. Proofs for the Theoretical Analysis

A.1. Single-Step Erasure Gain

Proof of Theorem 3.1. Under the assumptions of the theorem, Eq. (5) gives

$$\mu = \sum_{i=1}^{d_k} b_i k_i^2 \in [0, 1]. \tag{25}$$

The second column of $\mathbf{T}$ in Eq. (6) has norm $\sqrt{1+\gamma^2} \geq 1$. Together with the identity action on the orthogonal complement, this implies $\|\mathbf{A}\|_2 = \|\mathbf{T}\|_2$.

Let $\sigma_1 \geq \sigma_2 \geq 0$ be the singular values of $\mathbf{T}$. Their squared sum and product satisfy

$$\begin{aligned} \sigma_1^2 + \sigma_2^2 &= \|\mathbf{T}\|_F^2 = (1-\mu)^2 + \gamma^2 + 1, \\ \sigma_1 \sigma_2 &= |\det \mathbf{T}| = 1 - \mu. \end{aligned} \tag{26}$$

Consequently,

$$\begin{aligned} (\sigma_1 + \sigma_2)^2 &= (2-\mu)^2 + \gamma^2, \\ (\sigma_1 - \sigma_2)^2 &= \mu^2 + \gamma^2. \end{aligned} \tag{27}$$

Taking the nonnegative square roots yields

$$\|\mathbf{A}\|_2 = \sigma_1 = \frac{\sqrt{\mu^2+\gamma^2} + \sqrt{(2-\mu)^2+\gamma^2}}{2}, \tag{28}$$

which proves Eq. (7). When $\gamma = 0$, this expression equals 1 because $\mu \in [0, 1]$. When $\gamma > 0$, the two square roots strictly exceed μ and $2-\mu$, respectively, giving $\|\mathbf{A}\|_2 > 1$. Hence $\|\mathbf{A}\|_2 = 1$ if and only if $\gamma = 0$, which is equivalent to $\mathbf{a} = \mu\mathbf{k}$ and therefore to collinearity of $\mathbf{a}$ and $\mathbf{k}$. □

The same representation also verifies that $\mathbf{A}$ is non-normal when $\gamma > 0$, since

$$\mathbf{T}^\top \mathbf{T} - \mathbf{T}\mathbf{T}^\top = \begin{pmatrix} -\gamma^2 & \mu\gamma \\ \mu\gamma & \gamma^2 \end{pmatrix} \neq \mathbf{0} \quad \text{when } \gamma > 0. \tag{29}$$

For the recurrent state energy introduced in Section 3.1, the maximum amplification factor is

$$\sup_{\mathbf{S} \neq \mathbf{0}} \frac{E(\mathbf{A}\mathbf{S})}{E(\mathbf{S})} = \|\mathbf{A}\|_2^2. \tag{30}$$

The upper bound follows by applying the operator norm bound to each column of $\mathbf{S}$. Equality is attained when one column lies along a leading right singular direction of $\mathbf{A}$ and all other columns are zero. Thus, nonzero shear permits erasure alone to increase the energy of the recurrent state, without any new write.

A.2. Unbounded Growth Without Further Writes

Proof of Corollary 3.2. A periodic sequence of two alternating transitions realizes the alignment described in Section 3.1. We construct it on two key coordinates and one recurrent state column, then embed it in higher dimensions. For this example, choose $\mathbf{D}_t = d\mathbf{I}$ with $d = 0.99$ at every token.

Choose the following pairs of keys and erase gates,

$$\begin{aligned} \mathbf{k}_a &= \frac{1}{\sqrt{5}} \begin{pmatrix} 2 \\ 1 \end{pmatrix}, \quad \mathbf{b}_a = \frac{1}{5} \begin{pmatrix} 1 \\ 4 \end{pmatrix}, \\ \mathbf{k}_b &= \frac{1}{\sqrt{5}} \begin{pmatrix} 1 \\ 2 \end{pmatrix}, \quad \mathbf{b}_b = \frac{1}{5} \begin{pmatrix} 4 \\ 1 \end{pmatrix}. \end{aligned} \tag{31}$$

Both keys have unit norm, and all erase gate entries lie in $(0, 1)$. The corresponding erase operators are

$$\begin{aligned} \mathbf{A}_a &= \frac{1}{25} \begin{pmatrix} 21 & -8 \\ -2 & 21 \end{pmatrix}, \\ \mathbf{A}_b &= \frac{1}{25} \begin{pmatrix} 21 & -2 \\ -8 & 21 \end{pmatrix} = \mathbf{A}_a^\top. \end{aligned} \tag{32}$$

Hence the full transitions $\mathbf{M}_a = d\mathbf{A}_a$ and $\mathbf{M}_b = d\mathbf{A}_b$ satisfy $\mathbf{M}_b = \mathbf{M}_a^\top$. For $\mathbf{A}_a$, the geometric coefficients are $\mu = 8/25$ and $\gamma = 6/25$. Theorem 3.1 therefore gives

$$\lambda = d^2 \frac{19 + 6\sqrt{2}}{25} \approx 1.077533 > 1. \tag{33}$$

At the first token, choose the write

$$\mathbf{S}_1 = \frac{1}{\sqrt{100 + 10\sqrt{2}}} \begin{pmatrix} 7 \\ -(1+5\sqrt{2}) \end{pmatrix}, \qquad \|\mathbf{S}_1\|_2 = 1. \tag{34}$$

This write is realized by $\mathbf{k}_1 = \mathbf{S}_1$, $\mathbf{b}_1 = (1/2, 1/2)^\top$, $w_1 = 1/2$, and $v_1 = 2$. Indeed, the initial recurrent state is zero, so the first update reduces to $\mathbf{k}_1(w_1 v_1) = \mathbf{S}_1$. For every subsequent token $t \geq 2$, set $w_t = 1/2$ and $v_t = 0$, so that no further content is written while the write gate remains strictly between zero and one. Alternate the pairs in Eq. (31) so that

$$\mathbf{M}_{2n+2} = \mathbf{M}_a, \qquad \mathbf{M}_{2n+3} = \mathbf{M}_b, \qquad n \geq 0. \tag{35}$$

Direct multiplication gives

$$\mathbf{M}_b \mathbf{M}_a = \frac{d^2}{125} \begin{pmatrix} 89 & -42 \\ -42 & 101 \end{pmatrix}, \tag{36}$$

and the initial column in Eq. (34) satisfies $\mathbf{M}_b \mathbf{M}_a \mathbf{S}_1 = \lambda \mathbf{S}_1$. It follows by induction that

$$\begin{aligned} \mathbf{S}_{2n+1} &= \lambda^n \mathbf{S}_1, \\ \mathbf{S}_{2n+2} &= \lambda^n \mathbf{M}_a \mathbf{S}_1, \qquad n \geq 0. \end{aligned} \tag{37}$$

Moreover, the transpose pairing implies

$$\|\mathbf{M}_a\mathbf{S}_1\|_2^2 = \mathbf{S}_1^\top \mathbf{M}_b\mathbf{M}_a\mathbf{S}_1 = \lambda. \tag{38}$$

Since $\lambda = \sigma_{\max}(\mathbf{M}_a)^2$, this equality shows that $\mathbf{S}_1$ lies along a leading right singular direction of $\mathbf{M}_a$. The first transition maps it to the corresponding left singular direction, and the second maps it back. Each pair both amplifies the recurrent state and restores the direction required for the next pair. In particular, for every $t \geq 1$,

$$\|\mathbf{S}_t\|_2 = \lambda^{(t-1)/2}, \qquad E(\mathbf{S}_t) = \lambda^{t-1} \longrightarrow \infty. \tag{39}$$

This growth occurs entirely after writing ceases, despite the strict decay applied at every step.

For a recurrent state with more coordinates or columns, embed every key, including $\mathbf{k}_1$, in the same two key coordinates and set all remaining key entries to zero. Set the remaining erase gate entries and all write gate entries to $1/2$, retain $\mathbf{D}_t = d\mathbf{I}$, and place the initial value 2 in one value coordinate, with all other values zero. All subsequent value vectors are zero. The resulting recurrent state remains supported on the chosen two coordinates and one column, so

$$\|\mathbf{S}_t\|_F = \lambda^{(t-1)/2} \longrightarrow \infty. \tag{40}$$

All keys remain normalized and all erase, write, and decay gate entries lie in $(0, 1)$, so the embedded sequence satisfies the corollary. □

B. Details of Empirical Characterization

B.1. Training Protocol and Failure Accounting

The Pure Linear and Hybrid architectures, tokenizer, pretraining data, and shared optimization settings follow Section 5.1. Pure Linear uses GDN-2 in all 20 layers, while Hybrid repeats three GDN-2 layers followed by one full attention layer. Their parameter counts are approximately 1.595B and 1.554B, respectively. Pretraining starts from random initialization and uses the ProLong corpus (Gao et al., 2025) with next token supervision over all nonpadding tokens. The maximum packed sequence length is 65,536 tokens. The recurrent state propagates across documents within each pretraining pack and is reset at the start of the next pack.

Post-training uses the official training split of `nvidia/ChatQA2-Long-SFT-data` (Xu et al., 2025), with loss applied only to answer and EOS targets. The standalone post-training condition starts from random initialization, whereas post-training after pretraining starts from the final weights of the corresponding pretraining run. The latter resets the optimizer, learning rate schedule, training random number state, and step counter. We retain 143,367 complete QA examples that fit within 65,536 tokens, yielding approximately 2.170B nonpadding input tokens and 79.987M supervised targets. Examples exceeding the length limit are excluded without truncation. Each QA example is an independent segment within a pack, with positions, recurrent state, and convolution state reset at its boundary; full attention in Hybrid is also restricted to that segment. Thus, the 64K pack limit does not imply uninterrupted recurrent state propagation across separate QA examples. Both post-training conditions use one data epoch, a global batch of 16 packs, and 2,331 optimizer updates. The learning rate warms up to 3×10^{-4} over 100 updates, remains constant, and decays linearly to 3×10^{-5} over the final 233 updates. Other optimizer settings follow Section 5.1.

The four seed configurations form a 2×2 design crossing data order seeds $\{2026, 2027\}$ with initialization and training seeds $\{42, 43\}$. Seeds 1–4 correspond to the pairs $(2026, 42)$, $(2026, 43)$, $(2027, 42)$, and $(2027, 43)$, respectively. The packing seed is fixed at 3407, so changing the data seed changes pack order rather than the underlying packing. The same seed configuration identifies corresponding stages within each architecture. A failed pretraining run does not launch its dependent post-training stage, leaving three evaluated stages for Pure Linear and two for Hybrid in the final column of Table 1.

We count failures according to the loss NaN criterion in Section 3.2. Table 1 uses one designated trajectory per seed configuration and training stage. For Hybrid Seed 2, the initial standalone post-training trajectory completed all 2,331 updates, while a separately initialized repeat with the same seeds failed at update 471. The table uses the latter as its designated trajectory, giving 3/4 failures for Hybrid post-training; the earlier completion is disclosed here as an additional trajectory. Other repeats and failure replays do not contribute additional seed configurations, and resumption after an infrastructure interruption continues the same trajectory. The designated runs comprise 14 completed stages and 7 numerical failures, with 3 dependent post-training stages not launched.

B.2. Replay and Diagnostic Protocol

Figure 2 examines the Pure Linear Seed 4 pretraining failure at update 2,058. Replay starts from the last finite checkpoint, saved after update 2,057, and uses the triggering input sequence of 65,536 tokens, identified as sample 1,025,853 in the packed dataset. The recurrent state diagnostics select layer 4 and head 5, with both indices starting from zero. GDN and KDA are replayed using checkpoints at the same training step, the same input sample, and the corresponding layer and head. Each model supplies its own keys, gates, and decay factors.

The ordinary replay records recurrent state norms and the

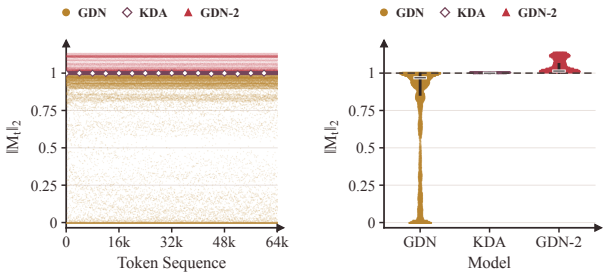

Figure 5. Spectral norms of recurrent state transitions. The full transition includes erasure and decay. Measurements use the same input sequence and selected layer and head for GDN, KDA, and GDN-2, with 65,536 transitions per model. The left panel shows norms by token position, and the right panel shows their distributions with medians and interquartile ranges. The dashed line marks $\|\mathbf{M}_t\|_2 = 1$.

occurrence of nonfinite values. State observations are sampled at intervals of 64 tokens, with the first nonfinite GDN-2 state observed at token 52,736. Nonfinite values then propagate through subsequent matrix computations to the logits and loss. The overflow marker denotes this observed position, rather than a threshold applied to the Frobenius norm. Training uses BF16 mixed precision, with recurrent state accumulation primarily in FP32 and states saved across chunks in BF16. The precision of the accumulation and saved states is therefore distinguished in the diagnosis.

For Figure 2(c), each model first processes 64 tokens with ordinary writes to obtain its initial recurrent state. From token 65 onward, the additive write term is set to zero while the keys, gates, and decay factors remain fixed to the values extracted from the corresponding ordinary replay. This isolates the propagation in Eq. (9) without changing the transition sequence. The offline diagnostic uses FP64 arithmetic with recurrent state rescaling and a separately accumulated logarithmic scale to track norms over the full sequence. The extended logarithmic curves describe this propagation diagnostic; overflow timing is taken from the native replay in Figure 2(b). Effective rank in Figure 2(d) is evaluated on finite states from ordinary replay, with scalar normalization before singular value computation. The ratio defined in the main text is invariant to this normalization and is also known as stable rank.

Figure 5 reports the spectral norms of the full transitions $\mathbf{M}_t$, including both erasure and decay. For GDN-2 in this sample and selected head, the maximum is 1.135 and 89.58% of the transitions have norm greater than one. These measurements identify expanding directions in individual transitions, while the trajectory in Figure 2(c) shows the accumulated amplification realized by their action on the recurrent state.

B.3. Recurrent State Magnitudes Across Seeds

Table 4 reports the largest recorded finite absolute recurrent state entry during each completed training trajectory. These

Table 4. Maximum recurrent state magnitudes across seeds. Entries give the largest recorded finite absolute recurrent state value during each completed training trajectory, rounded to four significant digits. ✗ marks a run terminated by loss NaN. An em dash denotes a stage that was not launched.

Model	Seed	Pretraining	Post-training	Pretraining → Post-training
Pure Linear	1	4.383×10^{5}	1.016×10^{6}	7.559×10^{11}
	2	3.157×10^{36}	✗	5.504×10^{3}
	3	1.660×10^{20}	8.048×10^{14}	1.464×10^{5}
	4	✗	1.047×10^{11}	—
Hybrid	1	9.408×10^{3}	✗	4.940×10^{2}
	2	✗	✗	—
	3	✗	3.707×10^{5}	—
	4	3.264×10^{11}	✗	1.590×10^{2}

maxima are taken over available monitoring records; the detailed failure replay in Appendix B.2 is reported separately. The seed mapping and designated trajectories follow Appendix B.1.

Extreme values occur even along training trajectories that ultimately complete. For example, Pure Linear Seed 2 completes pretraining after reaching a recorded magnitude of approximately 3.157×10^{36} during training. This value remains within the finite representable range of FP32, whose maximum is approximately 3.403×10^{38}. The absence of loss NaN therefore does not imply controlled recurrent state magnitudes throughout training.

C. State Bounds for Stable GDN-2

We use the update, normalization, and gate parameterization specified in Sections 4.1–4.2 throughout.

C.1. Write Norm and Additive Accumulation

The residual in Eq. (13) gives, for each coordinate,

$$(\mathbf{u}_t + \mathbf{r}_t)_i = \left[\sqrt{b_{t,i}} + \sqrt{1 - d_{t,i}^2}(1 - \sqrt{b_{t,i}})\right] k_{t,i}. \quad (41)$$

The coefficient in brackets lies in $[0, 1]$, because both $\sqrt{b_{t,i}}$ and $\sqrt{1 - d_{t,i}^2}$ lie in that interval. Hence $\|\mathbf{u}_t + \mathbf{r}_t\|_2 \leq \|\mathbf{k}_t\|_2 \leq 1$, and the Frobenius norm of an outer product yields

$$\begin{aligned} \|(\mathbf{u}_t + \mathbf{r}_t)\mathbf{z}_t^\top\|_F &= \|\mathbf{u}_t + \mathbf{r}_t\|_2 \|\mathbf{z}_t\|_2 \\ &\leq \|\mathbf{k}_t\|_2 \|\mathbf{z}_t\|_2 = \|\mathbf{k}_t \mathbf{z}_t^\top\|_F. \end{aligned} \quad (42)$$

Expanding Eq. (14), applying the triangle inequality, and using $\|\mathbf{\Phi}_{T:t+1}\|_2 \leq 1$ gives Eq. (16). This argument accounts for arbitrary alignment between successive writes.

C.2. Proof of the Recurrent State Norm Bound

Proof of Proposition 4.2. Fix a token and omit its index. Introduce the auxiliary quantities

$$\mathbf{f} = \mathbf{k} - \mathbf{u}, \qquad \mathbf{C} = (\mathbf{I} - \mathbf{D}^2)^{1/2}, \qquad \nu = \|\mathbf{u}\|_2^2. \tag{43}$$

Equation (13) gives $\mathbf{r} = \mathbf{C}\mathbf{f}$. Since $\mathbf{A}^{\mathrm{sym}}$ is symmetric and $\mathbf{I} - (\mathbf{A}^{\mathrm{sym}})^2 = (2-\nu)\mathbf{u}(\mathbf{u})^\top$,

$$\begin{aligned}\mathbf{I} - \mathbf{M}\mathbf{M}^\top &= \mathbf{I} - (\mathbf{A}^{\mathrm{sym}})^2 + \mathbf{A}^{\mathrm{sym}}(\mathbf{I} - \mathbf{D}^2)\mathbf{A}^{\mathrm{sym}} \\ &= (2-\nu)\mathbf{u}(\mathbf{u})^\top + \mathbf{A}^{\mathrm{sym}}\mathbf{C}^2\mathbf{A}^{\mathrm{sym}} = \mathbf{L}\mathbf{L}^\top,\end{aligned} \tag{44}$$

where the auxiliary matrix is $\mathbf{L} = [\sqrt{2-\nu}\,\mathbf{u} \quad \mathbf{A}^{\mathrm{sym}}\mathbf{C}]$. The write vector admits a factorization through the same matrix,

$$\mathbf{u} + \mathbf{r} = \mathbf{L}\boldsymbol{\xi}, \qquad \boldsymbol{\xi} = \begin{bmatrix} \dfrac{1 + (\mathbf{u})^\top \mathbf{C}\mathbf{f}}{\sqrt{2-\nu}} \\ \mathbf{f} \end{bmatrix}. \tag{45}$$

Indeed, the first block contributes $\mathbf{u}[1 + (\mathbf{u})^\top \mathbf{C}\mathbf{f}]$, while the second contributes $\mathbf{A}^{\mathrm{sym}}\mathbf{C}\mathbf{f} = \mathbf{C}\mathbf{f} - \mathbf{u}(\mathbf{u})^\top \mathbf{C}\mathbf{f}$. Their sum is $\mathbf{u} + \mathbf{r}$.

To bound $\|\boldsymbol{\xi}\|_2$, set $\eta = (\mathbf{u})^\top \mathbf{f}$ and $\tau = \|\mathbf{f}\|_2^2$. The coordinate products $u_i f_i = \sqrt{b_i}(1 - \sqrt{b_i})k_i^2$ are nonnegative. Since the diagonal entries of $\mathbf{C}$ lie in $[0, 1]$,

$$\begin{aligned} 0 \le (\mathbf{u})^\top \mathbf{C}\mathbf{f} \le \eta, \qquad \eta^2 \le \nu\tau, \\ \nu + 2\eta + \tau = \|\mathbf{k}\|_2^2 \le 1. \end{aligned} \tag{46}$$

The denominator in Eq. (45) is positive because $\nu \le 1$. Moreover,

$$\begin{aligned}&\left(\frac{3}{2} - \tau\right)(2-\nu) - (1+\eta)^2 \\ &\quad = 2(1 - \nu - 2\eta - \tau) + \frac{\nu}{2} + 2\eta \\ &\qquad + (\nu\tau - \eta^2) \ge 0.\end{aligned} \tag{47}$$

It follows that

$$\|\boldsymbol{\xi}\|_2^2 \le \tau + \frac{(1+\eta)^2}{2-\nu} \le \frac{3}{2}. \tag{48}$$

Combining Eqs. (44)–(48) gives

$$(\mathbf{u} + \mathbf{r})(\mathbf{u} + \mathbf{r})^\top = \mathbf{L}\boldsymbol{\xi}\boldsymbol{\xi}^\top\mathbf{L}^\top \preceq \frac{3}{2}(\mathbf{I} - \mathbf{M}\mathbf{M}^\top). \tag{49}$$

Thus, the block matrix $[\mathbf{M} \quad \sqrt{2/3}(\mathbf{u} + \mathbf{r})]$ has operator norm at most one. Restoring the time index, apply this contraction to the stacked recurrent state and gated value to obtain

$$\begin{aligned}\|\mathbf{S}_t\|_F^2 &= \left\| \begin{bmatrix}\mathbf{M}_t & \sqrt{2/3}(\mathbf{u}_t + \mathbf{r}_t)\end{bmatrix} \begin{bmatrix}\mathbf{S}_{t-1} \\ \sqrt{3/2}\,\mathbf{z}_t^\top\end{bmatrix} \right\|_F^2 \\ &\le \|\mathbf{S}_{t-1}\|_F^2 + \frac{3}{2}\|\mathbf{z}_t\|_2^2.\end{aligned} \tag{50}$$

Since $\mathbf{z}_t = \mathbf{w}_t \odot \mathbf{v}_t$ and the write gate has entries in $(0, 1)$, $\|\mathbf{z}_t\|_2 \le \|\mathbf{v}_t\|_2 \le V$. Summing the preceding inequality over tokens and using $\mathbf{S}_0 = \mathbf{0}$ gives

$$\begin{aligned}\|\mathbf{S}_T\|_F^2 &\le \frac{3}{2}\sum_{t=1}^{T}\|\mathbf{z}_t\|_2^2 \\ &\le \frac{3}{2}\sum_{t=1}^{T}\|\mathbf{v}_t\|_2^2 \le \frac{3}{2}TV^2.\end{aligned} \tag{51}$$

Taking square roots proves Eq. (17). □

C.3. Role of Decay Regulation

Removing $\sqrt{1 - \mathbf{d}_t^2}$ from the residual leaves the transition unchanged and makes the total write equal to $\mathbf{k}_t\mathbf{z}_t^\top$. Non-expansive transitions are therefore preserved, but the resulting recurrence need not admit an $O(\sqrt{T})$ norm bound with a constant independent of the gate values and sequence length. To see this, consider a scalar invariant direction with constant $k = 1$, $b = \varepsilon^2$, $d = 1 - \varepsilon^2$, and $z = 1$, starting from $S_0 = 0$. These choices can be embedded along one coordinate of any key dimension; for example, the gated value $z = 1$ is realized by $w = 1/2$ and $v = 2$. For $0 < \varepsilon < 1$, the unscaled residual gives

$$\begin{aligned} S_t = mS_{t-1} + 1, \qquad m = (1-\varepsilon^2)^2 < 1, \\ S_T = \sum_{j=0}^{T-1} m^j. \end{aligned} \tag{52}$$

For any $T \ge 4$, choose $\varepsilon = 1/T$. Bernoulli's inequality gives $m^{T-1} = (1 - T^{-2})^{2(T-1)} \ge 1 - 2(T-1)/T^2 \ge 1/2$, and hence $S_T \ge T/2$. Thus, the values remain bounded by $V = 2$, while $\|\mathbf{S}_T\|_F \ge T/2$. For any constant C independent of the gate values and sequence length, this family violates $\|\mathbf{S}_T\|_F \le CV\sqrt{T}$ for sufficiently large T. The growth comes from repeated writes under weak decay, rather than amplification through the transition products analyzed in Section 3.1. This separates the control of write accumulation provided by decay regulation from the non-expansive propagation provided by symmetric erasure.

C.4. A Bound under Uniform Strict Decay

If, in addition, $\|\mathbf{d}_t\|_\infty \le \rho < 1$ and $\|\mathbf{v}_t\|_2 \le V$ for all tokens, then $\|\mathbf{z}_t\|_2 \le V$. Proposition 4.1 and Eq. (42) therefore imply

$$\begin{aligned}\|\mathbf{S}_T\|_F &\le \rho^T\|\mathbf{S}_0\|_F + V\sum_{t=1}^{T}\rho^{T-t} \\ &= \rho^T\|\mathbf{S}_0\|_F + V\frac{1-\rho^T}{1-\rho} \le \|\mathbf{S}_0\|_F + \frac{V}{1-\rho}.\end{aligned} \tag{53}$$

This bound is independent of sequence length under the additional uniform decay condition. The $O(\sqrt{T})$ norm bound in Proposition 4.2 does not require a uniform gap between the decay factors and one.