arXiv:2611.06828v2 [cs.LG] 17 Nov 2026

Systolic Linear Attention: Co-Designing Tile-Grouped Gated Delta Attention for TPUs

Gagik Amirkhanyan
Google
agagik@google.com

Nuojin Cheng
Google
chengnuojin@google.com

Abstract

Gated-delta linear attention layers such as Gated DeltaNet (GDN), Kimi Delta Attention (KDA), and GDN-2 are tuned for GPUs, where small tensor-core tiles and fast vector registers make unpacked heads, per-channel decay rescaling, and per-value-head chunk solves cheap. On Google TPUs, which concentrate compute in wide 256×256 systolic matrix units backed by a slower vector unit, direct ports leave half of each matrix tile padded (for $d_k = 128$ heads) and repeat decay rescaling and chunk solves across every value head. We present **Systolic Linear Attention (SLA)**, which reshapes gated delta attention around the matrix unit by widening keys to 256 and sharing both keys and scalar write gates β_t across grouped value heads ($R\times$ fewer chunk solves), grouping decay gates into four 64-channel tiles, and fusing exact FP32 block-doubling chunk inversion and its backward reuse inside one on-chip kernel. With exact chunk inversion on TPU v7x, 2-stage `SLA-R4` (four value heads per key head) runs $2.1\times$–$2.6\times$ faster than open-source MaxText KDA and $1.4\times$–$2.4\times$ faster than open-source Tokamax KDA (and 1-kernel fused `SLA-R4` reaches $2.5\times$–$2.8\times$ and $1.5\times$–$2.4\times$ under unified single-chip protocols), while using 49%–61% less temporary high-bandwidth memory (HBM). Across a 2,048-chip TPU v7x pre-training study of sparse Mixture-of-Experts (MoE) models from 6.5B to 63.1B nominal parameters, SLA trains 11%–24% faster than MaxText KDA (4%–21% faster than Tokamax KDA), matches KDA perplexity on three paired seeds at 12.5B, and stays within 3% at other scales. Under controlled multi-seed equal-chip-hour pre-training against safe-clipped open-source Tokamax KDA on 64-chip TPU v7x slices ($n = 3$ paired seeds per arm), 1-kernel fused `SLA-R4` and fused `SLA-R4 + GDN-2` lower perplexity on all $3/3$ seeds at both 12.5B (-0.37 ± 0.11 and -0.41 ± 0.35) and 63.1B (-1.48 ± 0.21 and -1.41 ± 0.16), while $G = 4$ outperforms scalar $G = 1$ across $6/6$ paired 12.5B MoE and dense-LM seeds. Without clamping its decay gates, SLA also extrapolates from 8K training to 32K–64K context far better than unclamped KDA.

1 Introduction

Linear attention removes the quadratic compute and key-value (KV) cache of softmax attentions [17, 69] by storing the context in a fixed-size state matrix $S_t \in \mathbb{R}^{d_k \times d_v}$ (for key width d_k and value width d_v), updated at each token t with its key $k_t \in \mathbb{R}^{d_k}$ and value $v_t \in \mathbb{R}^{d_v}$ as [37]:

$$S_t = S_{t-1} + k_t v_t^\top.$$

Because this sum grows without bound, the state saturates on long sequences. Gated models such as Mamba-2 [16] and GLA [72] add a data-dependent forget (decay) gate $F_t \in (0, 1]^{d_k \times d_k}$,

$$S_t = F_t S_{t-1} + k_t v_t^\top,$$

Preprint.

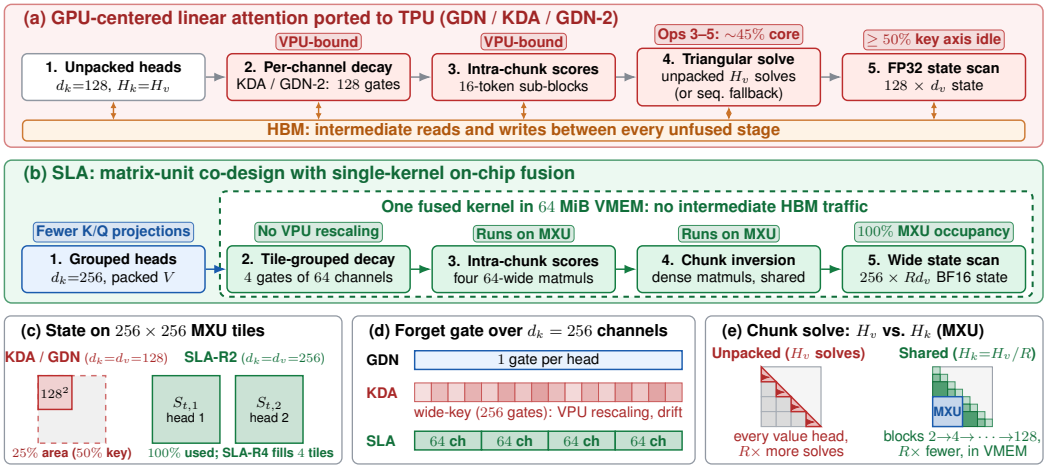

Figure 1: **Overview of SLA. (a)** Direct ports of GPU-tuned gated-delta layers (GDN [74], KDA [78], GDN-2 [29]) run as an unfused pipeline on TPUs: intermediates round-trip through HBM, per-channel decay rescaling in KDA and GDN-2 (and sequential triangular substitution in portable or FP32 fallback paths) runs on the VPU, 128-wide states leave half of each 256×256 MXU tile's key axis padded, and per-value-head write gates force H_v chunk solves per chunk. **(b)** SLA fuses stages 2–5 into one VMEM kernel with $R\times$ fewer solves. **(c)** A head with key and value widths $d_k = d_v = 128$ fills a quarter of an MXU tile's area (50% of the key axis), while SLA's wide packed state fills whole tiles. **(d)** Decay gates are grouped into four 64-channel tiles, so intra-chunk scores are four 64-wide MXU matmuls with no VPU rescaling. **(e)** Building on open-source MXU block solvers [1, 23], SLA fuses $R\times$-fewer exact FP32 block-doubling chunk solves and backward reuse inside VMEM.

but decay fades all stored keys together and hurts associative recall [5].

To overwrite individual keys selectively, DeltaNet [55, 73] introduced the delta rule with unit-norm key k_t ($\|k_t\|_2 = 1$), scalar write gate β_t, and $d_k \times d_k$ identity I_{d_k}:

$$S_t = (I_{d_k} - \beta_t k_t k_t^\top) S_{t-1} + \beta_t k_t v_t^\top .$$

Viewing such updates as online gradient descent, prior work also added heavy-ball momentum so that a single noisy update cannot overwrite established memories [10, 49]. Combining decay with the delta rule produced **Gated DeltaNet (GDN)** [74], adopted in Qwen3-Next [53]. **Kimi Delta Attention (KDA)** [78] then expanded GDN's scalar forget gate into a per-channel diagonal gate $F_t = \mathrm{diag}(f_t)$ with channel-wise decay vector $f_t \in (0, 1]^{d_k}$, and **GDN-2** [29] decoupled per-channel erase and write gates. Section 4 surveys further variants.

However, these architectures and their chunkwise WY training algorithms [12, 72, 74] were developed and tuned on NVIDIA GPUs (`fla` [71]). They map less well onto 2D systolic-array accelerators such as Google TPUs [22, 35, 36], for three reasons:

Small tiles and unpacked heads. GPU Tensor Cores operate on small fragments, so $d_k = d_v = 128$ with one key head per value head already reaches high occupancy. The TPU v6e/v7x Matrix Multiply Unit (MXU) is a 256×256 systolic array: a 128-wide key head pads half of every tile's key axis, and even when $d_k = 256$ (as at $D = 1792$), unpacked heads repeat projections, convolutions, and chunk solves for every value head.

Per-channel decay gates. In KDA and GDN-2, each key channel decays at its own rate. Kernels must rescale keys elementwise in 16-token sub-blocks to avoid float32 overflow. GPUs absorb this vector work cheaply; on TPUs, where matrix throughput exceeds vector throughput by roughly $128\times$, it stalls the kernel.

Per-value-head chunk solves and sequential fallbacks. The chunkwise WY algorithm [12, 74] partitions sequences into chunks of length C and solves a $C \times C$ unit lower-triangular system

per write-gate head and chunk. GPU kernels (`fla` [71]), XLA `solve_triangular` paths [1], and Tokamax's FP32 fallback [23] substitute sequentially on vector lanes. Open-source TPU kernels already express individual solves with MXU matrix products: Tokamax [23] uses a BF16 8-token Neumann series with Horner doubling ($C = 64$ for per-channel KDA, $C \in [16, 512]$ for scalar decay), and MaxText [1] uses exact block-doubling in its XLA GDN path. Yet existing architectures still assign independent write gates β_t to every value head, forcing H_v solves per chunk and unfused HBM round-trips (Ops 3–5 take about 45% of core stateless KDA time in Table 4).

Kernel tuning alone does not close this gap: fusing a GPU-shaped KDA layer into on-chip memory speeds it up by only $1.03\times$ on TPU v7x (Table 4), because the time goes to per-value-head chunk solves, per-channel decay rescaling, and padded MXU tiles rather than High-Bandwidth Memory (HBM) traffic alone. We therefore change the layer itself. **Systolic Linear Attention (SLA)** applies a block-diagonal, tile-grouped forget gate $F_t \in \mathbb{R}^{d_k \times d_k}$ to S_{t-1} and then performs the delta update with key $k_t \in \mathbb{R}^{d_k}$, packed value $v_t \in \mathbb{R}^{Rd_v}$ (stacking $R \in \{2, 4\}$ value heads per key head), and a write gate $\beta_t \in (0, 2)$ shared across the packed value heads:

$$S_t = F_t S_{t-1} + \beta_t k_t \big(v_t - (F_t S_{t-1})^\top k_t\big)^\top. \tag{1}$$

SLA defaults to $\mu = 0$ (no chunk-boundary momentum) and offers an optional Schur-guarded chunk-boundary momentum mode $\mu_c \le (1 - \varphi_c)/(1 + \varphi_c)$ kept in on-chip SRAM (VMEM); earlier 12.5B and 63.1B runs used fixed $\mu = 0.25$, while isolated multi-seed controls in Appendix I show no quality advantage over $\mu = 0$. SLA combines three changes (Section 2):

Wide keys and shared β_t across grouped value heads ($R\times$ solve reduction). SLA widens keys to $d_k = 256$ and shares both each key head (as in grouped value attention [52, 53]) and the scalar write gate $\beta_t \in (0, 2)$ across $R \in \{2, 4\}$ value heads (`SLA-R2`, `SLA-R4`). Unlike existing GVA linear layers that retain per-value-head (g, β_t) and still run H_v chunk solves, SLA packs the R value heads into full $256 \times Rd_v$ MXU tiles and cuts chunk solves by $R\times$ ($2.45\times$–$3.28\times$ faster core forward and $1.87\times$–$2.48\times$ faster full-layer F+B in an unconfounded same-shape control; Appendix Table 30).

Tile-grouped forget gate. SLA splits the 256 key channels into 4 tiles of 64 channels, each with one scalar decay rate. Intra-chunk scores then become four 64-wide MXU matmuls scaled by decay masks, with no VPU rescaling. This middle point between GDN's single rate per head ($G = 1$) and KDA's rate per channel ($G = d_k$) runs at near-scalar single-kernel VMEM speed while lowering language-modeling perplexity over $G = 1$ across all six paired 12.5B MoE and dense-LM seeds.

Fused VMEM exact inversion, backward reuse, and conditioning guarantee. Building on open-source MXU chunk solvers [1, 23], SLA fuses its $R\times$-fewer exact FP32 recursive block-doubling solves and custom-VJP backward reuse inside one VMEM kernel, proves a chunk-size-independent conditioning bound (Theorem 2.3), and shows that exact FP32 block-doubling holds 10^{-7} accuracy on real chunks through $C = 256$ while adding only 4.0%–7.7% to full-LM step time once values are packed (Appendix H).

We evaluate SLA on TPU v7x at the kernel, layer, serving, and context-parallel level, and across sparse Mixture-of-Experts (MoE) models [19, 57] from 6.5B (7.58B exact) to 63.1B nominal parameters on Dolma [60, 66] (Section 3; Table 1 summarizes the comparison with KDA). Building a single-kernel VMEM-fused per-channel KDA control separates the layer speedup against open-source Tokamax KDA into $1.28\times$–$1.41\times$ from 4:1 head packing, $1.10\times$–$1.12\times$ from VMEM fusion, and $1.07\times$–$1.20\times$ from $G = 4$ tile grouping (Appendix Table 35). At equal token count on 12.5B MoE, SLA matches KDA within seed noise and stays within about 0.1 perplexity of a wide-key per-channel KDA control while avoiding unclamped slow-channel long-context drift (which also appears in public `Kimi-Linear-48B-A3B` checkpoints and can be mitigated in KDA via zero-retraining gate clamping). SLA is therefore best understood per token as a hardware-efficient operating point complementary to richer update rules such as GDN-2 [29]: under equal chip-hours on 64-chip TPU v7x slices, both 1-kernel fused `SLA-R4` and a fused `SLA-R4 + GDN-2` hybrid beat safe-clipped open-source Tokamax KDA across all $3/3$ paired seeds at 12.5B and 63.1B (Table 9).

In the following sections, Section 2 reviews gated delta attention, its TPU bottlenecks, and SLA's design; Section 3 evaluates SLA from kernels and serving to MoE pre-training; and Section 4 covers related work.

Table 1: **SLA vs. KDA at a glance on TPU v7x.** Layer, prefill, equal-chip-hour, and 63.1B rows use `SLA-R4` and the 60.4B-token 12.5B rows use `SLA-R2`, all with exact inversion (`Exact Inv.`). Layer rows show 2-stage `SLA-R4` at $D = 1792$ (single-chip 1-kernel fused `SLA-R4`: 14.4 ms, 2.76× / 2.36×); at $D = 2048$, 2-stage `SLA-R4` is 2.1× / 1.5× faster and uses 49% / 52% less HBM. MoE throughput compares against MaxText KDA across nominal 6.5B (7.58B exact) to 63.1B scales (4%–21% vs. Tokamax KDA in Appendix Table 20b; the short 5B-token 31.5B [30.58B exact] `SLA-R2` run is 1.08×, 1.14× for `SLA-R4`). Equal-chip-hour rows compare 1-kernel fused iso-parameter `SLA-R4` ($\mu = 0$; brackets: fused `SLA-R4 + GDN-2`) against safe-clipped open-source Tokamax Mosaic KDA across $n = 3$ paired seeds on 64-chip TPU v7x slices. Clamping KDA's retention horizons to 8K raises its 64K `HotpotQA` score to 75.4%. Details in Tables 5–9.

Benchmark	KDA	**SLA**	**SLA vs. KDA**
Full layer, 32K fwd+bwd (2-stage; MaxText / Tokamax)	48.7 / 42.4 ms	20.2 ms	2.4× / 2.1× **faster**
Temporary HBM, 32K layer	9.1 / 9.3 GiB	3.7 GiB	60% / 61% **less**
131K prefill, 32-layer hybrid (Exact Inv., BF16)	2.66 s	1.62 s	1.64× **faster**
MoE training throughput, 6.5B–63.1B (vs. MaxText KDA)	1.00×	1.11×–1.24×	11%–24% **faster**
12.5B equal-chip-h PPL vs. Tokamax (42.7 chip-h, n=3)	19.52	19.15 [19.11]	−0.37 [−0.41] (3/3 **seeds)**
63.1B equal-chip-h PPL vs. Tokamax (37.3 chip-h, n=3)	26.56	25.08 [25.14]	−1.48 [−1.41] (3/3 **seeds)**
12.5B iso-token PPL vs. MaxText (60.4B tok, n=3)	15.60	15.58	−0.02 (within noise)
63.1B iso-token PPL vs. MaxText (120.8B tok, n=1)	12.60	12.45	−0.15
64K `HotpotQA`, 12.5B, no clamping	50.3%	77.9%	+27.6 **pts**

2 Hardware Bottlenecks and the SLA Design

To motivate SLA, we first review gated delta recurrences and their chunkwise WY training algorithm (Section 2.1), identify three hardware bottlenecks that GPU-centric designs hit on TPUs (Section 2.2), and then present SLA's mathematical formulation and chunk-boundary momentum (Section 2.3), its three algorithmic pillars and mathematical derivations (Section 2.4), and its single-kernel VMEM backward pass (Section 2.5), followed by empirical evaluation in Section 3.

2.1 Preliminaries: Gated Delta Recurrences and Chunkwise WY Parallelism

Given an input sequence $X = (x_1, \ldots, x_T) \in \mathbb{R}^{T \times D}$ of length T and model width D, a gated-delta linear attention layer with H_v value heads and H_k query-key heads maps each token $x_t \in \mathbb{R}^D$ ($t \in \{1, \ldots, T\}$) to L_2-normalized queries and keys $q_t, k_t \in \mathbb{R}^{d_k}$ ($\|q_t\|_2 = \|k_t\|_2 = 1$), values $v_t \in \mathbb{R}^{d_v}$, and head output $o_t \in \mathbb{R}^{d_v}$ (with $H_k = H_v$ in unpacked baselines). Rather than caching t past tokens, each head maintains a fixed-size associative state matrix $S_t \in \mathbb{R}^{d_k \times d_v}$ updated by applying a diagonal forget gate $F_t \in (0, 1]^{d_k \times d_k}$ followed by a rank-one delta update [10, 55, 65, 73] with scalar step size β_t:

Prior Work: Gated Delta Recurrence in GDN and KDA

$$
\begin{aligned}
S_t &= F_t S_{t-1} + \beta_t k_t \big(v_t - (F_t S_{t-1})^\top k_t\big)^\top \\
&= \big(I_{d_k} - \beta_t k_t k_t^\top\big) F_t S_{t-1} + \beta_t k_t v_t^\top, \\
o_t &= S_t^\top q_t.
\end{aligned} \tag{2}
$$

Here $(F_t S_{t-1})^\top k_t$ is the value currently retrieved for key k_t, and β_t sets how far it is moved toward v_t. Existing linear attentions differ only in (F_t, β_t) in Eq. (2), where we characterize F_t by its number of **decay groups** G (distinct decay rates across the d_k key channels of a head). Unweighted Linear Attention [37] sets $F_t = I_{d_k}$ without erasure ($S_t = S_{t-1} + k_t v_t^\top$), while DeltaNet [55, 73] and test-time learning [65, 70] retain $F_t = I_{d_k}$ with the rank-one delta update. Recent gated-delta architectures instantiate (F_t, β_t) using the logistic sigmoid $\sigma(\cdot)$ on scalar write-gate logit $\tilde{\beta}_t \in \mathbb{R}$ and $\mathrm{softplus}(u) = \ln(1 + e^u)$ on forget-gate logits $\tilde{f}_t \in \mathbb{R}^G$:

Gated DeltaNet (GDN). **GDN** [74] uses one scalar decay $\gamma_t \in (0, 1]$ per head ($G = 1$) together with a unit-interval scalar step size:

$$F_t = \gamma_t I_{d_k}, \qquad \beta_t = \sigma(\tilde{\beta}_t) \in (0, 1).$$

Kimi Delta Attention (KDA). **KDA** [78] assigns an independent decay rate to every key channel ($G = d_k$, $f_t = \exp(-\operatorname{softplus}(\tilde{f}_t))$):

$$F_t = \operatorname{diag}(f_t), \qquad f_t \in (0, 1]^{d_k}.$$

While standard GDN and KDA use $\beta_t = \sigma(\tilde{\beta}_t) \in (0, 1)$, negative-eigenvalue extensions [24] widen the step-size range to

$$\beta_t = 2\sigma(\tilde{\beta}_t) \in (0, 2),$$

so that $I_{d_k} - \beta_t k_t k_t^\top$ has eigenvalue $1 - \beta_t \in (-1, 1)$ along k_t (adopted by SLA in Section 2.3).

GDN-2. **GDN-2** [29] retains per-channel decay ($G = d_k$) and further replaces the scalar β_t with separate channel-wise erase and write gates.

Chunkwise WY parallelization. To saturate matrix units during training, chunkwise WY parallelization [12, 73, 74] partitions the sequence into T/C contiguous chunks of length C (Algorithm 1, written unified across sharing ratio $R = H_v/H_k$ and gate scale $s_\beta \in \{1, 2\}$, where unpacked GDN/KDA set $R = 1, s_\beta = 1$). Op 1 projects X to pre-convolution activations $\tilde{Q}, \tilde{K}, \tilde{V}$, gate logits $(\tilde{\beta}_t, \tilde{f}_t)_{t=1}^T$, and output-gate logits $(z_t)_{t=1}^T$ ($z_t \in \mathbb{R}^{Rd_v}$ per key head), followed by causal short convolutions, L_2 normalization into Q, K, and gate activations (Op 2). Within chunk $c \in \{1, \dots, T/C\}$ starting from boundary state $S_{c-1} \in \mathbb{R}^{d_k \times Rd_v}$ ($S_0 = 0$), let local token indices satisfy $1 \le j \le t \le C$, let $Q_c, K_c \in \mathbb{R}^{C \times d_k}$ and $V_c, O_c \in \mathbb{R}^{C \times Rd_v}$ stack the chunk's queries, keys, values, and readouts as rows, let $\beta_c = (\beta_1, \dots, \beta_C)^\top \in \mathbb{R}^C$ stack its write gates, and let I_C denote the $C \times C$ identity. Define

$$P_{t \leftarrow j} = \prod_{s=j+1}^{t} F_s \quad (P_{t \leftarrow t} = I_{d_k}), \qquad F_c = P_{C \leftarrow 0} = \prod_{s=1}^{C} F_s$$

as the cumulative diagonal decay from local token j to t and the total chunk decay, and let $\bar{K}_c, \bar{Q}_c, \hat{K}_c \in \mathbb{R}^{C \times d_k}$ denote the decayed key and query matrices whose t-th rows are $(P_{t \leftarrow 0} k_t)^\top$, $(P_{t \leftarrow 0} q_t)^\top$, and $(P_{C \leftarrow t} k_t)^\top$. Within a chunk, each token's write depends on what earlier tokens in the same chunk have already written, so the C writes cannot be computed independently. Collect them as rows of $E_c \in \mathbb{R}^{C \times Rd_v}$. Each row is what the token would write on its own, $\operatorname{diag}(\beta_c)(V_c - \bar{K}_c S_{c-1})$, minus the earlier tokens' writes, weighted by how strongly their decayed keys overlap with its own key. Using indicator $\mathbf{1}[\cdot] \in \{0, 1\}$, these overlaps define the strictly lower-triangular Gram matrix $M_c \in \mathbb{R}^{C \times C}$ (Op 3a, alongside lower-triangular intra-chunk query-key score matrix $L_c \in \mathbb{R}^{C \times C}$, Op 3b), so the chunk writes satisfy (derived in Section 2.4.2):

$$E_c = \operatorname{diag}(\beta_c)\,(I_C + M_c)^{-1}(V_c - \bar{K}_c S_{c-1}).$$

Computing $A_c = (I_C + M_c)^{-1}$ (Op 4) and forming the WY factors U_c, W_c (Op 5) runs in parallel across chunks, leaving a sequential scan over the T/C chunk boundaries for S_c and optional velocity buffer $Z_c \in \mathbb{R}^{d_k \times Rd_v}$ ($Z_0 = 0$, momentum $\mu \ge 0$; Op 6), parallel chunk readouts O_c (rows $o_t^\top$; Op 7a), and gated root-mean-square normalization ($\mathrm{RMSNorm}$, with elementwise product $\odot$) followed by the output projection to $Y \in \mathbb{R}^{T \times D}$ (Op 7b–c).

2.2 Three Hardware Bottlenecks on TPUs

TPUs pair a wide systolic MXU with a much slower VPU. On TPU v7x the MXU tile is 256×256 and the MXU:VPU throughput ratio is about $128\times$, compared with about $15\times$ on H100 (Table 2). This skew creates three bottlenecks for Algorithm 1:

Bottleneck 1: Narrow states and a fixed per-step scan cost. Op 6 carries the state across T/C chunk boundaries sequentially. On TPU v7x, each scan step costs a nearly fixed $\approx 12\,\mu$s regardless of chunk size: $8\times$ more FLOPs per step adds under 1% time (Figure 2). Scan cost is thus set by the number of steps and heads, not by matmul size. A 256-wide tile also processes a $d_k = 128$ state at the same cost as a $d_k = 256$ state, so 128-wide states leave half of each tile idle while paying the full per-step cost for every head.

Algorithm 1 Chunkwise WY Parallel Forward Pass

Require: Sequence $X \in \mathbb{R}^{T\times D}$, chunk size C, head ratio R, gate scale s_β, momentum μ ($(R, s_\beta, \mu) = (1, 1, 0)$ for GDN/KDA; $(R, 2, 0)$ recommended default with $R \in \{2, 4\}$ for SLA, or optional Schur-guarded $\mu \in (0, 0.25]$)

Ensure: Output $Y \in \mathbb{R}^{T\times D}$, final state $S_{T/C} \in \mathbb{R}^{d_k\times Rd_v}$

$\tilde{Q}, \tilde{K}, \tilde{V}, (\tilde{\beta}_t, \tilde{f}_t, z_t)_{t=1}^{T} \leftarrow \mathrm{LinearProj}(X)$ ▷ Op 1 (MXU)

$Q, K \leftarrow \mathrm{L2Norm}(\mathrm{Conv1D}(\tilde{Q}, \tilde{K}))$ ▷ Op 2a (VPU)

$V \leftarrow \mathrm{SiLU}(\mathrm{Conv1D}(\tilde{V}))$

$\beta_t \leftarrow s_\beta \sigma(\tilde{\beta}_t)$; $F_t \leftarrow \mathrm{diag}(e^{-\mathrm{softplus}(\tilde{f}_t)})$ ▷ Op 2b (VPU)

Partition Q, K, V, β, F into T/C chunks; $S_0 \leftarrow 0$, $Z_0 \leftarrow 0 \in \mathbb{R}^{d_k\times Rd_v}$

for all $c \in \{1, \dots, T/C\}$ **in parallel do**

  $P_{t\leftarrow j} \leftarrow \prod_{s=j+1}^{t} F_s \;\; \forall 0 \le j \le t \le C \quad (P_{t\leftarrow t} = I_{d_k})$

  $(\bar{K}_c, \bar{Q}_c, \hat{K}_c)_t \leftarrow P_{t\leftarrow 0}k_t, P_{t\leftarrow 0}q_t, P_{C\leftarrow t}k_t$ ▷ each row t

  $(M_c)_{t,j} \leftarrow \mathbf{1}[t > j]\, \beta_j k_t^\top P_{t\leftarrow j} k_j$ ▷ Op 3a: Gram

  $(L_c)_{t,j} \leftarrow \mathbf{1}[t \ge j]\, q_t^\top P_{t\leftarrow j} k_j$ ▷ Op 3b: Score

  $A_c \leftarrow (I_C + M_c)^{-1} \in \mathbb{R}^{C\times C}$ ▷ Op 4: Solve

  $U_c \leftarrow \mathrm{diag}(\beta_c) A_c V_c \in \mathbb{R}^{C\times Rd_v}$ ▷ Op 5a (MXU)

  $W_c \leftarrow \mathrm{diag}(\beta_c) A_c \bar{K}_c \in \mathbb{R}^{C\times d_k}$ ▷ Op 5b (MXU)

end for

for $c = 1$ **to** T/C **do**

  $E_c \leftarrow U_c - W_c S_{c-1}$ ▷ Op 6a (MXU)

  $Z_c \leftarrow \mu Z_{c-1} + \hat{K}_c^\top E_c$ ▷ Op 6b (MXU)

  $S_c \leftarrow F_c S_{c-1} + Z_c$

end for

for all $c \in \{1, \dots, T/C\}$ **in parallel do**

  $O_c \leftarrow \bar{Q}_c S_{c-1} + L_c E_c$ ▷ Op 7a (MXU)

end for

$Y \leftarrow \mathrm{OutProj}\big((\sigma(z_t) \odot \mathrm{RMSNorm}(o_t))_{t=1}^{T}\big)$ ▷ Op 7b–c

return $Y, S_{T/C}$

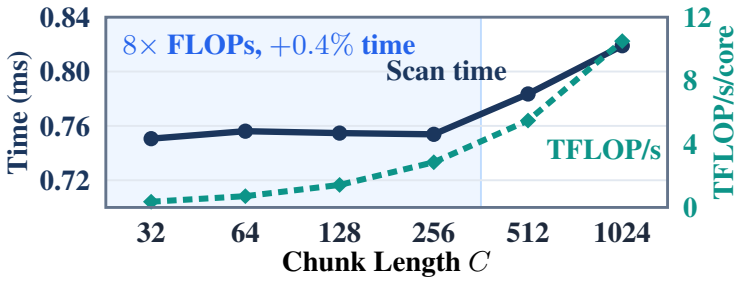

Figure 2: **Sequential state-scan cost on TPU v7x.** A 64-step scan of $(C, 256) \times (256, 256)$ matmuls takes nearly constant time for $C \in \{32, \dots, 256\}$, so a fixed per-step overhead, not matmul FLOPs, dominates the scan.

Bottleneck 2: Per-channel decay. KDA [78] and GDN-2 [29] give every key channel its own forget gate ($F_t = \mathrm{diag}(f_t)$). Let $\Lambda_{t,i} = \sum_{s=1}^{t} \log f_{s,i}$ be the cumulative decay of channel i within a chunk. Each Gram entry then weights channels differently:

$$(M_c)_{t,j} = \mathbf{1}[t > j]\, \beta_j \sum_i e^{\Lambda_{t,i} - \Lambda_{j,i}} k_{t,i} k_{j,i}.$$

To use the MXU, kernels fold $e^{\Lambda_{t,i}}$ and $e^{-\Lambda_{j,i}}$ into the keys, but $e^{-\Lambda_{j,i}}$ overflows FP32 for long chunks. KDA kernels therefore split each chunk into 16-token sub-blocks and re-center the keys on the VPU before every sub-block matmul, which stalls the kernel. In our KDA runs, about 1% of channels also learn near-unit forget gates, which hurts length extrapolation unless the gates are clamped (Section 3.5). GDN's scalar gate ($G = 1$) avoids both issues, at the cost of a single memory horizon per head.

Table 2: **Documented hardware specifications per chip** for NVIDIA H100 SXM and four Google TPU generations [6, 22, 36]. MXU:VPU is the ratio of peak BF16 matrix throughput to FP32 vector throughput; Ridge is the BF16 arithmetic-intensity ridge point (FLOPs/byte). *8-bit peak is FP8 on H100 and TPU v7x and INT8 on TPU v4–v6e.

Accelerator	Cores	Array Tile	BF16 TF/s	8-bit TF/s*	VPU TF/s	MXU:VPU	HBM (GB)	HBM GB/s	Ridge
H100 SXM	132 SM	16×16	989	1,979	67.0	14.8×	80	3,350	295
TPU v4	2	128×128	275	275	3.2	85.9×	32	1,200	229
TPU v5p	2	128×128	459	918	5.4	85.0×	95	2,765	166
TPU v6e	1	256×256	918	1,836	7.2	127.5×	32	1,600	574
TPU v7x	2	256×256	2,307	4,614	18.0	128.2×	192	7,373	313

Bottleneck 3: Per-value-head chunk solves and sequential fallbacks. Computing $A_c = (I_C + M_c)^{-1}$ (Op 4) by sequential substitution runs on vector lanes in GPU kernels (`fla` [71]), MaxText's XLA `solve_triangular` path [1], and open-source Tokamax's FP32 fallback [23]. Even when open-source TPU kernels move individual solves onto the MXU (Tokamax's BF16 Neumann–Horner path, validated for per-channel KDA at $C = 64$ [23], or MaxText's XLA block-doubling path for GDN [1]), per-value-head write gates still require H_v separate $C \times C$ solves per chunk and unfused HBM round-trips. Repeated across every chunk and all $H_v = 8$ value heads, Ops 3–5 account for about 45% of core intra-chunk and scan time in our stateless functional KDA audit (Table 4), keeping the practical chunk size near 64–128.

2.3 SLA: Addressing the Three Bottlenecks

Guided by the three systolic bottlenecks above, we now formulate SLA, which reshapes the state geometry, decay structure, and chunk inversion around the 256×256 systolic tile:

Definition 2.1 (SLA Recurrence)**.** Consider a linear attention layer with H_v value heads of width d_v and $H_k = H_v/R$ query-key heads of width $d_k = 256$ indexed by $h_k \in \{1, \dots, H_k\}$, where each query-key head is shared across $R \in \{2, 4\}$ value heads. At token t, each query-key head computes:

- **Shared query, key, and write gate:** L_2-normalized $q_t, k_t \in \mathbb{R}^{256}$ with $\|q_t\|_2 = \|k_t\|_2 = 1$, and a scalar step size $\beta_t = 2\sigma(\tilde{\beta}_t) \in (0, 2)$ ($s_\beta = 2$) that doubles the $(0, 1)$ range of standard GDN/KDA;
- **Packed values and output gates:** R value vectors $v_{t,1}, \dots, v_{t,R} \in \mathbb{R}^{d_v}$ stacked as $v_t = [v_{t,1}; \dots; v_{t,R}] \in \mathbb{R}^{D_{v,\text{pack}}}$ with packed value width $D_{v,\text{pack}} = Rd_v$, and output-gate vector $z_t = [z_{t,1}; \dots; z_{t,R}] \in \mathbb{R}^{D_{v,\text{pack}}}$ with slices $z_{t,m} \in \mathbb{R}^{d_v}$ ($m \in \{1, \dots, R\}$);
- **Four-tile forget gate:** a block-diagonal gate $F_t \in (0, 1]^{256 \times 256}$ that splits the 256 key channels into $G = 4$ contiguous tiles of $b = 64$ channels, with forget-gate logit $\tilde{f}_t = (\tilde{f}_{t,1}, \dots, \tilde{f}_{t,4}) \in \mathbb{R}^4$, all-ones vector $\mathbf{1}_{64} \in \mathbb{R}^{64}$, and one scalar decay $\gamma_{t,g} = \exp(-\operatorname{softplus}(\tilde{f}_{t,g})) \in (0, 1]$ on each tile $g \in \{1, 2, 3, 4\}$:

$$F_t = \operatorname{diag}(\gamma_{t,1}\mathbf{1}_{64},\ \gamma_{t,2}\mathbf{1}_{64},\ \gamma_{t,3}\mathbf{1}_{64},\ \gamma_{t,4}\mathbf{1}_{64}). \tag{3}$$

(i) Intra-chunk token recurrence. Maintaining a packed state $S_t = [S_{t,1}, \dots, S_{t,R}] \in \mathbb{R}^{256 \times Rd_v}$ initialized at $S_0 = 0$ and restarted at each chunk from the boundary state S_{c-1} of (ii), with slice $S_{t,m} \in \mathbb{R}^{256 \times d_v}$ and output slice $o_{t,m} \in \mathbb{R}^{d_v}$ ($o_t = [o_{t,1}; \dots; o_{t,R}]$) for value head $m \in \{1, \dots, R\}$, SLA decays the state, writes the prediction error, and normalizes each value head:

SLA Token-Level Recurrence ($d_k = 256,\ G = 4,\ b = 64$)

$$S_t = F_t S_{t-1} + \beta_t k_t \big(v_t - (F_t S_{t-1})^\top k_t\big)^\top, \tag{4}$$

$$o_{t,m} = \sigma(z_{t,m}) \odot \operatorname{RMSNorm}\big(S_{t,m}^\top q_t\big). \tag{5}$$

Here $(F_t S_{t-1})^\top k_t$ is the value currently stored for k_t, so Eq. (4) writes only the prediction error, and one rank-one write updates all R value heads.

(ii) Optional chunk-boundary momentum. Let $\tilde{S}_c$ be the state after unrolling Eq. (4) over the C tokens of chunk c from S_{c-1}, let $F_c = \prod_{t \in c} F_t$, and let the reset indicator $\kappa_c \in \{0, 1\}$ equal 0 if chunk c contains a document boundary (which also zeros cross-document carry F_c) and 1 otherwise. When enabled, an optional velocity buffer $Z_c \in \mathbb{R}^{256 \times D_{v,\text{pack}}}$ ($Z_0 = 0$) smooths boundary updates via

$$Z_c = \mu_c \kappa_c Z_{c-1} + (\tilde{S}_c - F_c S_{c-1}), \qquad S_c = F_c S_{c-1} + Z_c, \tag{6}$$

with $\mu_c = 0$ (recovering $S_c = \tilde{S}_c$) as the recommended default. When momentum is enabled ($\mu \in (0, 0.25]$), the per-chunk Schur guard $\mu_c = \min\big(\mu, (1 - \varphi_c)/(1 + \varphi_c)\big)$ with $\varphi_c = \prod_{t \in c} \max_g \gamma_{t,g}$ enforces unconditional chunk-to-chunk stability. (Our earlier 12.5B and 63.1B MoE runs in Section 3 were trained with fixed $\mu = 0.25$, which stayed well inside the stable region; controlled ablations below motivate demoting momentum from the default.)

Notation and operator taxonomy. Unless redefined locally, symbols follow Algorithm 1 and Definition 2.1; Appendix Table 13 collects them for quick reference (symbols used only inside one section, such as the roofline constants in Section 3.1, are defined where they first appear). Table 3 contrasts the algebraic operators and systolic hardware mappings of GDN [74], KDA [78], GDN-2 [29], **SLA** (`SLA-R2` / `SLA-R4`), and `SPD-SLA-W`, an optional extension of SLA that preconditions chunks with Σ_c to form whitened keys $\tilde{K}_c$ (rows $\tilde{k}_t^\top$) under step-size margin $\varepsilon_\beta > 0$ (Appendix A); `SPD-SLA-W` is not used in any core SLA result except the middle bar of Figure 7a and the conditioning analysis of Appendix A.

Table 3: Side-by-side algebraic and systolic hardware comparison across prior GPU gated-delta recurrences and our SLA family.

Design Axis	**GDN** [74]	**KDA** [78]	**GDN-2** [29]	**SLA** (`SLA-R2` / `SLA-R4`) (Ours)	`SPD-SLA-W` (Appendix A)
Key/Value Heads	Unpacked $H_k = H_v$	Unpacked $H_k = H_v$	Unpacked $H_k = H_v$	Packed GVA $H_k = H_v/R$ ($R \in \{2, 4\}$)	Packed GVA $H_k = H_v/R$
State Tile & Dtype	$S^{(h_k)} \in \mathbb{R}^{128 \times d_v}$ FP32 (50% pad)	$S^{(h_k)} \in \mathbb{R}^{128 \times d_v}$ FP32 (50% pad)	$S^{(h_k)} \in \mathbb{R}^{d_k \times d_v}$ FP32	$S \in \mathbb{R}^{256 \times D_{v,\text{pack}}}$ BF16 / Row-FP8 (no padding)	$S \in \mathbb{R}^{256 \times D_{v,\text{pack}}}$ BF16 / Row-FP8
Forget Gate F_t	Scalar $\gamma_t I_{d_k}$ ($G = 1$)	Channel $\mathrm{diag}(f_t)$ ($G = d_k$)	Decoupled ch. ($G = d_k$)	Tile-grouped $\bigoplus_{g=1}^{4} \gamma_{t,g} I_{64}$ ($G = 4$)	Tile-grouped $\bigoplus_{g=1}^{4} \gamma_{t,g} I_{64}$ ($G = 4$)
Step Size β_t	$(0, 1)$ per H_v	$(0, 1)$ per H_v	Ch. vector per H_v	$(0, 2)$ scalar per H_k (neg. eig.)	Clamped $(0, (2 - \varepsilon_\beta)/k_t^\top \tilde{k}_t]$
Chunk Gram M_c	$1\times$ MXU + scalar mask	16-row VPU pivots ($\lvert\Lambda_{j,i}\rvert > 88.7$)	Decoupled ch. broadcast	$G \times [C, b] \times [b, C]$ pure MXU ($b \geq 64$, 0 ovf.)	Oblique $k_t^\top P_{t \leftarrow j} \tilde{k}_j$ MXU
Chunk Inverse A_c	$H_v\times$ solves (seq. `fla`; XLA block-doub. MaxText; BF16 MXU Tokamax)	$H_v\times$ solves (MXU block-doub. Pallas; BF16 Neumann C=64 Tokamax)	$H_v\times$ solves (seq. in-block)	$H_k = H_v/R$ solves in VMEM: Exact FP32 (C=128) or 1-Step Neumann (App. H)	$H_k\times$ MXU inverse + `Polar-2+` (5 MM)
Momentum / 2nd-Ord.	None ($\mu = 0$, Euclidean)	None ($\mu = 0$, Euclidean)	None ($\mu = 0$, Euclidean)	Optional ($\mu = 0$ default; Schur-guarded $Z_c = \mu_c \kappa_c Z_{c-1} + \Delta S_c$)	Whitened keys $\tilde{K}_c$ (Eq. (14))
Output Norm	Head RMSNorm	Head RMSNorm	Head RMSNorm	Per-value-head $\mathrm{RMSNorm}$	Per-value-head $\mathrm{RMSNorm}$
VMEM / HBM Trips	6 HBM trips (unfused)	6 HBM trips (unfused wrapper)	6 HBM trips (6.8–14.9 GiB)	Fused SLA: 1 kernel in 64 MiB VMEM + custom VJP	Fits 64 MiB VMEM
Context Parallelism	Prefix tree, one $d_k \times d_k$ map per head	Prefix tree, one $d_k \times d_k$ map per head	Prefix tree, one $d_k \times d_k$ map per head	Prefix tree, one $[256, 256]$ map per key head ($2\times$ with momentum)	Same as SLA

Why chunk-level momentum, implementation, and caveats. Unlike token-level momentum [10], applying heavy-ball momentum [49] once per chunk boundary (Op 6b of Algorithm 1) leaves the intra-chunk WY solve unchanged. Specifically, in Op 6b, the net chunk state increment in Eq. (6) is

$$\Delta S_c = \tilde{S}_c - F_c S_{c-1} = \hat{K}_c^\top E_c.$$

The velocity buffer $Z_c \in \mathbb{R}^{d_k \times D_{v,\text{pack}}}$ has the same shape as S_c and stays in VMEM inside the fused kernel, adding zero extra HBM state traffic. On packed multi-document streams, intra-chunk Gram and score matrices apply an exact segment mask $\mathbf{1}[\mathrm{seg}_t = \mathrm{seg}_j]$, only the final document segment

writes to ΔS_c, and any document reset anywhere inside chunk c zeros both incoming state carry F_c and $\kappa_c = 0$, preventing mid-chunk velocity leakage. For a fixed chunk with a scalar gate ($G = 1$, chunk decay $\varphi \in (0, 1)$) and a real eigenvalue $\vartheta \in [-\varphi, \varphi]$ of the linear map from S_{c-1} to $\tilde{S}_c$, the lifted block $\left[\begin{smallmatrix} \vartheta & \mu \\ \vartheta - \varphi & \mu \end{smallmatrix}\right]$ has characteristic polynomial $p(\lambda) = \lambda^2 - (\vartheta + \mu)\lambda + \mu\varphi$ and meets Jury's Schur stability conditions ($\mu\varphi < 1$, $p(1) \geq (1 - \varphi)(1 - \mu) > 0$, $p(-1) > 0$) for all $\mu \in [0, 1)$. This argument does not cover $G = 4$, complex eigenvalues, or chunks that change over time, and we can construct unstable cases there at $\mu = 0.25$. Mapping the lifted eigenvalue spectrum on fixed repeated $G = 1$ chunks at $\mu = 0.25$ shows that complex-eigenvalue growth occurs only when high chunk retention ($\varphi > 0.773$, $\tau \gtrsim 500$ tokens at $C = 128$) coincides with near-reflection writes ($\beta > 1.73$ on a two-write chunk; random dense chunks are 0% unstable for $\bar{\beta} \leq 1.4$). A general sufficient stability condition is the per-chunk guard $\mu_c \leq (1 - \varphi_c)/(1 + \varphi_c)$ with $\varphi_c = \prod_{t \in c} \max_g \gamma_{t,g}$ (which at fixed $\mu = 0.25$ covers $\varphi_c < 0.6$). Our trained 12.5B `SLA-R4` models operate well inside the stable interior with mean write gates $\bar{\beta} \in [0.251, 0.348]$, and 24-layer state norms contract cleanly from 8K to 32K in the 63.1B model (Appendix F). With `Exact Inv.` and the same chunk boundaries, autoregressive decoding reproduces training up to rounding by updating Z_c once every C tokens. Under context parallelism, momentum couples (S_c, Z_c) into a $2d_k \times 2d_k$ affine map (Eq. (13)); Table 6 therefore reports both the uncoupled ($\mu = 0$, 256^2) and coupled ($\mu = 0.25$, 512^2) tree scans. Because momentum introduces chunk-size dependence and doubles the context-parallel transition, we evaluate both $\mu = 0$ and guarded μ_c controls directly in Section 3.

Appendix Table 14 compares $\mu = 0$ and $\mu = 0.25$ on a 12.5B-total (0.79B-active) MoE trained on 15.1B Dolma 3.5 tokens (7,200 steps, 2.1M tokens/step, 128 TPU v4 chips, seed 42). With 4:1 key sharing ($R = 4$), momentum lowers loss by 0.019–0.023 nats (though that $\mu = 0.25$ run used an 8,200-step cosine schedule evaluated at step 7,199); with 2:1 sharing ($R = 2$), it changes loss by at most $+0.004$ nats, and an input-dependent μ degrades loss by $+0.026$ nats. To isolate momentum without schedule or fusion confounders, we ran a 49-point Schur stability sweep ($\varphi \in [0.4, 0.99]$, $\beta \in [0.3, 1.99]$) and an isolated 3-seed dense LM ablation (274.7M params, 104.9M FineWeb-Edu tokens, identical `SLA-R4` kernel; Appendix I, Table 37). On the 49-point grid, fixed $\mu = 0.25$ exhibits 5 unstable (φ, β) pairs with spectral radius up to 1.090 (224.7× 64-chunk norm growth), whereas the per-chunk guard $\mu_c = \min(\mu, (1 - \varphi_c)/(1 + \varphi_c))$ eliminates every unstable pair ($5 \to 0$). In the 3-seed LM ablation, trained gates stay deep inside the stable interior ($\bar{\beta} \in [0.76, 0.79]$, $p_{95}(\beta) \leq 1.20$, $\bar{\varphi}_c \leq 0.057$, $p_{95}(\varphi_c) \leq 0.26$, guard activation 0%), and $\mu = 0$ yields the lowest mean validation perplexity (46.49 vs. 46.57 fixed and 46.53 guarded; paired $+0.09 \pm 0.36$, $t = +0.41$) and 16K perplexity (45.90 vs. 46.27 and 46.19). Because chunk momentum adds chunk-size dependence, doubles context-parallel transition size to 512×512 (Section 3.3), and requires the 64 MiB VMEM cap at $D = 1792$ without improving isolated perplexity, we demote momentum from SLA's default ($\mu = 0$ recommended; Schur-guarded μ_c optional) and note transparently where our earlier 12.5B/63.1B MoE checkpoints used fixed $\mu = 0.25$.

2.4 The Three Algorithmic Pillars of SLA

Each pillar targets one of the bottlenecks in Section 2.2 (Figure 1c–e; detailed diagrams in Figure 3).

2.4.1 Pillar 1: Wide Keys with Grouped Value Heads

A key dimension below 256 still occupies a full 256-wide tile on TPU v6e/v7x, so SLA widens keys to $d_k = 256$, which costs no extra cycles in the inter-chunk scan. Prior linear grouped-value layers (such as Qwen3-Next [53] and `fla` [71]) share query-key heads across value heads at $d_k = 128$, yet keep independent decay and write gates (g, β_t) per value head, which still forces H_v separate Gram matrices and H_v chunk solves per chunk. SLA instead shares both the 256-wide query-key head and the scalar write gate $\beta_t \in (0, 2)$ across $R = H_v/H_k \in \{2, 4\}$ packed value heads (together with per-value-head gated normalization [52], Eq. (5)). Packing these R value heads side by side forms a contiguous $S_t \in \mathbb{R}^{256 \times Rd_v}$ state that fills whole MXU tiles without padding, while cutting QK projections, short convolutions, M_c, L_c, and the chunk solve A_c by $R\times$. On the identical exact FP32 $C = 128$ kernel ($D = 2048$), packing $R = 4$ value heads per key head runs 2.92× faster forward (2.73× F+B) than $R = 1$; holding head geometry fixed at $H_k = 2, d_k = 256, H_v = 8$ and isolating only the shared write gate β_t (2 solves vs. 8 solves per chunk) yields 2.45×–3.28× faster core forward, 1.87×–2.48× faster full-layer F+B, and 1.68× faster small-LM pre-training

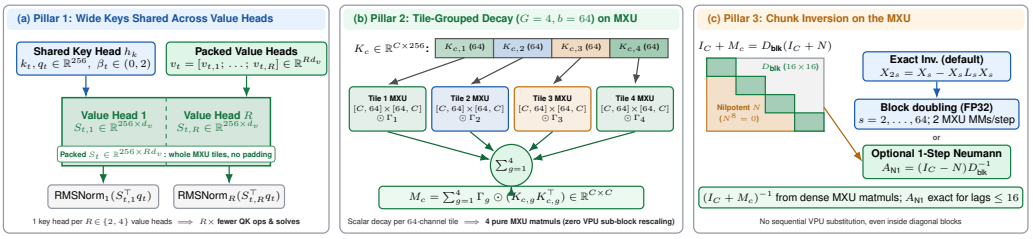

Figure 3: **Detailed architectural diagram of the three mathematical pillars of SLA. (a) Pillar 1:** Widening the key dimension to $d_k = 256$ and sharing each key head h_k and write gate $\beta_t \in (0, 2)$ across $R \in \{2, 4\}$ value heads forms a dense $256 \times D_{v,\text{pack}}$ state that fills whole 256×256 MXU tiles without padding while cutting QK projections, short convolutions, and chunk solves by $R\times$. **(b) Pillar 2:** Partitioning the $d_k = 256$ key channels into $G = 4$ contiguous tiles of $b = 64$ channels decomposes the intra-chunk coupling matrix M_c into four 64-wide MXU matrix multiplications weighted by scalar decay masks Γ_g, eliminating VPU sub-block rescaling. **(c) Pillar 3:** Because β_t is shared across R value heads, SLA runs $R\times$ fewer chunk solves inside VMEM: default FP32 recursive block-doubling (`Exact Inv.`, building on open-source MXU solvers [1, 23]) with custom-VJP reuse, or the optional 1-step Neumann inverse ($I_C + M_c = D_{\text{blk}}(I_C + N)$ with 16×16 diagonal blocks, $N^8 = 0$ at $C = 128$; Section 2.4.3).

(Appendix H). Head sharing itself is prior work; our contribution is widening d_k to the systolic tile and sharing β_t across the R packed value heads so one chunk solve serves all R heads.

Derivation of GVA systolic invariance. Fix total value width $D_v = H_v d_v$ and set $H_k = H_v/R$ with $D_{v,\text{pack}} = Rd_v$. Each key head $h_k \in \{1, \ldots, H_k\}$ maintains state $S^{(h_k)} \in \mathbb{R}^{d_k \times D_{v,\text{pack}}}$, so total state elements are invariant to R:

$$H_k d_k (Rd_v) = H_v d_k d_v.$$

Inter-chunk state scan FLOPs per token are likewise invariant to R:

$$H_k \big(2d_k D_{v,\text{pack}} \text{ [erase]} + 2d_k D_{v,\text{pack}} \text{ [write]} + 2d_k D_{v,\text{pack}} \text{ [readout]}\big) = 6H_v d_k d_v.$$

Because β_t is scalar per key head h_k, QK projections ($2DH_k d_k$ parameters), QK causal `Conv1D` (for kernel width $K_{\text{conv}} = 4$, $2K_{\text{conv}} H_k d_k$ weights), chunk Gram matmuls ($4CH_k d_k$ FLOPs/token), and $[C, C]$ WY/Neumann solves (H_k systems per chunk) scale linearly with $H_k = H_v/R$ and shrink by $R\times$. Reshaping $\tilde{S}_t^\top q_t \in \mathbb{R}^{Rd_v}$ into R slices of width d_v and applying independent $\mathrm{RMSNorm}$ to all $H_v = H_k R$ value heads prior to the sigmoid output gate $\sigma(z_t) \in (0, 1)^{Rd_v}$ (slices $\sigma(z_{t,m})$ in Eq. (5)) retains H_v independent head-scale degrees of freedom while keeping $[256, D_{v,\text{pack}}]$ MXU tiles contiguous. □

2.4.2 Pillar 2: Tile-Grouped Forget Gate and Packed WY Equivalence

Bottleneck 2 arises because the decay cannot be pulled out of key–key inner products when every channel has its own rate. SLA splits each query and key into the same $G = 4$ groups of $b = 64$ channels as the forget gate in Eq. (3), $q_t = [q_{t,1}; \ldots; q_{t,4}]$ and $k_t = [k_{t,1}; \ldots; k_{t,4}]$. Between tokens j and t, group g decays by a single number $\prod_{s=j+1}^t \gamma_{s,g} \leq 1$, which moves outside the group's inner product. Stacking group-g queries and keys as $Q_{c,g}, K_{c,g} \in \mathbb{R}^{C \times 64}$ and defining causal masks $(\Psi_g)_{t,j} = \mathbf{1}[t \geq j] \prod_{s=j+1}^t \gamma_{s,g}$ and $(\Gamma_g)_{t,j} = \mathbf{1}[t > j]\, \beta_j (\Psi_g)_{t,j}$, both intra-chunk matrices factor across tiles as

$$M_c = \sum_{g=1}^4 \Gamma_g \odot \big(K_{c,g} K_{c,g}^\top\big), \qquad L_c = \sum_{g=1}^4 \Psi_g \odot \big(Q_{c,g} K_{c,g}^\top\big),$$

where each term is a $[C, 64] \times [64, C]$ MXU matmul followed by an elementwise scaling by decay factors in $[0, 1]$, computed only for $t \geq j$, so there is no overflow and no need for VPU sub-blocks. This places $G = 4$ at the multi-timescale quality and single-kernel VMEM speed Pareto knee: scalar decay ($G = 1$, as in GDN) runs at nearly the same single-kernel VMEM speed as $G = 4$ (2.7%–4.0%

difference; Table 35) because either 1 or 4 gate vectors per head reside entirely in VMEM vector registers without materializing $[B, H_k, T/C, C, C, d_k]$ pair-decay tensors in HBM, but $G = 1$ limits each head to a single retention timescale. On 12.5B MoE ($n = 6$ paired seeds; Table 10), $G = 4$ improves validation perplexity on all $6/6$ seeds (ΔPPL $= -0.34, -0.06, -0.06, -0.66, -0.07, -0.21$; mean -0.23 ± 0.24, two-sided $p = 0.060$, exact sign test $p = 0.031$, while $G = 1$ leads unclamped synthetic 32K 8-Needle at 59.6% vs. 35.4%); to separate $G = 4$ from $G = 1$ cleanly without MoE routing variance, our 6-seed paired dense LM control (274.7M params, 104.9M tokens, identical `SLA-R4` kernel; Appendix I, Table 37) shows $G = 4$ beating $G = 1$ on $6/6$ seeds across every metric (-1.15 ± 0.49 Val PPL, $t = -5.78$; -1.75 ± 0.61 16K tail PPL, $t = -7.01$), capturing 78.5% of the $G = 1 \to G = 256$ gain while beating unclamped $G = 256$ at 16K. Conversely, on a synthetic 2-layer multi-timescale associative recall diagnostic, $G = 256$ solves the task reliably while $G \in \{1, 2, 4\}$ remain bimodal or underfit (Table 37), so $G \in \{1, 4\}$ acts as a clean hardware-expressivity knob with $G = 4$ as the language-modeling default. Architecturally, unpacked GDN ($H_k = 8, G = 1$) gives $H_k G = 8$ decay timescales per layer, whereas 4:1 head sharing ($H_k = 2$) with $G = 1$ leaves only 2; setting $G = 4$ ($b = 64$) restores $H_k G = 8$ timescales per layer in `SLA-R4` (16 in `SLA-R2`) at near-scalar single-kernel VMEM speed.

Isolating this VMEM-resident decayed score stage at $d_k = 256$ on TPU v7x quantifies why $G = 4$ ($b = 64$) remains inexpensive while finer sub-tiles degrade sharply. Analytically, at TPU v7x's 128× MXU:VPU throughput skew, forming G causal decay masks costs $128GC^2$ MXU-equivalent ops against $2C^2 d_k$ Gram matmul FLOPs; these balance at $G^\star = 2d_k/128 = 4$ for $d_k = 256$ ($G^\star = 8$ at $d_k = 512$), whereas per-channel decay ($G = d_k$) spends 98.5% of intra-chunk work on the VPU. Under a decaying-memory capacity model with log-uniform lag demand from 8K to 1M tokens, unequal tile widths gain under 0.1% and $G > 4$ gains under 0.3%, so hardware balance rather than recall capacity pins $G = 4$. At $C = 256$, an undecayed $[256, 256] \times [256, 256]$ matmul takes 0.303 μs per chunk (1.00×), whereas masked group sums at $G = 1$ ($b = 256$), $G = 2$ ($b = 128$), $G = 4$ ($b = 64$), and $G = 16$ ($b = 16$) take 0.306 μs (1.01×), 0.324 μs (1.07×), 0.461 μs (1.52×), and 1.269 μs (4.19×), respectively (0.197, 0.201, 0.209, 0.310, and 0.730 μs at $C = 128$). Because score construction accounts for only 10%–15% of fused-layer FLOPs once $R \in \{2, 4\}$ value heads are packed into the downstream chunk-update and readout matmuls, the 1.52× stage cost at $G = 4$ adds under 4% to full-layer latency, whereas narrower sub-tiles ($G = 16$) incur steep narrow-reduction and GC^2 VPU masking overhead on the 256×256 MXU.

Proof of tile-grouped packed WY equivalence. Consider chunk $c \in \{1, \dots, T/C\}$ for a single key head $h_k \in \{1, \dots, H_k\}$ serving $R = H_v/H_k$ packed value heads of total width $D_{v,\text{pack}} = Rd_v$, starting from incoming state $\tilde{S}_0 = S_{c-1} \in \mathbb{R}^{d_k \times D_{v,\text{pack}}}$. Each token $t \in \{1, \dots, C\}$ has $G = 4$ tile-grouped diagonal decay $F_t = \bigoplus_{g=1}^{4} \gamma_{t,g} I_{64}$ (diagonal entries in $(0, 1]$), query and key $q_t, k_t \in \mathbb{R}^{d_k}$, packed value $v_t \in \mathbb{R}^{D_{v,\text{pack}}}$, and scalar step size $\beta_t \in (0, 2)$ shared across the R packed value heads of key head h_k.

Let $P_{t \leftarrow j} = \prod_{s=j+1}^{t} F_s$ denote the cumulative diagonal transition from step j to step t ($P_{t \leftarrow t} = I_{d_k}$). Unrolling Eq. (4) ($\bar{S}_t = F_t \tilde{S}_{t-1}$, $\hat{v}_t = \bar{S}_t^\top k_t$, $\tilde{S}_t = \bar{S}_t + \beta_t k_t (v_t - \hat{v}_t)^\top$) from $\tilde{S}_0 = S_{c-1}$ gives

$$\tilde{S}_t = P_{t \leftarrow 0} S_{c-1} + \sum_{j=1}^{t} \beta_j \left(P_{t \leftarrow j} k_j\right) u_j^\top, \tag{7}$$

where $u_t = v_t - \hat{v}_t \in \mathbb{R}^{D_{v,\text{pack}}}$ is the post-forget delta residual at local step t. Multiplying $\bar{S}_t = F_t \tilde{S}_{t-1}$ on the left by $k_t^\top$ substitutes $F_t P_{t-1 \leftarrow j} = P_{t \leftarrow j}$ into every prior term:

$$u_t^\top = v_t^\top - k_t^\top P_{t \leftarrow 0} S_{c-1} - \sum_{j=1}^{t-1} \beta_j \left(k_t^\top P_{t \leftarrow j} k_j\right) u_j^\top. \tag{8}$$

With tile cumulative log-decay $\Lambda_{t,g} = \sum_{s=1}^{t} \log \gamma_{s,g}$, the inner product $k_t^\top P_{t \leftarrow j} k_j = \sum_{g=1}^{4} e^{\Lambda_{t,g} - \Lambda_{j,g}} \langle k_{t,g}, k_{j,g} \rangle$ evaluates via the four masked $[C, 64] \times [64, C]$ MXU matmuls above without VPU sub-blocks. Because β_j is scalar per key head h_k, the strictly lower-triangular coupling matrix $M_c \in \mathbb{R}^{C \times C}$ defined by

$$(M_c)_{t,j} = \mathbf{1}[t > j] \, \beta_j \left(k_t^\top P_{t \leftarrow j} k_j\right) \tag{9}$$

is shared across all $D_{v,\text{pack}}$ value columns and requires inverting a single $C \times C$ unit lower-triangular system $A_c = (I_C + M_c)^{-1}$ per key head.

Stacking chunk values $V_c \in \mathbb{R}^{C \times D_{v,\text{pack}}}$ (row t equal to $v_t^\top$), decayed keys $\bar{K}_c \in \mathbb{R}^{C \times d_k}$ (row t equal to $(P_{t\leftarrow 0} k_t)^\top$), and residuals $u_{1:C} \in \mathbb{R}^{C \times D_{v,\text{pack}}}$ (row t equal to $u_t^\top$) yields β-scaled WY factors and chunk writes:

$$U_c = \operatorname{diag}(\beta_c) A_c V_c \in \mathbb{R}^{C \times D_{v,\text{pack}}}, \qquad W_c = \operatorname{diag}(\beta_c) A_c \bar{K}_c \in \mathbb{R}^{C \times d_k},$$
$$E_c = \operatorname{diag}(\beta_c)\, u_{1:C} = U_c - W_c S_{c-1}.$$

Letting the carried keys $\hat{K}_c \in \mathbb{R}^{C \times d_k}$ have row j equal to $(P_{C\leftarrow j} k_j)^\top$ and $F_c = P_{C\leftarrow 0}$, the end-of-chunk state and chunk increment satisfy

$$\tilde{S}_C = F_c S_{c-1} + \hat{K}_c^\top (U_c - W_c S_{c-1}), \qquad \Delta S_c = \tilde{S}_C - F_c S_{c-1} = \hat{K}_c^\top E_c \in \mathbb{R}^{d_k \times D_{v,\text{pack}}}.$$

(When optional oblique key whitening $\tilde{k}_j$ is enabled in `SPD-SLA-W`, Appendix A, replacing the write key k_j with $\tilde{k}_j$ in M_c and $\hat{K}_c$ yields the identical WY dual equations.) Finally, the chunk-level velocity $Z_c = \mu \kappa_c Z_{c-1} + \Delta S_c$ and $S_c = F_c S_{c-1} + Z_c$ (Eq. (6), with reset indicator $\kappa_c \in \{0, 1\}$ equal to 0 if chunk c contains a document boundary) act solely at chunk boundaries, preserving timestep key-value pairing inside every chunk and matching token-by-token decoding that accumulates ΔS_c over C steps before applying Eq. (6). □

Intra-tile conditioning and orthogonal rotation absorption. Because each group's decay is a scalar multiple of the identity, it preserves conditioning within each tile and commutes with rotary rotations inside the group:

Observation 2.2 (Intra-Tile Conditioning and Orthogonal Rotation Absorption)**.** Within each group g, the decay block $\gamma_{t,g} I_{d_k/G}$ has unit condition number $\kappa_2(\gamma_{t,g} I_{d_k/G}) = 1$ (preserving intra-tile cosine similarity) and commutes with any rotation acting within each tile, $\text{Rot} \in \text{SO}(d_k/G)^G$, such as Givens or rotary rotations, leaving the grouped Gram computation of Pillar 2 unchanged.

Proof. For a per-channel diagonal decay $F_t = \operatorname{diag}(f_t)$ with $f_{t,i} \in (0, 1)$ and cumulative log-decay $\Lambda_{t,i} = \sum_{s=1}^{t} \log f_{s,i}$, the chunk decay $F_c = P_{C\leftarrow 0} = \operatorname{diag}\big((\exp(\Lambda_{C,i}))_{i=1}^{d_k}\big)$ has singular values $\exp(\Lambda_{C,i})$ and condition number $\kappa_2(F_c) = \exp(\max_i \Lambda_{C,i} - \min_i \Lambda_{C,i})$; write $\Delta\Lambda_C = \max_i \Lambda_{C,i} - \min_i \Lambda_{C,i}$. In BF16 (7 mantissa bits, machine epsilon $2^{-7} \approx 7.8 \times 10^{-3}$, unit roundoff $2^{-8} \approx 3.9 \times 10^{-3}$, exponent range $\approx [2^{-126}, 2^{127}]$), relative channel scaling across $P_{t\leftarrow 0} P_{j\leftarrow 0}^{-1}$ loses all mantissa significance across fast vs. slow channels once $\Delta\Lambda_C \gtrsim 5.5$ in BF16 ($\kappa_2 \approx 2^8$) or $\Delta\Lambda_C \gtrsim 16$ in FP32 accumulators ($\kappa_2 \approx 8.9 \times 10^6 \approx 2^{23}$), and $\exp(\pm\Lambda_{j,i})$ overflows or underflows FP32 when $|\Lambda_{j,i}| > 88.7$. Under tile-grouped decay $F_t = \bigoplus_{g=1}^{4} \gamma_{t,g} I_{64}$, each 64-channel tile block is a scalar multiple of I_{64} with $\kappa_2(\gamma_{t,g} I_{64}) = 1$, preserving exact intra-tile cosine similarity.

For any 2×2 block-diagonal (Givens/rotary) rotation with per-pair angles $\theta_t = (\theta_{t,1}, \ldots, \theta_{t,d_k/2}) \in \mathbb{R}^{d_k/2}$ indexed by pair $l \in \{1, \ldots, d_k/2\}$,

$$\text{Rot}(\theta_t) = \bigoplus_{l=1}^{d_k/2} \begin{pmatrix} \cos\theta_{t,l} & -\sin\theta_{t,l} \\ \sin\theta_{t,l} & \cos\theta_{t,l} \end{pmatrix} \in \text{SO}(2)^{d_k/2},$$

the cumulative rotation $\prod_{t=1}^{C} \text{Rot}(\theta_t) = \text{Rot}(\Theta_C)$ with $\Theta_t = \sum_{s=1}^{t} \theta_s$ has all singular values equal to 1 ($\kappa_2 = 1$). Since $\text{Rot}(\theta_s)$ commutes with $G = 4$ scalar tile multiples (64 channels = 32 contiguous 2×2 pairs per tile), for tile-grouped F_s we have

$$k_t^\top \Big(\prod_{s=j+1}^{t} F_s \text{Rot}(\theta_s) \Big) k_j = \big(\text{Rot}(\Theta_t)^\top k_t\big)^\top P_{t\leftarrow j} \big(\text{Rot}(\Theta_j)^\top k_j\big).$$

Replacing (q_t, k_t) with the rotated queries and keys $(\text{Rot}(\Theta_t)^\top q_t,\ \text{Rot}(\Theta_t)^\top k_t)$ before the chunk kernel reduces the intra-chunk transition to pure $G = 4$ real tile decay, preserving the unchanged systolic WY kernel up to a final boundary unrotation $\text{Rot}(\Theta_C)$ on outgoing state. □

2.4.3 Pillar 3: Fused MXU Chunk Inversion and Conditioning Certificate

Existing MXU solvers [Prior / Existing] **and SLA's fused integration.** While GPU kernels (`fla` [71]), MaxText's XLA `solve_triangular` path, and open-source Tokamax's FP32 fallback compute $A_c = (I_C + M_c)^{-1}$ by sequential substitution, open-source TPU kernels already express unit lower-triangular inversion with MXU matrix multiplications. For BF16 inputs, Tokamax [23] inverts 8-token diagonal blocks via a Neumann series and combines blocks by Horner doubling with FP32 accumulation (validated at $C = 64$ for per-channel KDA, and $C \in [16, 512]$ for scalar decay). Open-source MaxText [1] implements exact recursive block-doubling in its XLA GDN path (and our Pallas KDA baseline uses the same block-doubling formulation). SLA adopts FP32 recursive block-doubling (`Exact Inv.`, doubling 2×2 blocks six times to $C = 128$, Figure 1e) and contributes three elements on top: (1) executing $R\times$ fewer exact solves directly inside a single VMEM kernel with custom-VJP inverse reuse (Section 2.5); (2) a chunk-size-independent conditioning certificate (Theorem 2.3); and (3) TPU v7x numerical verification showing that FP32 block-doubling holds worst-case relative error at 1.4×10^{-7}–3.0×10^{-7} across $C \in \{32, 64, 128, 256\}$ on real Qwen3.5 and high-coherence SLA chunks, whereas Tokamax-style BF16 blocked Neumann reaches 10.6–12.7 relative error at $C = 128$ and unblocked Neumann products overflow (Appendix H). Let $H_{t,j} = \sum_{g=1}^{G} e^{-|\Lambda_{t,g} - \Lambda_{j,g}|} \langle k_{t,g}, k_{j,g} \rangle$ denote the symmetric decayed key Gram matrix, so $M_c = \mathrm{tril}_{-1}(H)\,\mathrm{diag}(\beta_c)$. Theorem 2.3 certifies that exact block-doubling is unconditionally well-conditioned at any chunk size or key coherence:

Theorem 2.3 (Chunk-WY Conditioning and Recursive Sub-Block Certificate)**.** *For unit keys $\|k_t\|_2 = 1$, non-increasing cumulative log-decays $\Lambda_{t,g}$ ($G \in \{1, \ldots, d_k\}$), and $D = \mathrm{diag}(\beta_c)$ with $\beta_t \in (0, \beta_{\max}]$, $\beta_{\max} < 2$, the Laplace-Gram Schur product yields $H \succeq 0$ with unit diagonal $H_{t,t} = 1$, and $B = D^{1/2}(I_C + M_c)D^{-1/2}$ satisfies $B + B^\top = 2I_C - D + D^{1/2} H D^{1/2} \succeq (2 - \beta_{\max}) I_C$. Consequently,*

$$\left\|B^{-1}\right\|_2 \;\le\; \frac{2}{2 - \beta_{\max}}, \qquad \left\|D(I_C + M_c)^{-1}\right\|_2 \;\le\; \frac{2\beta_{\max}}{2 - \beta_{\max}},$$

strictly independent of chunk size C, key coherence, $\beta_{\min}$, and $\|M_c\|_2$. Moreover, because every contiguous index subset $S \subseteq \{1, \ldots, C\}$ induces a positive semidefinite principal submatrix $H_{S,S} \succeq 0$ with unit diagonal, every intermediate block-doubling factor X_s ($s = 2, 4, \ldots, C/2$) in symmetric B-coordinates obeys the identical bound $\left\|B_{S,S}^{-1}\right\|_2 \le \frac{2}{2-\beta_{\max}}$, so each doubling step $X_{2s} = X_s - X_s L_s X_s$ amplifies local rounding perturbations by at most $\mathcal{O}\big((2 - \beta_{\max})^{-2} \|L_s\|_2\big)$ rather than exponentially in C.

Proof. For each decay group $g \in \{1, \ldots, G\}$, $\Lambda_{t,g} = \sum_{s=1}^{t} \log \gamma_{s,g}$ is monotonically non-increasing in t, so $\Lambda_{t,g} - \Lambda_{j,g} = -|\Lambda_{t,g} - \Lambda_{j,g}|$ for all $t \ge j$. On $\mathbb{R}$, the Laplace kernel $\mathcal{K}_{\text{Lap}}(x, y) = e^{-|x-y|} = \frac{1}{\pi} \int_{-\infty}^{\infty} \frac{\cos(\omega(x-y))}{1+\omega^2}\, d\omega$ is positive semidefinite, hence $E_g = [e^{-|\Lambda_{t,g} - \Lambda_{j,g}|}]_{t,j=1}^{C} \succeq 0$. Since each group key Gram matrix $K_{c,g} K_{c,g}^\top \succeq 0$, the Schur product theorem gives $E_g \odot (K_{c,g} K_{c,g}^\top) \succeq 0$. Summing over $g = 1, \ldots, G$ proves $H = \sum_{g=1}^{G} E_g \odot (K_{c,g} K_{c,g}^\top) \succeq 0$, with $H_{t,t} = \sum_{g=1}^{G} \|k_{t,g}\|_2^2 = \|k_t\|_2^2 = 1$ (covering scalar $G = 1$, tile-grouped $G = 4$, and per-channel $G = d_k$ decay).

Writing $M_c = \mathrm{tril}_{-1}(H) D$, conjugating $I_C + M_c$ by $D^{1/2}$ yields $B = D^{1/2}(I_C + M_c)D^{-1/2} = I_C + D^{1/2}\,\mathrm{tril}_{-1}(H) D^{1/2}$. Because $\mathrm{tril}_{-1}(H) + \mathrm{tril}_{-1}(H)^\top = H - I_C$, symmetrizing B gives

$$B + B^\top = 2I_C + D^{1/2}(H - I_C)D^{1/2} = 2I_C - D + D^{1/2} H D^{1/2} \;\succeq\; (2 - \beta_{\max}) I_C,$$

since $D^{1/2} H D^{1/2} \succeq 0$ and $2I_C - D \succeq (2 - \beta_{\max}) I_C$. For any unit vector $u \in \mathbb{R}^C$, Cauchy-Schwarz yields $\|Bu\|_2 \ge u^\top B u = \frac{1}{2} u^\top (B + B^\top) u \ge 1 - \frac{1}{2}\beta_{\max} > 0$, establishing $\sigma_{\min}(B) \ge 1 - \frac{1}{2}\beta_{\max}$ and $\left\|B^{-1}\right\|_2 \le \frac{2}{2-\beta_{\max}}$. The WY write operator $D(I_C + M_c)^{-1} = D^{1/2} B^{-1} D^{1/2}$ applied in Ops 5a–5b therefore satisfies $\left\|D(I_C + M_c)^{-1}\right\|_2 \le \left\|D^{1/2}\right\|_2^2 \left\|B^{-1}\right\|_2 \le \frac{2\beta_{\max}}{2-\beta_{\max}}$, independent of C, key coherence, $\beta_{\min}$, and $\|M_c\|_2$. Restricting to any contiguous block $S \subseteq \{1, \ldots, C\}$ preserves $H_{S,S} \succeq 0$ and $(H_{S,S})_{t,t} = 1$, so the identical symmetrization gives $B_{S,S} + B_{S,S}^\top \succeq (2 - \beta_{\max}) I_{|S|}$ and $\left\|B_{S,S}^{-1}\right\|_2 \le \frac{2}{2-\beta_{\max}}$ for every intermediate diagonal inverse X_s formed during block-doubling. Across 600 random and coherent configurations ($G \in \{1, 4, 256\}$, $C \in \{64, 128, 256\}$), the worst empirical ratio $\left\|B^{-1}\right\|_2 / [2/(2 - \beta_{\max})]$ is $0.973 \le 1$.

Two practical caveats separate the certificate from the default kernel implementation. First, $\beta_t = 2\sigma(\tilde{\beta}_t) \in (0, 2)$ has no hard architectural ceiling $\beta_{\max} < 2$ unless clamped: with identical keys and $\beta_t = 1.99$, $\left\|(I_C + M_c)^{-1}\right\|_2$ reaches 146 (still bounded by $2/(2 - 1.99) = 200$), whereas trained 12.5B SLA layers stay well below 2 ($\bar{\beta} \in [0.251, 0.348]$). Second, our production kernel inverts the unscaled $I_C + M_c$ rather than the symmetric B: with $C = 256$ identical keys and constant $\beta_t = 1.9$, $\|M_c\|_2 = 309.05$ yet $\left\|(I_C + M_c)^{-1}\right\|_2 = \left\|B^{-1}\right\|_2 = 19.87 \leq 20$, whereas if $\beta_{1:128} = 1.9$ and $\beta_{129:256} = 10^{-8}$, unscaled $\left\|(I_C + M_c)^{-1}\right\|_2$ reaches $50.37 > 20$ via $D^{-1/2}B^{-1}D^{1/2}$ skew while symmetric $\left\|B^{-1}\right\|_2 = 19.54 \leq 20$ and $\left\|D(I_C + M_c)^{-1}\right\|_2 = 37.13 \leq 38.0$ obey the certificate. Scaling keys symmetrically by $\sqrt{\beta_t}$ before Op 4 inverts B directly on the MXU, holding 6-step block-doubling relative error at 5.2×10^{-6} in FP32 and 3.6% with BF16 operands even at $\|M_c\|_2 = 309$. Viewed iteratively, the intra-chunk delta rule executes one successive over-relaxation (Kaczmarz) sweep on H with relaxation $\beta_t \leq \beta_{\max} < 2$. □

SLA also offers a cheaper approximation, `1-Step Neumann`, which keeps only the first correction term between blocks; the ablation runs and prefill benchmarks use a matmul-free row-scaled variant of it. On a cluster of b aligned keys ($H = \mathbf{1}_b\mathbf{1}_b^\top$), one stationary Jacobi sweep $I_b - \beta H$ has spectral radius $|b\beta - 1| > 1$ whenever $b\beta > 2$. Row-sum Gershgorin damping stabilizes `A3/A5` at $+0.22$ pre-training perplexity (Section 3.5), whereas swapping `Exact Inv.` in at test time shifts 64K QA by at most 0.01 points (2.5×10^{-5} nats 14-layer tail KL). Pre-training and context-parallel runs therefore use `Exact Inv.`, except ablation arms `A3` and `A5`; prefill and tree-verification benchmarks use `1-Step Neumann`, and single-layer benchmarks report both.

Block-doubling and block-Neumann factorization. Let D_{blk} be the block-diagonal part of $I_C + M_c$ (its $C/16$ unit lower-triangular 16×16 diagonal blocks) and let $N = D_{\text{blk}}^{-1}(I_C + M_c - D_{\text{blk}})$ (local to Pillar 3), so that $I_C + M_c = D_{\text{blk}}(I_C + N)$ with N strictly block-lower-triangular and $N^{C/16} = 0$. Algebraically, at $C = 128$, the exact inverse and its 1-step block-Neumann truncation $A_{\text{N1}} \in \mathbb{R}^{C\times C}$ take the form

$$\begin{aligned} A_{\text{N1}} &= (I_C - N)\, D_{\text{blk}}^{-1}, \\ (I_C + M_c)^{-1} &= \sum_{s=0}^{7} (-N)^s D_{\text{blk}}^{-1} = (I_C + N^4)(I_C + N^2)(I_C - N)\, D_{\text{blk}}^{-1}. \end{aligned} \tag{10}$$

Rather than evaluating the expanded polynomial product $(I_C + N^4)(I_C + N^2)(I_C - N)D_{\text{blk}}^{-1}$, which amplifies powers of N in finite precision (diverging on 22%–23% of real Qwen3.5 $C = 128$ chunks; Appendix H), our Pallas kernel (`Exact Inv.`) builds on open-source MXU block-doubling [1] via the numerically stable **recursive Schur/block-doubling recurrence** in FP32 (`Precision.HIGHEST`):

$$X_{2s} = X_s - X_s L_s X_s \qquad (s = 2, 4, 8, 16, 32, 64),$$

where X_s is the block-diagonal matrix of already-inverted $s \times s$ diagonal blocks of $I_C + M_c$ ($X_{128} = A_c$) and L_s keeps only the strictly lower-left $s \times s$ coupling block inside each $2s \times 2s$ diagonal block (zero elsewhere), so $L_s X_s L_s = 0$ and the update is exact. Concretely, at each step $s \to 2s$, given the inverse of the two $s \times s$ diagonal halves ($X_{\text{top}}, X_{\text{bot}}$) and their strictly lower-left $s \times s$ coupling block L_{off}, the lower-left block of X_{2s} is computed via two MXU matmuls as $-X_{\text{bot}} L_{\text{off}} X_{\text{top}}$.

Variants of `1-Step Neumann` and truncation properties. Two different first-order approximations appear under the name `1-Step Neumann`, and we keep them distinct:

1. the *block* truncation $A_{\text{N1}} = (I_C - N)D_{\text{blk}}^{-1}$ with block width $B_{\text{blk}} = 16$, analyzed below and evaluated in our standalone lag/skew benchmark; and
2. the matmul-free *row-scaled* approximation $A_{\text{N1}}^{\text{row}} \in \mathbb{R}^{C\times C}$, with row-scaling vector $r = (r_1, \ldots, r_C)^\top \in \mathbb{R}^C$, used by the pre-training ablation arms `A3` and `A5` (Table 8b) and the prefill benchmarks:

$$A_{\text{N1}}^{\text{row}} = I_C - \operatorname{diag}(r)^{-1} M_c, \qquad r_i = \max\!\big(1,\, 1.25\, \|(M_c)_{i,:}\|_1\big).$$

The exactness property below (zero error for lags $t - j \leq 16$) holds for (1) but not for (2). Partition unit lower-triangular $I_C + M_c = D_{\text{blk}} + M_{\text{off}} = D_{\text{blk}}(I_C + N)$, where $D_{\text{blk}} = \operatorname{diag}(D_1, \ldots, D_m)$ consists

of $m = C/B_{\text{blk}}$ unit lower-triangular diagonal blocks $D_a \in \mathbb{R}^{B_{\text{blk}} \times B_{\text{blk}}}$ ($a \in \{1, \dots, m\}$) and M_{off} is zero on and above the block diagonal. Then $N = D_{\text{blk}}^{-1} M_{\text{off}}$ is strictly block-lower-triangular (its (a, b) block satisfies $N_{a,b} = 0$ for $a \le b$, hence $(N^s)_{a,b} = 0$ for all $a - b < s$ and $N^m = 0$). Comparing $(I_C + M_c)^{-1} = \sum_{s=0}^{m-1} (-N)^s D_{\text{blk}}^{-1}$ against $A_{\text{N1}} = (I_C - N) D_{\text{blk}}^{-1} = D_{\text{blk}}^{-1} - D_{\text{blk}}^{-1} M_{\text{off}} D_{\text{blk}}^{-1}$, the residual

$$(I_C + M_c)^{-1} - A_{\text{N1}} = \Big(\sum_{s=2}^{m-1} (-1)^s N^s \Big) D_{\text{blk}}^{-1}$$

vanishes identically on all block entries (a, b) with $a - b \le 1$, namely every diagonal block (a, a) and every first sub-diagonal block $(a + 1, a)$. Any two token indices $1 \le j \le t \le C$ with $t - j \le B_{\text{blk}}$ satisfy $\lceil t/B_{\text{blk}} \rceil - \lceil j/B_{\text{blk}} \rceil \in \{0, 1\}$ and therefore lie in block (a, a) or $(a + 1, a)$, incurring zero truncation error.

In finite-precision BF16/FP32 arithmetic, exact WY training vs. stateful step-by-step decoding achieves measured output KL divergence $< 10^{-6}$ per layer (3.4×10^{-5} across 4-layer BF16 stacks when chunk momentum syncs every C steps), while under `1-Step Neumann` ($B_{\text{blk}} = 16, C = 128$), truncation at lags $t - j > 16$ exhibits relative Frobenius error 1.6×10^{-5} in short-horizon `4L:1A` layers up to 4.9×10^{-2} in long-horizon `15L:1A` layers, yielding single-layer output KL $< 10^{-5}$ against exact recurrent decoding (and zero discrepancy if chunked prefill uses `1-Step Neumann`). Training with the row-scaled approximation (2) nonetheless changes the learned model: in the pre-training ablation (Table 8b, `A2`→`A3`) it raises 11-set perplexity by 0.22. □

Why exact inversion remains the default once values are packed. Once Pillar 1 packs $R = 4$ value heads per key head (2 solves per chunk), exact FP32 block-doubling adds only 4%–8% to full-LM pre-training step time over matmul-free 1-step approximations while maintaining the lowest mean perplexity and tightest seed spread (Appendix H). Replacing exact inversion zero-shot on exact-trained checkpoints with low-order global series fails because deep layers exhibit strongly correlated keys ($\|M_c\|_2 > 1$ in 74%–83% of Qwen3.5 chunks). Appendix H details exploratory approximate solvers for inference, including banded exact 16×16 blocks (matching exact zero-shot MMLU within 0.00–0.42 points at $N = 1{,}425$), a signed erase-aware 1-step operator (`eaadj`) exact for collinear keys under scalar decay, and $C = 256$ compilation with `kVectorMove`-safe exact block-doubling.

2.5 Single-Kernel VMEM Execution and Custom-VJP Backward Pass

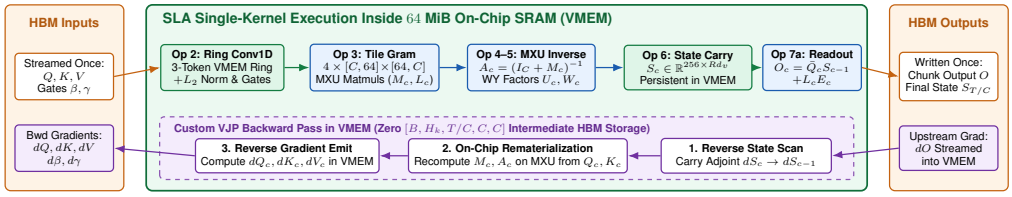

Figure 4: **Single-kernel VMEM execution and custom VJP backward pass in SLA. Top (Forward):** Streamed Q, K, V and gates enter the 64 MiB VMEM scratchpad once; a 3-token ring buffer executes causal `Conv1D` and L_2 normalization (Op 2), followed in-place by 4-tile MXU Gram contractions (Op 3), MXU chunk inversion and WY factors (Ops 4–5), persistent $[256, D_{v,\text{pack}}]$ state updates (Op 6), and chunk readout (Op 7a) before writing O to HBM. **Bottom (Backward):** During the reverse state scan, the custom VJP kernel recomputes intra-chunk matrices M_c, A_c on the MXU directly inside VMEM, eliminating $[B, H_k, T/C, C, C]$ intermediate HBM writes for batch size B.

Algorithm 2 details the custom-VJP backward pass of fused SLA (illustrated in Figure 4, bottom), where for a scalar training loss $\mathcal{L}$ we write $dM_c = \partial \mathcal{L}/\partial M_c$ (and analogously $dQ_c, dK_c, dV_c, \dots$) for the loss gradient (adjoint) with respect to each tensor. Let $\gamma_c \in (0, 1]^{C \times 4}$ stack chunk c's tile decays $\gamma_{t,g}$ (with tile column $\gamma_{c,g} \in (0, 1]^C$), and let $dO_c = \partial \mathcal{L}/\partial O_c \in \mathbb{R}^{C \times Rd_v}$ denote the upstream gradient of chunk c's readout after backpropagating through the per-value-head $\mathrm{RMSNorm}$ and output gate in Eq. (5). For batch size B, a standard autodiff tape stores the intra-chunk matrices $M_c, L_c, A_c \in \mathbb{R}^{C \times C}$ from the forward pass in HBM, consuming $3BH_k(T/C)C^2$ words. Because SLA's Pillar 2 and Pillar 3 evaluate both the $G = 4$ Gram matrices and the chunk inverse $A_c = (I_C + M_c)^{-1}$ as

MXU matmuls, our backward kernel instead saves only the chunk-boundary states $S_{c-1} \in \mathbb{R}^{d_k \times Rd_v}$ and recomputes M_c, L_c, A_c, and the unscaled writes $\tilde{E}_c = A_c(V_c - \bar{K}_c S_{c-1}) \in \mathbb{R}^{C \times Rd_v}$ on-the-fly inside VMEM during the reverse scan over $c = T/C, \ldots, 1$.

Whereas unfused baselines recompute or store separate per-value-head solves for the transposed system $(I_C + M_c)^{-\top}$ in the backward pass, SLA already holds its $R\times$-fewer dense inverses A_c in VMEM. Differentiating $E_c = \mathrm{diag}(\beta_c)A_c(V_c - \bar{K}_c S_{c-1}) = \mathrm{diag}(\beta_c)\tilde{E}_c$ through $A_c = (I_C + M_c)^{-1}$ reduces the inverse-solve backpropagation to two MXU matmuls for the scaled adjoint $D_c \in \mathbb{R}^{C \times Rd_v}$ and Gram adjoint $dM_c \in \mathbb{R}^{C \times C}$, shared across all R packed value heads:

$$D_c = A_c^\top \,\mathrm{diag}(\beta_c)\, dE_c \in \mathbb{R}^{C \times Rd_v}, \qquad dM_c = -\,\mathrm{tril}_{-1}\big(D_c \tilde{E}_c^\top\big) \in \mathbb{R}^{C \times C}, \tag{11}$$

where $\mathrm{tril}_{-1}(\cdot)$ and $\mathrm{tril}_0(\cdot)$ mask to strictly ($t > j$) and weakly ($t \geq j$) lower-triangular entries, $\mathrm{rowsum}(\cdot)$ and $\mathrm{colsum}(\cdot)$ sum matrix rows and columns, and $p^{\text{in}}_{c,g}, p^{\text{out}}_{c,g} \in (0,1]^C$ denote the tile-g boundary-decay vectors $(p^{\text{in}}_{c,g})_t = \prod_{s=1}^{t} \gamma_{s,g}$ and $(p^{\text{out}}_{c,g})_t = \prod_{s=t+1}^{C} \gamma_{s,g}$. From D_c, the value and decayed-key gradients follow immediately as $dV_c = D_c$ and $d\bar{K}_c = -D_c S_{c-1}^\top$.

Algorithm 2 Fused SLA Custom-VJP Backward Pass (Per Key Head in VMEM)

Require: Saved boundary states $\{S_{c-1}\}_{c=1}^{T/C}$, chunk inputs $\{Q_c, K_c, V_c, \beta_c, \gamma_c\}$, upstream readout gradients $\{dO_c\}_{c=1}^{T/C}$
Ensure: Input gradients $\{dQ_c, dK_c, dV_c, d\beta_c, d\gamma_c\}_{c=1}^{T/C}$

Initialize reverse carry $dS_{T/C} \leftarrow 0,\ dZ_{T/C} \leftarrow 0 \in \mathbb{R}^{d_k \times Rd_v}$ in VMEM
for $c = T/C$ **down to** 1 **do**
  // 1. On-chip VMEM recomputation (Ops 3–5 on MXU; zero $[C, C]$ HBM reads)
  Recompute $p^{\text{in}}_{c,g}, p^{\text{out}}_{c,g}$, decayed $(\bar{K}_c, \bar{Q}_c, \hat{K}_c)$, $G = 4$ masks (Γ_g, Ψ_g), and M_c, L_c via 64-wide MXU matmuls
  Recompute $A_c \leftarrow (I_C + M_c)^{-1}$ on MXU; form $\tilde{E}_c \leftarrow A_c(V_c - \bar{K}_c S_{c-1})$ and $E_c \leftarrow \mathrm{diag}(\beta_c)\tilde{E}_c$
  // 2. Reverse state and momentum scan adjoints (Ops 6–7a)
  $\widetilde{dZ}_c \leftarrow dZ_c + dS_c; \quad dZ_{c-1} \leftarrow \mu\kappa_c \widetilde{dZ}_c$ ▷ velocity carry ($\kappa_c = 0$ at doc reset)
  $dE_c \leftarrow \hat{K}_c \widetilde{dZ}_c + L_c^\top dO_c \in \mathbb{R}^{C \times Rd_v}; \quad d\hat{K}_c \leftarrow E_c \widetilde{dZ}_c^\top \in \mathbb{R}^{C \times d_k}$ ▷ MXU
  $dL_c \leftarrow \mathrm{tril}_0(dO_c E_c^\top) \in \mathbb{R}^{C \times C}; \quad d\bar{Q}_c \leftarrow dO_c S_{c-1}^\top \in \mathbb{R}^{C \times d_k}$ ▷ MXU
  // 3. Transposed WY inverse and value/key adjoints on MXU (Ops 4–5)
  $D_c \leftarrow A_c^\top \,\mathrm{diag}(\beta_c)\, dE_c \in \mathbb{R}^{C \times Rd_v}; \quad dV_c \leftarrow D_c; \quad d\bar{K}_c \leftarrow -D_c S_{c-1}^\top$ ▷ MXU
  $dM_c \leftarrow -\,\mathrm{tril}_{-1}(D_c \tilde{E}_c^\top) \in \mathbb{R}^{C \times C}$ ▷ contracts Rd_v packed value cols on MXU
  $dS_{c-1} \leftarrow F_c^\top dS_c - \bar{K}_c^\top D_c + \bar{Q}_c^\top dO_c \in \mathbb{R}^{d_k \times Rd_v}$ ▷ prev. chunk state carry ($\bar{K}_c^\top D_c = W_c^\top dE_c$)
  // 4. Tile-grouped Q, K and gate adjoints ($G = 4$ tiles of $b = 64$ channels)
  for $g \in \{1, 2, 3, 4\}$ **do**
    $\widetilde{dM}_{c,g} \leftarrow dM_c \odot \Gamma_g; \quad \widetilde{dL}_{c,g} \leftarrow dL_c \odot \Psi_g$
    $dQ_{c,g} \leftarrow \widetilde{dL}_{c,g} K_{c,g} + \mathrm{diag}(p^{\text{in}}_{c,g})(d\bar{Q}_c)_{:,g}$ ▷ 64-wide MXU matmul; $(\cdot)_{:,g}$ is tile-g slice
    $dK_{c,g} \leftarrow (\widetilde{dM}_{c,g} + \widetilde{dM}_{c,g}^\top) K_{c,g} + \widetilde{dL}_{c,g}^\top Q_{c,g} + \mathrm{diag}(p^{\text{in}}_{c,g})(d\bar{K}_c)_{:,g} + \mathrm{diag}(p^{\text{out}}_{c,g})(d\hat{K}_c)_{:,g}$
  end for ▷ concatenate $dQ_c = [dQ_{c,1}, \ldots, dQ_{c,4}],\ dK_c = [dK_{c,1}, \ldots, dK_{c,4}]$
  $d\beta_c \leftarrow \mathrm{rowsum}(dE_c \odot \tilde{E}_c) + \mathrm{colsum}\big(dM_c \odot \sum_{g=1}^{4} \Psi_g \odot (K_{c,g} K_{c,g}^\top)\big)^\top$
  Accumulate columns $d\gamma_{c,g}$ of $d\gamma_c$ via reverse prefix sum over $dM_c, dL_c, d\bar{K}_c, d\bar{Q}_c, d\hat{K}_c$, and $dS_c S_{c-1}^\top$
end for
return $\{dQ_c, dK_c, dV_c, d\beta_c, d\gamma_c\}_{c=1}^{T/C}$

3 Empirical Results

We evaluate SLA on TPU v7x (Ironwood [6, 22]) at four levels: single-layer kernels and multi-generation rooflines (Section 3.1), long-context serving (Section 3.2), context-parallel training and end-to-end 12.5B MoE cube serving (Section 3.3), and multi-scale MoE pre-training throughput and quality (Sections 3.4–3.5). We use two configurations with $d_k = 256$, $G = 4$ decay tiles, chunk size $C = 128$, and $H_v = 8$ value heads: `SLA-R2` shares each query-key head across 2 value heads ($H_q = H_k = 4$, where H_q is the query-head count) and `SLA-R4` across 4 ($H_q = H_k = 2$); layer shapes are written $H_k \times d_k \times Rd_v$. `xL:1A` denotes x linear attention layers per full softmax attention layer, so `0L:1A` is a pure softmax attention Transformer. SLA keeps the baseline's value width, $d_v = 256$ at model width $D = 2048$ and 512 at $D = 1792$, so SLA's BF16 state has twice the entries of FP32 KDA with $d_k = 128$ in the same memory, and the same entries as the 63.1B KDA with $d_k = 256$ in half the memory. `SLA Iso-State` instead halves d_v to match KDA's element count in half the bytes. Precision explains only a small part of the speedup: switching MaxText KDA's state

from FP32 to BF16 makes it just 2%–8% faster (Table 5, footnote). `SLA Tree-CP` (Table 6c) and all SLA pre-training runs use `Exact Inv.`, except ablation arms `A3` and `A5` (Table 8b); 32-layer prefill and tree verification (Table 6a–b) use `1-Step Neumann`, and single-layer benchmarks report both (Tables 4–5).

3.1 Kernel Performance and Multi-Generation TPU Rooflines

Table 4: **Seven-operator forward audit (ms) on TPU v7x** ($D = 2048$, $T = 8{,}192$). Unfused columns report isolated per-stage forward medians; fused brackets report single-kernel Ops 2–7a latency.

Stage Latency ($D = 2048, T = 8\text{K}$)	KDA (Unf.)	KDA (Fus.)	SLA-R4 (Unf.)	**SLA-R4 (Fused)**
Op 1: QKV + Gate Proj (MXU)	0.251	0.251	0.168	0.168
Op 2: Causal Conv4 + QK Norm	0.562		0.341	
Op 3–5: Gram + WY Inversion	1.184	Fused	0.418	Fused
Op 6: Chunk State Scan (S_c)	1.248	3.104	0.614	1.112
Op 7a: Intra + State Readout	0.219		0.219	
Op 7b–7c: Out Norm + Out Proj	0.255	0.255	0.245	0.245
Measured Full Layer Fwd (ms)	3.719	3.610	2.005	**1.525**
Analytical Roofline Fwd (ms)	2.890	2.740	1.380	**0.867**
Measured Full Layer F+B (ms)	7.108	6.912	3.982	**2.992**
F+B Speedup vs. KDA (Unf.)	1.00×	1.03×	1.79×	2.38×

Stateless functional harness (`1-Step Neumann` for `SLA-R4`). Fused: one kernel for Ops 2–7a; full NNX module latencies in Table 5.

Table 5: **Full-layer TPU v7x latency (ms, Ops 1–7c, $B = 1$), state bytes, 32K temporary HBM, and F+B speedups** vs. `MaxText Pallas KDA` ($C = 128$, used in pre-training and serving) and `Tokamax Mosaic KDA` ($C = 64$). Core-kernel and iso-shape full-layer attributions are in Appendix Tables 16–17.

Geometry	Full Layer ($H_k \times d_k \times Rd_v$, Chunk C)	State	8K F+B	vs. Max.	vs. Tok.	32K F+B	vs. Max.	vs. Tok.	32K Temp HBM
(a) $D = 1792$ (63.1B MoE width)	MaxText Pallas KDA ($8 \times 256 \times 512, C = 128$, FP32)*	4.0 MiB	14.46	1.00×	0.91×	48.66	1.00×	0.87×	9,302 MiB
	Tokamax Mosaic KDA ($8 \times 256 \times 512, C = 64$, FP32)	4.0 MiB	13.22	1.09×	1.00×	42.38	1.15×	1.00×	9,483 MiB
	GDN Scalar Unpacked ($G = 1, 8 \times 256 \times 512, C = 128$)	4.0 MiB	13.78	1.05×	0.96×	48.39	1.01×	0.88×	9,217 MiB
	SLA-R2 Unfused ($G = 1, 4 \times 256 \times 1024, C = 128$, BF16)	2.0 MiB	10.88	1.33×	1.21×	32.46	1.50×	1.31×	6,922 MiB
	SLA-R2 Unfused ($G = 4, 4 \times 256 \times 1024, C = 128$, BF16)	2.0 MiB	10.56	1.37×	1.25×	32.56	1.49×	1.30×	6,922 MiB
	SLA-R4 Unfused ($G = 4, 2 \times 256 \times 2048, C = 128$, BF16)	2.0 MiB	7.71	1.88×	1.72×	26.24	1.85×	1.62×	5,782 MiB
	SLA-R4 2-Stage, Exact Inv. ($2 \times 256 \times 2048$, BF16)	2.0 **MiB**	**5.57**	2.60×	2.37×	**20.24**	2.40×	2.09×	3,738 **MiB**
	SLA-R4 2-Stage, 1-Step Neumann ($2 \times 256 \times 2048$, BF16)	2.0 **MiB**	**5.25**	2.75×	2.52×	**18.79**	2.59×	2.26×	3,738 **MiB**
	1-Chip Pre-Alloc. Ref.: MaxText KDA / Tokamax KDA†	4.0 MiB	10.42 / 8.70	1.00×†	1.00×†	39.81 / 34.01	1.00×†	1.00×†	9,302 MiB
	SLA-R4 1-Kernel, Scalar $G = 1$ **Control** ($2 \times 256 \times 2048$, BF16)†	2.0 **MiB**	**3.73**	2.79×†	2.33×†	**14.05**	2.83×†	2.42×†	3,738 **MiB**
	SLA-R4 1-Kernel, Exact Inv., $G = 4$ ($2 \times 256 \times 2048$, BF16)†	2.0 **MiB**	**3.89**	2.68×†	2.24×†	**14.41**	2.76×†	2.36×†	3,738 **MiB**
(b) $D = 2048$ (12.5B linear value width)	MaxText Pallas KDA ($8 \times 128 \times 256, C = 128$, FP32)*	1.0 MiB	9.53	1.00×	0.65×	32.45	1.00×	0.74×	4,432 MiB
	Tokamax Mosaic KDA ($8 \times 128 \times 256, C = 64$, FP32)	1.0 MiB	6.16	1.55×	1.00×	24.06	1.35×	1.00×	4,727 MiB
	GDN Scalar Unpacked ($G = 1, 8 \times 128 \times 256, C = 128$)	1.0 MiB	10.24	0.93×	0.60×	31.86	1.02×	0.76×	4,360 MiB
	SLA-R2 Unfused ($G = 1, 4 \times 256 \times 512, C = 128$, BF16)	1.0 MiB	7.25	1.31×	0.85×	25.28	1.28×	0.95×	4,260 MiB
	SLA-R4 Unfused ($G = 4, 2 \times 256 \times 1024, C = 128$, BF16)	1.0 MiB	5.22	1.82×	1.18×	19.68	1.65×	1.22×	3,472 MiB
	SLA-R2 2-Stage, 1-Step ($4 \times 256 \times 512, C = 128$, BF16)	1.0 MiB	4.77	2.00×	1.29×	17.29	1.88×	1.39×	3,092 MiB
	SLA-R4 2-Stage, Exact Inv. ($2 \times 256 \times 1024$, BF16)	1.0 **MiB**	**4.54**	2.10×	1.36×	**15.72**	2.06×	1.53×	2,263 **MiB**
	SLA-R4 2-Stage, 1-Step Neumann ($2 \times 256 \times 1024$, BF16)	1.0 **MiB**	**3.99**	2.39×	1.55×	**14.82**	2.19×	1.62×	2,263 **MiB**
	1-Chip Pre-Alloc. Ref.: MaxText KDA / Tokamax KDA†	1.0 MiB	6.87 / 4.26	1.00×†	1.00×†	26.69 / 16.28	1.00×†	1.00×†	4,432 MiB
	SLA-R4 1-Kernel, Scalar $G = 1$ **Control** ($2 \times 256 \times 1024$, BF16)†	1.0 **MiB**	**2.63**	2.61×†	1.62×†	**9.20**	2.90×†	1.77×†	2,263 **MiB**
	SLA-R4 1-Kernel, Exact Inv., $G = 4$ ($2 \times 256 \times 1024$, BF16)†	1.0 **MiB**	**2.74**	2.51×†	1.55×†	**9.59**	2.78×†	1.70×†	2,263 **MiB**

*Switching MaxText KDA's state from FP32 to BF16 changes 8K/32K F+B only slightly to $13.91/46.40$ ms ($D = 1792$) and $8.80/31.72$ ms ($D = 2048$). Tokamax Mosaic KDA receives BF16 QKV inputs (its MXU Neumann–Horner path); the FP32 label denotes its FP32 recurrent state (Appendix H audits Tokamax BF16 vs. FP32 QKV inputs directly). Row-scaled FP8 halves SLA state bytes further (1.0 MiB and 0.5 MiB). Separate 8K Fwd/Bwd splits for `SLA-R4 2-Stage Exact`: $2.88/2.69$ ms ($D = 1792$) and $2.44/2.10$ ms ($D = 2048$) vs. MaxText $6.41/8.04$ and $4.74/4.79$ ms (Tokamax: $5.90/7.33$ and $3.57/2.59$ ms). Undaggered rows measured on container `maxtext_e174258_20261008` (Tokamax 0.0.14) with $B = 1$/core concurrently across a 64-chip (128-core) slice and `-xla_tpu_scoped_vmem_limit_kib=65472` (64 MiB VMEM budget, required by fused `SLA-R4` at $D = 1792$ with $\mu = 0.25$; XLA's default 16 MiB cap suffices at $\mu = 0$ and $D = 2048$). †Daggered 1-kernel rows and their matched single-chip pre-allocated baseline references are measured under the unified single-chip protocol (3 independent process restarts, 64 MiB VMEM, Appendix Table 20a), so every daggered speedup ratio divides directly against the italicized reference row immediately above it (e.g. $10.42/3.89 = 2.68\times$ and $8.70/3.89 = 2.24\times$; iso-shape controls in Appendix Table 17; GDN-2 in Section 3.5).

Fusion and stage breakdown. Because every intra-chunk stage of SLA runs on the MXU, the fused kernel keeps Ops 2–7a in the 64 MiB VMEM scratchpad and recomputes M_c and A_c in the

backward pass instead of storing them in HBM (Section 2.5 and Figure 4). Table 4 shows where the time goes. Fusing GPU-shaped KDA gains only $1.03\times$. Reshaping the layer without fusion (unfused `SLA-R4`) gives $1.79\times$, primarily because 4:1 key sharing cuts projection, convolution, Gram, and solve work while packing full MXU tiles. Fusion then brings the total to $2.38\times$ (with the 1-step Neumann inverse; Table 5 reports both inverses).

Comparison with open-source TPU baselines. Table 5 compares full SLA layers with MaxText Pallas KDA [1] ($C = 128$, used in our pre-training and serving runs), Tokamax Mosaic KDA [23] (the fastest baseline; its per-channel KDA kernel is validated at $C = 64$, running on its BF16-input MXU Neumann–Horner path), and unpacked scalar GDN [74]. With the exact inverse, 2-stage `SLA-R4` is $2.1\times$–$2.6\times$ faster than open-source MaxText KDA and $1.4\times$–$2.4\times$ faster than open-source Tokamax KDA across the 64-chip slice (and 1-kernel fused `SLA-R4` reaches $2.51\times$–$2.78\times$ vs. MaxText and $1.55\times$–$2.36\times$ vs. Tokamax under the unified single-chip protocol; $1.9\times$–$2.5\times$ vs. BF16-state MaxText KDA), while using 49%–61% less temporary HBM. Note that at $D = 1792$ the KDA baseline already uses $d_k = 256$, so key-tile padding elimination applies to $D = 2048$ ($d_k = 128 \to 256$), whereas at $D = 1792$ SLA's speedup comes from 4:1 head and write-gate sharing, tile-grouped decay, and single-kernel VMEM fusion. Specifically, at $D = 1792, 32\text{K}$, the 20.24 ms (20.237 ms) latency in Table 5a corresponds to the 2-stage chunk kernel with $\mu = 0.25$ (20.30 ± 0.23 ms on an isolated chip in Appendix Table 17), which is already $1.69\times$ faster than single-chip Tokamax Mosaic KDA (34.31 ± 0.13 ms) and $2.09\times$ faster under Table 5's 64-chip concurrent timing (42.38 ms), whereas the 1-kernel fused variant (14.08 ± 0.30 ms, $\mu = 0$) is $2.44\times$ faster than Tokamax Mosaic KDA at $D = 1792$. Core-kernel controls (Appendix Table 16b) show that beating Tokamax requires 4:1 head sharing: at $R = 2$ without sharing four value heads, $d_k = 256$ SLA cores run at $0.67\times$–$0.77\times$ of Tokamax. Conversely, Appendix Table 17 holds the full layer fixed at the `SLA-R4` shape ($H_k = 2, d_k = 256, R = 4$, $n = 3$ restarts) to separate each design choice: 4:1 key sharing accounts for $1.71\times$–$1.87\times$ on the MaxText KDA kernel and single-kernel VMEM fusion adds $1.21\times$–$1.41\times$, whereas BF16 vs. FP32 state adds $1.00\times$–$1.01\times$ (providing $2\times$ memory and VMEM headroom instead). Grouping decay from $G = d_k$ to $G = 4$ speeds up pure XLA by $1.62\times$–$2.45\times$ and Pallas chunk kernels by $1.04\times$–$1.12\times$ while running at the same single-kernel VMEM speed (0%–9%) as a scalar gate ($G = 1$, 14.08 vs. 14.16 ms at $D = 1792, 32\text{K}$), because 1 or 4 gate vectors both fit in VMEM vector registers without $[B, H_k, T/C, C, C, d_k]$ HBM materialization; however, $G = 1$ restricts each head to one retention timescale and trails $G = 4$ on 11-set validation perplexity across all six paired seeds (-0.23 ± 0.24 PPL, $6/6$ wins; Section 3.5), making $G = 4$ the multi-timescale quality and single-kernel VMEM speed Pareto knee. Packing $R = 4$ values into Tokamax alone gives only $1.01\times$–$1.10\times$ at $D = 2048$ and overflows 64 MiB VMEM in backward at $D = 1792$; fused `SLA-R4` remains $1.34\times$–$1.60\times$ faster than iso-shape Tokamax and $1.27\times$–$1.58\times$ faster than the fastest iso-shape KDA Pallas kernel. To make the attribution against open-source Tokamax exact, we also implemented and verified a single-kernel VMEM-fused per-channel ($G = d_k$) KDA kernel at both $R = 1$ and $R = 4$ (forward cosine 0.999995 and all five VJP gradients ≥ 0.999984 vs. FP32 reference; Appendix I, Table 35). Across $D \in \{1792, 2048\}$ and $T \in \{8\text{K}, 32\text{K}\}$ on single-chip TPU v7x, fused `SLA-R4` ($\mu = 0$) beats Tokamax $R = 1$ by $1.52\times$–$1.84\times$, factoring cleanly into $1.28\times$–$1.41\times$ from 4:1 head packing, $1.10\times$–$1.12\times$ from VMEM fusion at matched $R = 4$, and $1.07\times$–$1.20\times$ from tile grouping $G = d_k \to G = 4$ ($1.20\times$ at $D = 2048, 32\text{K}$, while $G = 1$ stays within 2.7%–4.0% of $G = 4$). Crucially, VMEM-fusing unpacked $R = 1$ per-channel KDA runs *slower* than Tokamax $R = 1$ ($0.44\times$–$0.69\times$), confirming that 4:1 value packing is what makes single-kernel VMEM fusion profitable.

Multi-generation TPU rooflines and FP8 skew. On an identical portable JAX/Pallas kernel stack ($B = 1, T = 8{,}192, D = 2048$), `SLA-R4`'s single-core forward speedup rises from $4.5\times$ on TPU v4 to $6.7\times$ on TPU v7x (Figure 5 and Appendix Table 18). Across these generations the MXU tile widens from 128 to 256 and the MXU:VPU throughput ratio α_{mv} rises from about $85\times$ to $128\times$. Because Tokamax Mosaic validates per-channel KDA on v6e/v7x at $C = 64$, this cross-generation sweep holds the software stack fixed against a portable FP32 KDA kernel ($C = 64$; 6.99 ms on v7x), whereas Appendix Table 16b compares core kernels against optimized MaxText (2.67 ms fwd) and Tokamax (1.07 ms fwd; `SLA-R4`: 1.05 ms) on v7x.

To understand this scaling on a single TensorCore with systolic tile width $W \in \{128, 256\}$ and single-core peaks $\mathcal{P}_{\text{MXU}}, \mathcal{P}_{\text{VPU}}$ (Table 2 divided by `Cores`), we model single-core forward latency

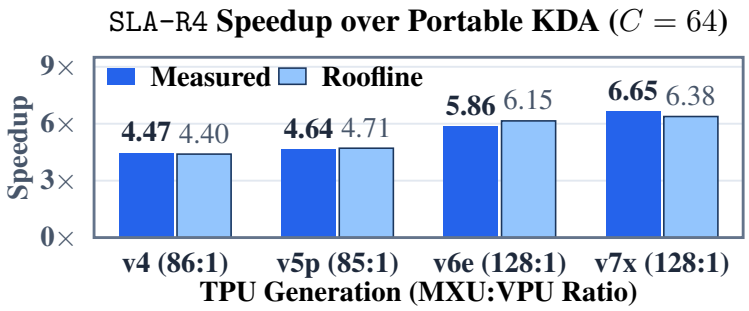

Figure 5: **Forward speedup of `SLA-R4` across TPU generations** over a portable FP32 KDA baseline, with leave-one-out roofline predictions (3.1% error). Not comparable to the MaxText/Tokamax-normalized speedups in Table 5.

$T_{\text{fwd}}(C)$ at chunk size C as the sum of inter-chunk scan, VPU solve, and intra-chunk MXU terms:

$$T_{\text{fwd}}(C) \approx H_k \frac{T}{C}\left(\tau_{\text{pipe}} + \frac{F_{\text{step}}}{\eta_{\text{sys}}\mathcal{P}_{\text{MXU}}}\right) + \frac{\gamma_{\text{inv}} H_k T C}{\mathcal{P}_{\text{VPU}}} + \frac{\Phi_{\text{MXU}}(C)}{\mathcal{P}_{\text{MXU}}}. \tag{12}$$

In the scan term, $(\tilde{d}_k, \tilde{D}_{v,\text{pack}})$ pads $(d_k, D_{v,\text{pack}})$ to multiples of W, $F_{\text{step}} = 2\max(C, W)\,\tilde{d}_k \tilde{D}_{v,\text{pack}}$ is the FLOP count per scan matmul at effective single-core utilization $\eta_{\text{sys}} \approx 8.6\%$, and $\tau_{\text{pipe}} \approx 1.5$–$2.5\,\mu$s is the in-VMEM Pallas pipeline fill/drain cost per head and chunk (below the top-level XLA `lax.scan` HBM-barrier cost $\tau_{\text{sync}} \approx 12\,\mu$s in Figure 2). In the remaining terms, γ_{inv} counts VPU operations per triangular-solve entry (unrelated to the decay rates $\gamma_{t,g}$) and $\Phi_{\text{MXU}}(C) = \mathcal{O}(H_k T C d_k + H_v T C d_v)$ counts intra-chunk matmul FLOPs. The four generations in Table 18 form two clusters. On 128×128-tile TPUs (v4 and v5p; no KDA key-tile padding, $\alpha_{\text{mv}} \approx 85\times$), `SLA-R4` is already $4.47\times$–$4.64\times$ faster. On 256×256-tile TPUs (v6e and v7x; 50% KDA key-tile padding at $d_k = 128$, $\alpha_{\text{mv}} \approx 128\times$), the speedup rises by a further $1.3\times$–$1.46\times$ to $5.86\times$–$6.65\times$. Tile width and α_{mv} change together between the clusters, so these measurements cannot separate their effects (a reduced-form leave-one-out fit is shown for reference in Table 18). On an isolated FP8 MXU microbenchmark, `SLA-R4` is $1.69\times$ faster than the standalone JAX/Pallas $C = 128$ KDA core kernel in native E4M3 FP8 mode, where α_{mv} doubles to $256\times$ (1.280 vs. 2.168 ms). This matches its $1.67\times$ BF16 speedup on the same harness (1.181 vs. 1.975 ms), which excludes MaxText's segment-reset and L_2-norm wrapper (included in the 2.668 ms row of Appendix Table 16b). To test out-of-sample predictiveness beyond these four generations, we locked a priori analytical roofline predictions for 4 anchor shapes and 18 strictly held-out single-chip TPU v7x designs before measuring them (0 failures; Appendix I, Table 34). Zero-parameter peak rooflines achieve Spearman rank correlation $\rho = 0.893$ on full-layer F+B (0.874 full forward); a single anchor-calibrated efficiency factor ($\eta_{\text{cal}} = 0.339$, anchor $R^2 = 0.987$) yields 19.8% MAPE (18.6% MedAPE) across all 18 held-out shapes and 13.9% MAPE (12.5% MedAPE) inside the target packed systolic regime ($R \geq 4, d_k \geq 256, n = 12$). On the held-out ($R = 8, d_k = 512, C = 256, d_v = 256, T = 8\text{K}$) design, the model predicts 2.56 ms vs. 2.47 ms measured (+3.5% error), whereas unpacked $R = 1$ shapes run 42%–53% slower than the packed efficiency extrapolation due to unmodeled VPU and scheduling overhead.

3.2 Long-Context Serving

Prefill latency and prefix caching. In 32-layer hybrid stacks with SwiGLU FFNs [56] (Table 6a), `15L:1A SLA-R4` (one softmax layer per 16) reaches a 131K time-to-first-token (TTFT) of 1.67 s, compared with 2.66 s for the same hybrid layout with KDA and 7.14 s for full grouped-query attention (GQA) [2]. Most of the gain over GQA comes from the hybrid layout itself. In the same `15L:1A` layout, `SLA-R4` with FP8 state and the `1-Step Neumann` inverse is $1.6\times$ faster than FP32 KDA; for zero-shot deployment on exact-trained weights, `Exact Inv.` (BF16 131K TTFT 1,623 ms) or banded exact blocks avoid approximation drift without retraining (Appendix H). A cached SLA state is also much smaller than a KV cache (0.5–1 MiB per layer versus 64 MiB at 32K), which makes multi-turn prefix caching cheap.

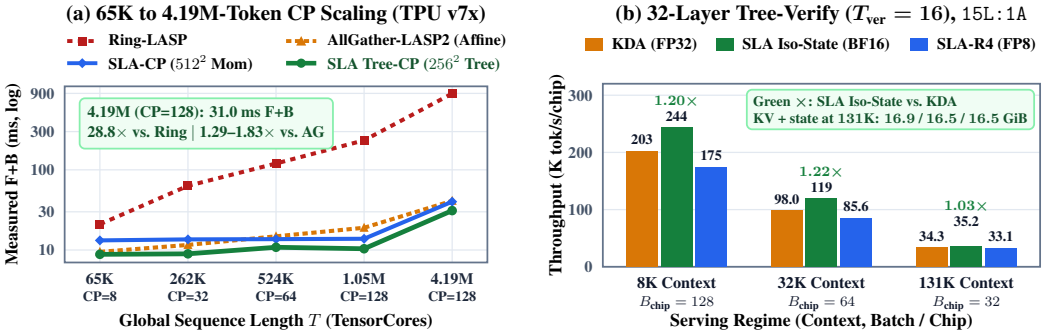

Figure 6: **Context parallelism and serving on TPU v7x. (a)** Forward-backward scan latency across 8–128 TensorCores (65K–4.19M tokens, Section 3.3); the callout highlights the no-momentum `SLA Tree-CP` curve ($\mu = 0$), while trained `SLA-CP` ($\mu = 0.25$) takes 40.1 ms at 4.19M ($0.85\times$–$1.37\times$ vs. AllGather across 262K–4.19M; Table 6c). **(b)** Single-chip 32-layer attention-layer tree-verification throughput for the `15L:1A` hybrid at $T_{\text{ver}} = 16$ (Table 6b). At equal state element count (BF16 vs. FP32), `SLA Iso-State` is $1.03\times$–$1.22\times$ faster than KDA; `SLA-R4` with FP8 state reaches $0.86\times$–$0.97\times$ of KDA's throughput.

Table 6: **Single-chip 32-layer prefill latency (a)**, 32-layer attention-layer serving throughput with speculative verification of $T_{\text{ver}} \in \{8, 16\}$ drafted tokens (b; `OOM` $>$ 190 GiB/chip), and context-parallel F+B scan latency across $P \in \{32, 64, 128\}$ TensorCores (c) on TPU v7x ($D = 2048$; 12.5B MoE cube serving in Appendix Table 21).

Regime	Architecture / Method (32L, $D = 2048$)	8K TTFT	32K TTFT	131K TTFT (vs. KDA)	32K Cache/Lyr
(a) 32L Prefill Full TTFT ($B = 1$)	`0L:1A` Dense GQA ($H_q = 16$, $H_{kv} = 4$ KV heads)	214.7 ms	761.1 ms	7,139 ms ($0.37\times$)	64.0 MiB KV
	`15L:1A` MaxText KDA ($8 \times 128 \times 256$, FP32)	265.7 ms	689.1 ms	2,664 ms ($1.00\times$)	1.0 MiB FP32
	`7L:1A` SLA-R4 ($2 \times 256 \times 1024$, BF16, 1-Step)	204.3 ms	439.8 ms	2,034 ms ($1.31\times$)	1.0 MiB BF16
	`15L:1A` **SLA-R4** ($2 \times 256 \times 1024$, FP8, 1-Step)	203.6 **ms**	416.9 **ms**	1,669 **ms** ($1.60\times$)	0.5 **MiB FP8**
	`15L:1A` **SLA-R4** ($2 \times 256 \times 1024$, BF16, **Exact Inv.**)	–	–	1,623 **ms** ($1.64\times$)	1.0 MiB BF16
Regime	Attention-Layer Serving (`tok/s/chip`)†	8K, $B=128$	32K, $B=64$	131K, $B=32$	KV+State (131K)
(b) Serving Tree Verify ($T_{\text{ver}} \in \{8, 16\}$)	`0L:1A` Dense GQA ($T_{\text{ver}} = 8$)	25.0K	OOM	OOM	256.0 GiB
	`7L:1A` MaxText KDA (FP32, $T_{\text{ver}} = 8$)	117.7K	42.5K	12.9K	32.88 GiB
	`7L:1A` SLA-R4 ($2\times$ El., $T_{\text{ver}} = 8$)	84.3K	38.4K	12.1K	32.88 GiB
	`7L:1A` MaxText KDA (FP32, $T_{\text{ver}} = 16$)	152.9K	64.3K	OOM	32.88 GiB
	`7L:1A` SLA Iso-State (BF16, $T_{\text{ver}} = 16$)	162.1K ($1.06\times$)	68.6K ($1.07\times$)	OOM	32.44 GiB
	`15L:1A` MaxText KDA (FP32, $T_{\text{ver}} = 16$)	202.5K	98.0K	34.3K	16.94 GiB
	`15L:1A` SLA Iso-State (BF16, $T_{\text{ver}} = 16$)	243.8**K** ($1.20\times$)	119.3**K** ($1.22\times$)	35.2**K** ($1.03\times$)	16.47 **GiB**
	`15L:1A` SLA-R4 FP8 (E4M3, $T_{\text{ver}} = 16$)	174.7K	85.6K	33.1K	16.47 **GiB**
Regime	Context-Parallel Scan (F+B ms)	$P=32$ (262K)	$P=64$ (524K)	$P=128$ (1.05M)	$P=128$ (4.19M)
(c) Context Parallelism (32–128 Cores)	Ring-LASP Sequential Ring ($P-1$ Hops)	63.20	120.12	233.47	894.41
	AllGather-LASP2 Collective Baseline	11.60	14.84	19.00	41.06
	SLA Tree-CP ($\mu = 0.25$, trained, 512^2)	13.59	13.71	13.84	40.08
	SLA Tree-CP ($\mu = 0$, no mom., 256^2)	8.97	10.84	10.40	31.03
	Tree-CP vs. AllGather ($\mu = 0.25$ / $\mu = 0$)	$0.85\times$ / $1.29\times$	$1.08\times$ / $1.37\times$	$1.37\times$ / $1.83\times$	$1.02\times$ / $1.32\times$

†Panel (b) reports attention layers only (2×95 GiB/core HBM cap); with SwiGLU FFNs included at $T_{\text{ver}} = 16$, `SLA Iso-State` is 2%–7% faster end-to-end than matched-cadence KDA. Panels (a)–(b) use `1-Step Neumann` (at 131K BF16 prefill, `Exact Inv.` reaches 1,623 ms [$4.40\times$ vs. Dense] vs. 1,477 ms for `1-Step Neumann` with 0.99998 14-layer cosine and 2.5×10^{-5} nats tail KL); panel (c) uses exact inversion (`Exact Inv.`). On exact-trained checkpoints without retraining, zero-shot prefill should use `Exact Inv.` or banded exact blocks (Appendix H), whereas matmul-free `1-Step Neumann` requires training with it or short KL healing.

Speculative verification. Single-token decoding underuses the MXU, so we verify $T_{\text{ver}} \in \{8, 16\}$ drafted tokens per step [39] with the same block inversion (Figure 6b and Table 6b). At equal state element count (`SLA Iso-State`, BF16 vs. FP32), SLA is $1.03\times$–$1.22\times$ faster than KDA on the attention layers, and 2%–7% faster once FFNs are included. `SLA-R4`, which stores twice KDA's state elements in the same bytes, is slower than KDA at every context length ($0.72\times$–$0.94\times$). Table 21c reports a 12.5B MoE speculative-decoding setting in which SLA's greedy acceptance length on out-of-distribution 32K/64K `HotpotQA` is $2.2\times$ that of unclamped KDA and $1.17\times$–$1.33\times$ that of clamped KDA ($1.06\times$–$1.10\times$ under stochastic sampling); KDA is slightly higher in-domain at 8K.

3.3 Context-Parallel Training and End-to-End 12.5B MoE Cube Serving

Context parallelism with a prefix-tree scan. Head-sharded Ulysses [33] needs at least as many key heads as devices, which SLA's $H_k \in \{2, 4\}$ do not provide, and ring-based LASP [62] passes the state through $P-1$ sequential hops across P devices. For any gated delta rule with scalar β_t, a chunk's state update is an affine map $S \mapsto JS + B$ with a $d_k \times d_k$ transition $J \in \mathbb{R}^{d_k \times d_k}$ and offset $B \in \mathbb{R}^{d_k \times D_{v,\text{pack}}}$, so P devices can reduce their local chunks independently and combine them in $\lceil \log_2 P \rceil$ prefix-tree steps [13, 43]. This construction is not specific to SLA, and our KDA baseline uses the same tree. SLA's advantage here is that one 256×256 transition and $G = 4$ group scalars serve all R value heads of a key head, so fewer transitions are exchanged; chunk momentum ($\mu = 0.25$) doubles the carried state and enlarges the transition to 512×512. Figure 6a and Table 6c report both settings across 8–128 TensorCores: on the isolated context-parallel scan across 64 chips, `SLA Tree-CP` achieves both a $1.09\times$ speedup over `KDA-CP` (22.58 ms vs. 24.64 ms) and a $7.65\times$ speedup over dense `RingCP` (22.58 ms vs. 172.71 ms). With momentum, as trained, the tree scan is $0.85\times$–$1.37\times$ vs. an AllGather-based LASP-2 baseline [63], and without momentum it is $1.29\times$–$1.83\times$ faster. In end-to-end 12.5B MoE training at 1M context ($L = 1{,}048{,}576$ on 64 TPU v7x chips, Table 21a), `15L:1A SLA Tree-CP` runs in 4,870.6 ms with momentum (4,778.1 ms without momentum; 4,737.6 ms in the uncoupled sweep), achieving a $1.09\times$ speedup ($1.11\times$–$1.13\times$ without momentum) over `15L:1A KDA-CP` (5,307.1–5,337.6 ms) with 6.31 GiB/core lower peak HBM (25.50 vs. 31.81 GiB/core) and a $7.50\times$–$7.75\times$ ($7.65\times$ without momentum) speedup over `0L:1A` dense `RingCP` (36,530.6–36,725.0 ms).

Affine monoid derivation for uncoupled and coupled `SLA Tree-CP`. Shard $p \in \{1, \dots, P\}$ holds L_{loc} consecutive chunks; quantities of its c-th local chunk ($c \in \{1, \dots, L_{\text{loc}}\}$) carry the subscript (p, c), with shard input state $S_{p,0}$ and shard output state $S_{p,L_{\text{loc}}}$. Substituting $E_{p,c} = U_{p,c} - W_{p,c} S_{p,c-1}$ (Section 2.4.2) into $\tilde{S}_C = F_{p,c} S_{p,c-1} + \hat{K}_{p,c}^\top E_{p,c}$ gives per-chunk transition $J_{p,c}^{\text{step}} \in \mathbb{R}^{d_k \times d_k}$ and offset $B_{p,c} \in \mathbb{R}^{d_k \times D_{v,\text{pack}}}$ such that, without momentum, each chunk is the affine map

$$\begin{aligned} S_{p,c} &= J_{p,c}^{\text{step}} S_{p,c-1} + B_{p,c}, \\ J_{p,c}^{\text{step}} &= F_{p,c} - \hat{K}_{p,c}^\top W_{p,c}, \qquad B_{p,c} = \hat{K}_{p,c}^\top U_{p,c}; \end{aligned}$$

$J_{p,c}^{\text{step}}$ does not depend on the values, so one transition serves all $D_{v,\text{pack}}$ value columns of a key head (this holds for any gated delta rule, since the transition only involves keys and gates). With $\mu = 0$, the L_{loc} local chunks compose into a single shard map (J_p, B_p) satisfying $S_{p,L_{\text{loc}}} = J_p S_{p,0} + B_p$, where $J_p = \prod_{c=L_{\text{loc}}}^{1} J_{p,c}^{\text{step}}$ (ordered from $c = 1$ on the right to $c = L_{\text{loc}}$ on the left) and B_p is the shard output $S_{p,L_{\text{loc}}}$ obtained from $S_{p,0} = 0$. (With $\mu > 0$, the same construction applies to the lifted maps below.) Writing shard maps as pairs (J_a, B_a) and (J_b, B_b) for shards a and b with composition (apply a, then b)

$$(J_b, B_b) \circ (J_a, B_a) = (J_b J_a,\ J_b B_a + B_b),$$

for any third shard d with map (J_d, B_d), matrix associativity and left-distributivity yield

$$\begin{aligned} \big((J_d, B_d) \circ (J_b, B_b)\big) \circ (J_a, B_a) &= (J_d J_b J_a,\ J_d J_b B_a + J_d B_b + B_d) \\ &= (J_d, B_d) \circ \big((J_b, B_b) \circ (J_a, B_a)\big), \end{aligned}$$

so the P shard maps can be combined by any parallel prefix scan in $\lceil \log_2 P \rceil$ rounds.

When the velocity $Z_c \in \mathbb{R}^{d_k \times D_{v,\text{pack}}}$ is carried across shard boundaries ($\mu > 0$, `SLA Tree-CP+Mom`; shard subscript p dropped, with per-chunk transition J_c^{step} and offset B_c), we have $\Delta S_c = \tilde{S}_C - F_c S_{c-1} = (J_c^{\text{step}} - F_c) S_{c-1} + B_c$. Substituting into Eq. (6) ($Z_c = \mu \kappa_c Z_{c-1} + \Delta S_c$, $S_c = F_c S_{c-1} + Z_c$) defines a lifted $2d_k \times 2d_k$ transition $\tilde{J}_c \in \mathbb{R}^{2d_k \times 2d_k}$ and offset $\tilde{B}_c \in \mathbb{R}^{2d_k \times D_{v,\text{pack}}}$ updating the stacked carry $[S_c; Z_c] \in \mathbb{R}^{2d_k \times D_{v,\text{pack}}}$ via the exact block-affine map

$$\begin{bmatrix} S_c \\ Z_c \end{bmatrix} = \underbrace{\begin{bmatrix} J_c^{\text{step}} & \mu\kappa_c I_{d_k} \\ J_c^{\text{step}} - F_c & \mu\kappa_c I_{d_k} \end{bmatrix}}_{\tilde{J}_c} \begin{bmatrix} S_{c-1} \\ Z_{c-1} \end{bmatrix} + \underbrace{\begin{bmatrix} B_c \\ B_c \end{bmatrix}}_{\tilde{B}_c}. \tag{13}$$

For consecutive shards a and b, the shard-reduced lifted maps $(\tilde{J}_a, \tilde{B}_a)$ and $(\tilde{J}_b, \tilde{B}_b)$ obey the identical associative composition law $(\tilde{J}_b, \tilde{B}_b) \circ (\tilde{J}_a, \tilde{B}_a) = (\tilde{J}_b \tilde{J}_a, \tilde{J}_b \tilde{B}_a + \tilde{B}_b)$, with $[512, 512]$ transitions for $d_k = 256$. □

3.4 MoE Pre-Training Setup, Architecture, and Throughput

All pre-training runs within each cohort share identical optimizer, tokenization, MoE routing, and batch-ordering configurations so that the attention layer design is the sole independent variable:

- **Optimizer and Schedule:** AdamW [42] (moment coefficients $(0.9, 0.95)$, numerical stability constant $\epsilon = 10^{-8}$, weight decay 0.1 with embedding/norm exclusion), global gradient norm clip 1.0, router load-balancing weight 0.01, z-loss 10^{-5}, peak learning rate 3.0×10^{-4} with 5% linear warmup and cosine decay to 3.0×10^{-5} (plus tuned LR $= 4.5 \times 10^{-4}$ with RoPE base frequency $\theta_{\max} = 5 \times 10^5$ [61] for fair `0L:1A` GQA+RoPE baselines across all scales). `+MS` denotes multi-scale depth-staggered decay bias initialization across linear layers.
- **Corpus, FP8 State PTQ, and QA Scoring Protocol:** Real Dolma 3.5 [60, 66, 67] streams (training sequence length $T_{\text{train}} = 8{,}192$) with deterministic `OlmoCoreOrderSampler` batch parity and document-boundary EOS resets (`eos_id = 100257`). In-loop validation tracks geometric-mean PPL across all 11 OLMo v3 validation sets. Post-training evaluation covers: (1) a disjoint 16.78M-token (512×32K) Dolma 3 split (`4K, 8K, 16K, 32K` full and `Tail-4K` PPL, the latter on the last 4K tokens of each sequence; true row-scaled E4M3 FP8 recurrent-state PTQ shifts 8K and 32K NLL by only $+0.00001$ and $+0.00006$ nats with 0.00% NIAH loss across SLA checkpoints, whereas an earlier post-layer activation-quantization probe altered 32K Tail PPL by -0.07 to -0.54 and is excluded from state-PTQ claims), (2) multi-needle Passkey (`MK-NIAH` 4/8-Needle, $N = 256$–512 evaluation items), (3) 9-task zero-shot `lm-eval` [21] across 11,793 items (1 SE $\approx$ 0.65 pts), and (4) real-text 4-candidate `RULER` QA [32] on `SQuAD` [54] and `HotpotQA` [75] at 8K, 32K, 64K (25% chance): unless labeled $N = 512$ (1 SE $\approx$ 2.0 pts), tables report raw log-likelihood `cand4_acc` at $N = 128$ (1 SE $\approx$ 4.0 pts; byte-length-normalized `cand4_acc_norm` agrees within $\pm$2–5 pts, e.g. 82.0/76.6% vs. 76.6/82.0% on clamped 63.1B KDA).
- **Four Sparse MoE Scales (10.7× Active Span; 20 tok/active param ≈ 1× Chinchilla [31]):** Across two MXU-aligned backbone width tiers ($D = 1024$ and $D = 1792$), we evaluate (1) 0.35**B active /** 6.5**B label** (7.58B exact total including embeddings, Table 15; decoder depth $L = 10$, $D = 1024$, `4L:1A`, $k = 4$, 15.10B and 30.20B tok), (2) 0.79**B active /** 12.5**B total** ($L = 16$, $D = 1024$, `7L:1A`, `15L:1A`, & `Pure-Linear` ∞:1, $k = 8$, 15.10B and 60.40B tok), (3) 1.80**B active /** 31.5**B label** (30.58B exact total, Table 15; $L = 15$, $D = 1792$, `4L:1A`, $k = 8$, 5.03B and 36.20B tok $\approx 1.0\times$ Chinchilla), and (4) 3.74**B active /** 63.1**B total** ($L = 30$, $D = 1792$, `4L:1A` & `15L:1A`, $k = 8$, 5.03B/30.2B RoPE on 64 chips and 120.80B NoPE on 512 chips).

Hybrid `LatentMoE` backbone. All four model scales belong to a single hybrid linear-softmax `LatentMoE` decoder family implemented in open-source MaxText [1], interleaving x linear delta layers with one gated GQA [2, 52] layer per $(x{+}1)$-layer cycle (using RoPE [61], or NoPE with Scalable Softmax [45] at 63.1B), with a dense SwiGLU [56] FFN on Layer 0 and one shared expert plus $E = 256$ routed latent experts on Layers $1, \ldots, L-1$. Model widths stay divisible by 256 ($D = 1024$ at 6.5B and 12.5B; $D = 1792$ at 31.5B and 63.1B), while every SLA layer fixes $d_k = 256$ so recurrent state updates never pad systolic tiles. In `Iso-Param SLA-R4`, widening the shared expert by $+256$ channels compensates for the mixer parameters saved by 4:1 key sharing; complete initialization, normalization, routing, and per-scale widths are given in Appendix D (Table 15).

End-to-end MoE training throughput vs. MaxText KDA and Tokamax Mosaic KDA. We pre-train four token-choice sparse MoE families (6.5B–63.1B total and 0.35B–3.74B active parameters; 256 routed experts, top-8 routing, one shared expert [19, 57]) with AdamW [42] on packed Dolma 3.5 [60, 66, 67] at $T_{\text{train}} = 8{,}192$, using 64–512 TPU v7x chips per run across a 2,048-chip cluster (Table 25). On 12.5B (`7L:1A`, 16L), `SLA-R4` achieves a 1.63× end-to-end single-microbatch training speedup over MaxText Pallas KDA (469.9 ms vs. 766.8 ms) and a 1.24× speedup under 4-step gradient accumulation (1,800 ms vs. 2,225 ms at 8.39M tok/step; 1.26× at 1,686.0 vs. 2,119.7 ms in XProf Table 23), and across all matched scales SLA trains 11%–24% (1.11×–1.24× with GA $= 4$; 1.08× for the short 5B-token 31.5B `SLA-R2` run; up to 1.63× single-microbatch) faster than MaxText KDA (Tables 7 and 8), including 1.11×–1.15× at 63.1B on 512 chips. To benchmark end-to-end 12.5B (`7L:1A`) step time directly against open-source Tokamax Mosaic KDA (`tokamax` 0.0.14, $C = 64$), Appendix Table 19 evaluates the full 16-layer 12.5B model (14 linear layers + 2 global GQA layers + 16 `LatentMoE` layers with $k = 8$ active experts + AdamW):

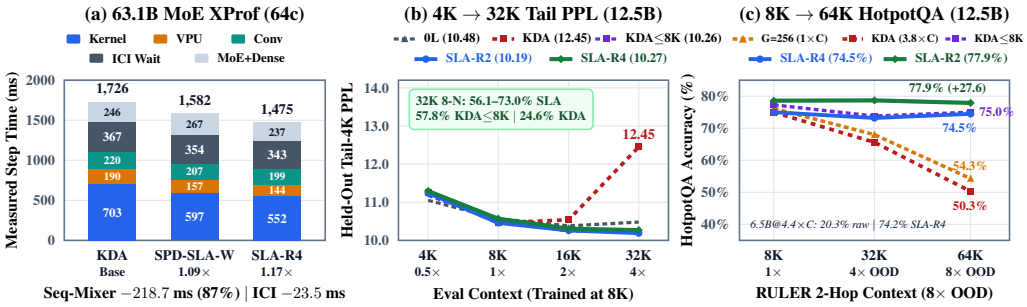

Figure 7: **(a)** Per-step XProf breakdown for 63.1B MoE on 64 TPU v7x chips across `KDA Base`, `SPD-SLA-W (R2)` (`SLA-R2` with optional in-chunk key whitening, Appendix A), and `SLA-R4` (Table 22); most of SLA's saving is in linear attention ops. **(b)** Held-out Dolma tail perplexity from 4K to 32K for 12.5B MoEs at 60.4B tokens (curves: seed `s0`; callout: 3-seed means from Table 8a). **(c)** Zero-shot 8K → 64K RULER `HotpotQA` at 1× and 3.8× Chinchilla tokens (C). Curves are seed means (3 seeds; 2 for $G = 256$; Table 26), except KDA≤8K, which clamps the `s0` KDA model's horizons to 8K without retraining.

Table 7: **Multi-scale SLA vs. MaxText KDA pre-training speed and quality** across 6.5B–63.1B sparse `LatentMoE` models (64–512 TPU v7x chips per run, $T_{\text{train}} = 8{,}192$, `Exact Inv.`). Complete 45-row ladder and per-subtask `lm-eval` breakdowns are in Appendix Tables 24–26.

Scale & Horizon	Linear Mixer	Step ms	vs. KDA	11-Set PPL	`lm-eval`	32K Tail	32K 8-N	64K `SQuAD / Hotpot`
6.5**B** (0.35**B act.**) 4L:1A, 30.2B tok (64c)	MaxText KDA Base (FP32)	351.5	1.00×	**16.16**	49.7%	620.9	16.8%	32.0%/20.3%
	SLA-R2 (BF16)	**315.0**	1.12×	16.25	49.5%	13.53	21.1%	56.2%/51.6%
	Iso-Param SLA-R4 (BF16)	**297.9**	1.18×	16.33	49.1%	**11.09**	45.3%	70.3%/74.2%
12.5**B** (0.79**B act.**) 7L:1A, 15.1B tok (64c)	MaxText KDA Base ($n = 3$)	556.0	1.00×	17.64	47.8%	13.28	18.1%	57.3%/56.0%
	SLA-R2 Unfused ($n = 3$)*	**497.3**	1.12×	**17.53**	48.6%	**11.23**	52.3%	69.5% /73.7%
	SLA-R2 Fused ($n = 2$)*	**481.8**	1.15×	17.81	48.2%	11.36	58.0%	69.1%/ 74.6%
12.5**B** (0.79**B act.**) 7L:1A, 60.4B tok (64c)	MaxText KDA Base ($n = 3$)	2,225	1.00×	15.60	50.4%	12.45	24.6%	54.9%/50.3%
	KDA + `TrainMaxTau=8K` ($n = 2$)	2,251	0.99×	15.80	50.3%	10.75	49.6%	75.8% /75.4%
	SLA-R2 (BF16, $n = 3$)*	**1,919**	1.16×	**15.58**	50.9%	**10.19**	56.6%	75.5%/ 77.9%
	Iso-Param SLA-R4 ($n = 3$)*†	**1,800**	1.24×	15.66	50.8%	10.27	56.1%	74.7%/74.5%
31.5**B** (1.80**B act.**) 4L:1A, 5.03B tok (64c)	MaxText KDA Base (FP32)	807.7	1.00×	**21.23**	–	47.16	12.9%	–
	SLA-R2 Fused (BF16)	**747.2**	1.08×	21.44	–	**13.88**	33.2%	–
	SLA-R4 Fused (BF16)	**707.4**	1.14×	21.71	–	14.20	46.1%	–
63.1**B** (3.74**B act.**) 4L:1A, 120.8B tok (512c, NoPE)‡	MaxText KDA Base (FP32)	2,165	1.00×	12.60	57.8%	353.4	15.6%	27.7%/27.0%
	$+\tau_{\max} = 8$K Clamp + `SSMax`	–	–	12.60	57.8%	**8.04**	74.2%	90.6%/86.5%
	SLA-R2 (0.5× Bytes)	**1,953**	1.11×	12.58	57.5%	13.78	18.6%	55.7%/56.8%
	+Zero-Clamp `SSMax`	–	–	12.58	57.5%	8.45	84.4%	–
	SLA-R4 (0.5× Bytes) + `SSMax`†	**1,878**	1.15×	**12.45**	56.7%	8.41	83.2%	75.6%/79.1%

*Paired seed ΔPPL vs. KDA: -0.10 ± 0.08 (12.5B/15.1B, $n = 3$; fused $n = 2$: 17.81 vs. 17.89), -0.02 ± 0.11 (`SLA-R2`, 60.4B), and $+0.06 \pm 0.26$ (`SLA-R4`, 60.4B). Zero-shot row-scaled FP8 state changes `lm-eval` by at most 0.25 pts (below one standard error) across all 20 evaluated models (Appendix Table 26). †Deferring the fused conv `AllReduce` across gradient accumulation gives 1,812 ms (12.5B) and 1,843 ms (63.1B). ‡On 512 chips (63.1B), throughput rises from 15,135 tok/s/chip (KDA) to 16,779 (`SLA-R2`) and 17,445 (`SLA-R4`); pre-`SSMax` `SLA-R4` is 22.78 Tail PPL, 18.0% 8-N, and 51.6%/45.1% QA (clamp-only KDA without `SSMax`: 7.70 Tail, 48.4% 8-N, 76.6%/82.0% QA at $N = 128$; all from Appendix Table 24).

fused 1-kernel `SLA-R4` is 1.68× (23.61 ms, $\mu = 0.25$) to 1.73× (22.85 ms, $\mu = 0$) faster on the 14-layer linear mixer F+B than Tokamax Mosaic KDA (39.62 ms; 2.65×–2.74× vs. MaxText KDA 62.50 ms), delivers a 1.24× (66.01 vs. 82.15 ms, $\mu = 0.25$) to 1.26× (65.24 ms, $\mu = 0$) and 1.33× (61.83 ms, 1-step Neumann) end-to-end 12.5B step speedup over Tokamax Mosaic KDA on single-chip active-parameter steps (1.64×–1.75× vs. MaxText KDA 108.06 ms), delivers 1.10×–1.11× (121.59/120.79 ms vs. 134.06 ms) on 8-chip `EP=8` including MoE dispatch/combine all-to-all (1.87× vs. MaxText FP32 227.00 ms), and achieves a 1.14× (1,800 vs. 2,049 ms at GA = 4) to 1.42× (469.9 vs. 668.4 ms single-microbatch; 1.18×–1.74× with 1-kernel $\mu = 0$) end-to-end 12.5B speedup over Tokamax Mosaic KDA on 64-chip TPU v7x. Under a unified single-protocol 4-scale benchmark across all four MoE rungs (6.5B, 12.5B, 31.5B, and 63.1B at $T_{\text{train}} = 8{,}192$, 3 process restarts; Appendix Table 20b), 1-kernel `SLA-R4` (`Exact Inv.`) trains 4.0%–21.4% faster (1.040×–1.214×) than open-source Tokamax Mosaic KDA (262.2 vs. 272.6 ms at 6.5B, 330.4 vs. 348.5 ms at

12.5B, 495.2 vs. 572.7 ms at 31.5B, and 1,015.7 vs. 1,232.9 ms at 63.1B) and by $1.166\times$–$1.224\times$ relative to MaxText Pallas KDA. At native $T_{\text{train}} = 32\text{K}$ with 4.19M tokens/step (EP = 128), SLA fits in per-core HBM (≤ 94.74 GiB) while KDA (96.74 GiB) and full attention (95.48 GiB) exceed the 95 GiB cap; in 8K $\rightarrow$ 32K continued pre-training at 12.5B, SLA is $1.08\times$–$1.15\times$ faster than KDA within 0.004–0.029 nats in 32K cross-entropy (Table 25). Hardware `XProf` traces on 64-chip 63.1B MoE attribute 96.6% of SLA's step-time saving to linear attention ops and reduced ICI wait (Figure 7a; Appendix F).

3.5 Pre-Training Quality and Long-Context Extrapolation

Table 8: **Controlled 12.5B MoE (0.79B active, 16L, $T_{\text{train}} = 8{,}192$) baselines, factor-isolation ladder (`A1`–`A6`), and 90-min Iso-Wall comparison** at (a) 60.4B tokens ($3.8\times$ Chinchilla [31]) and (b) 15.1B tokens ($1.0\times$ Chinchilla; GDN-2 uses our unfused TPU port; seeds `s0`/`s1`/`s2` use distinct data offsets, so PPL is compared on matched seeds).

Cohort	Architecture (7,200 Steps)	Step ms	vs. KDA	11-Set PPL	8K Tail	32K Tail	32K 8-N	64K `SQuAD / Hotpot`
	0L:1A Full-Attn GQA+RoPE (`s0`)	2,314	0.96×	16.13	10.57	10.48	12.1%	58.6%/53.1%
	7L:1A MaxText KDA Base ($n = 3$)*	$2{,}225 \pm 29$	1.00×	15.60 ± 0.58	10.46	12.45 ± 1.11	24.6%	54.9%/50.3%
	+Zero-Retrain $\tau_{\max} = 8\text{K}$ Clamp (`s0`)	–	–	16.17	10.53	10.26	57.8%	**78.1**%/75.0%
(a) v7x	7L:1A KDA + `TrainMaxTau=8K` ($n = 2$)	$2{,}251 \pm 11$	0.99×	15.80 ± 0.33	10.46	10.75 ± 0.21	49.6%	75.8%/75.4%
60.4B Tok	**7L:1A SLA-R2** (BF16, $n = 3$)*	$\mathbf{1{,}919 \pm 37}$	1.16×	15.58 ± 0.51	10.46	$\mathbf{10.19 \pm 0.10}$	56.6%	75.5%/**77.9**%
(3.8× Ch.)	**7L:1A SLA-R4** (Iso-Param, $n = 3$)*†	$\mathbf{1{,}800 \pm 39}$	1.24×	15.66 ± 0.41	10.57	10.27 ± 0.04	56.1%	74.7%/74.5%
	7L:1A SLA-R4 (−20% mixer, `s0`)	1,819	1.22×	16.20	10.81	10.41	**73.0**%	72.7%/75.0%
	15L:1A MaxText KDA Base (FP32, `s0`)	2,152	1.00×	15.50	10.46	11.27	15.2%	62.5%/60.2%
	15L:1A SLA-R2 (Uniform, BF16, `s0`)	**1,941**	1.11×	**15.44**	**10.43**	**10.19**	41.0%	72.7%/70.3%
	7L:1A MaxText KDA Base ($n = 3$)§	556.0	1.00×	17.64 (17.89)	11.57	13.28	18.1%	57.3%/56.0%
	7L:1A SLA-R2 (Unf., `Exact Inv.`, $n = 3$)§	**497.3**	1.12×	**17.53** (17.75)	11.54	**11.23**	52.3%	69.5%/73.7%
	7L:1A SLA-R2 (Fus., `Exact Inv.`, $n = 2$)	**481.8**	1.15×	17.81	11.70	11.36	58.0%	69.1%/74.6%
	7L:1A GDN-2 (Unf., 93m, 15.1B, $n = 2$)	778.1	0.71×	**17.59**	**11.44**	11.86	19.3%	62.1%/62.1%
(b) 15.1**B**	**7L:1A SLA-R2** (Fus., 90m Iso-Wall, 23.1B tok)	**488.3**	1.14×	16.66 (+53% tok)	10.84	10.98	52.3%	65.6%/70.3%
(1.0× Ch.	7L:1A `A1`: Wide-Key R=2 KDA (G=256, n=2)	524.8	1.06×	17.69	11.61	13.39	23.8%	55.9%/54.3%
& 90m Iso)	7L:1A `A2`: $+G = 4$ Decay (`Exact Inv.`, $n = 2$)	514.1	1.08×	17.75	11.65	11.34	46.1%	70.7%/73.8%
	7L:1A `A3`: +1-Step Neumann Inv. ($n = 2$)	497.2	1.12×	17.97	11.69	11.36	53.5%	71.9%/71.9%
	7L:1A `A5`: SLA-R2 (1-Step, $\mu = 0.25$, `s0`)	**467.1**	1.19×	18.02	11.87	11.45	**60.9**%	**74.2**%/**83.6**%
	7L:1A `A6`: SLA-R2 + FP8 Out-QAT (`s0`)	520.5	1.07×	18.01	11.80	11.37	59.8%	**74.2**%/75.8%
	15L:1A SLA-R2 (Uniform, BF16, $n = 3$)	**486.7**	1.14×	17.61	11.60	11.42	34.8%	69.0%/65.1%

*Paired 3-seed ΔPPL vs. KDA: `SLA-R2` -0.02 ± 0.11, `SLA-R4` $+0.06 \pm 0.26$. †Deferred conv `AllReduce` across `GA=4`: 1,812 ms (Appendix Table 23). §Parentheses: `s0,s1` mean matched to $n = 2$ rows (paired 3-seed `SLA-R2` vs. KDA: -0.10 ± 0.08 PPL, winning 3/3 seeds). On `s0`: KDA 18.07, `A1` 17.85, `A2` 17.87, fused `SLA-R2` 17.91, `A3` 18.10, `A5` 18.02, `A6` 18.01. Omitted label `A4` is the single-seed `A3` + BF16-state run (`s0`: 476.0 ms, 1.17×; 11.86/12.04 8K/32K Tail, 41.0% 8-N, 58.6%/58.6% 64K QA), excluded from the table because its final 11-set validation log was not archived (8K full Dolma PPL: 12.23 vs. 12.18 for `A3` on `s0`).

In-domain quality vs. KDA. At 12.5B parameters and 60.4B tokens with three paired seeds (Table 8a), `SLA-R2` (-0.02 ± 0.11) and `SLA-R4` ($+0.06 \pm 0.26$) match KDA within seed noise on 11-set Dolma perplexity and zero-shot `lm-eval` [21] (Table 26); single-seed 63.1B runs likewise match closely (12.58 and 12.45 vs. 12.60, Table 7). At 6.5B, `SLA-R2` trails KDA by 0.20 perplexity over three seeds (19.08 vs. 18.88), and short 5B-token runs at 31.5B and 63.1B trail by 0.1–0.5 (Table 25). In the 15.1B-token cohort (Table 8b), seeds `s0`, `s1`, and `s2` use distinct data offsets that shift perplexity by up to 0.9, so we compare only on matched seeds. Over all three seeds, `SLA-R2` (`Exact Inv.`, BF16 state, $\mu = 0.25$) slightly improves on KDA Base on every seed (17.533 vs. 17.638; paired difference -0.10 ± 0.08) and matches the same-kernel $G = 4$ FP32 control without momentum (17.538 on `s0..s2`; 17.590 across $n = 6$, Table 10).

Factor isolation (`A1`–`A6`). On matched seeds `s0,s1` (Table 8b), widening KDA's keys to $d_k = 256$ with $R = 2$ head sharing (`A1`, $[4, 256, 512]$, $G = 256$: 524.8 ms, 17.694 PPL) isolates head sharing from decay grouping. Switching from $G = 256$ to $G = 4$ with MXU block-doubling (`A2`: 514.1 ms), adding BF16 state and momentum (497.3 ms unfused), and fusing into VMEM (481.8 ms) saves 43.0 ms total while staying within 0.06–0.11 PPL of `A1` (17.749–17.806) and lifting 32K 8-Needle from 23.8% to 46.1%–58.0% and 64K `HotpotQA` from 54.3% to 73.7%–74.6%. The optional 1-step Neumann inverse is the one approximation with a clear quality cost: `A3` rises to 17.970 (+0.22 over `A2`), and on `s0` the fused Neumann arm `A5` reaches 18.016 versus 17.913 for fused `Exact Inv.`, which is why all flagship pre-training runs use exact inversion (`Exact Inv.`).

Comparison with GDN-2 and multi-scale equal-chip-hour pre-training vs. Tokamax KDA. GDN-2 [29] decouples the scalar delta gate into channel-wise erase and write gates while keeping

per-channel decay. On TPU v7x ($B = 1$, 8K/32K F+B), our unfused XLA port of per-channel GDN-2 takes 63.86/251.39 ms at $D = 1792$ (15,217 MiB temporary HBM) and 34.15/132.83 ms at $D = 2048$ (isolated Mamba-3 MIMO $R = 4$ kernel for reference: 9.27 ms), raising end-to-end MoE step time by 1.23×–1.61× (1.349 vs. 1.093 s on 128-chip TPU v4; 778.1 ± 3.9 vs. 481.8 ms for fused `SLA-R2` on 64-chip TPU v7x). Per token, GDN-2 reaches the best 8K perplexity in the 15.1B-token cohort (17.592 ± 0.188 vs. 17.806 for fused `SLA-R2` and 17.893 for KDA on `s0,s1`; 17.459 vs. 17.913 on `s0`; Table 8b), though its unclamped per-channel gate extrapolates worse (19.3% vs. 58.0% 32K 8-Needle). Under an equal ≈90-minute pure-step budget on `s0` (5,371 s for 11,000 fused `SLA-R2` steps vs. 5,582 s for 7,200 GDN-2 steps), SLA processes 53% more tokens and reaches 16.660 perplexity versus 17.459. Because this speed gap reflects our unfused TPU port of GDN-2 as much as architecture, we built a single-kernel VMEM-fused `SLA-R4 + GDN-2` mixer (combining SLA's $R = 4$ wide-key tile-grouped mapping with decoupled erase/write gates) and compared both 1-kernel fused `SLA-R4` ($\mu = 0$) and fused `SLA-R4 + GDN-2` ($\mu = 0$) directly against open-source Tokamax Mosaic KDA ($C = 64$, safe-clipped $g \geq -3.5$ to prevent float32 `exp2` backward overflow). On single-chip 760M–822M dense LMs ($T = 4$K FineWeb-Edu; Appendix Table 38), fused `SLA-R4` (1.28×) and `SLA-R4 + GDN-2` (1.23×) beat safe-clipped Tokamax KDA by -14.45 ($t = -43.7$) and -13.70 ($t = -19.5$) Val PPL at 180 s ($n = 3$) and by -9.68 and -9.93 at 360 s ($n = 2$). Table 9 scales this comparison to 64-chip TPU v7x sparse `LatentMoE` pre-training on packed Dolma 3.5 ($T_{\text{train}} = 8{,}192$, 2.10M tok/step, EP = 128, $n = 3$ paired seeds per arm; per-seed breakdown in Appendix Table 39). Each run calibrates its steady-state step time after warm-up and executes the exact step budget fitting a fixed post-compile wall-clock window (40.0 min = 42.67 chip-h at 12.5B; 35.0 min = 37.33 chip-h at 63.1B), while logging matched iso-token 11-set OLMo v3 validation perplexity along the way. At 12.5B (`7L:1A`), fused iso-parameter `SLA-R4` (455.2 ms/step, 1.125×, +13.3% tokens) matches Tokamax KDA within routing noise at step 4,000 ($+0.11 \pm 0.14$ PPL, $p = 0.29$) while fused `SLA-R4 + GDN-2` (479.2 ms/step, 1.069×, +7.7% tokens) edges both per token (19.90 vs. 20.05 and 20.16); at equal chip-hours, both beat Tokamax KDA on all 3/3 seeds, with `SLA-R4` lowering 11-set PPL by -0.37 ± 0.11 ($t = -6.02$, two-sided $p = 0.027$), raising 9-task `lm-eval` by $+0.92 \pm 0.31$ points ($t = +5.16, p = 0.036$), and lowering 32K Dolma tail PPL by -1.81 ± 0.68 ($p = 0.044$), while `SLA-R4 + GDN-2` reaches 19.11 ± 0.33 PPL (-0.41 ± 0.35, 3/3 seed wins). At 63.1B (`4L:1A`), fused `SLA-R4` (1,508.3 ms, 1.125×, +12.9% tokens) and `SLA-R4 + GDN-2` (1,548.7 ms, 1.096×, +9.1% tokens) already pull ahead of $g \geq -3.5$ clipped Tokamax KDA by step 1,200 (-0.82 ± 0.18 and -0.94 ± 0.22 PPL, $p = 0.016$ and 0.018, 3/3 wins) and widen their equal-chip-hour advantage to -1.48 ± 0.21 PPL ($t = -12.08, p = 0.0068$, 3/3 wins) and -1.41 ± 0.16 PPL ($t = -15.01, p = 0.0044$, 3/3 wins), alongside a -3.35 ± 0.88 reduction in 32K tail PPL ($p = 0.022$) and +16.1 points on 32K 8-Needle ($p = 0.050$).

Table 9: **Multi-scale equal-chip-hour and iso-token MoE pre-training vs. open-source Tokamax Mosaic KDA on 64-chip TPU v7x** (128 TensorCores, EP = 128, $T_{\text{train}} = 8{,}192$, 2.10M tok/step, packed Dolma 3.5, $n = 3$ paired seeds `s0..s2` per arm, sample std ddof = 1). Step counts are calibrated per run for an equal post-compile training budget (40.0 min = 42.67 chip-h at 12.5B; 35.0 min = 37.33 chip-h at 63.1B). Per-seed logs and disclosures are in Appendix Table 39.

Scale & Budget	Linear Mixer (1-Kernel `Exact Inv.`)*	Measured Step p_{50} (ms)	Mean Steps (Tokens)	Iso-Token PPL (Early / Late)	Eq.-Chip-h 11-Set PPL	Paired ΔPPL (t, two-sided p)	8K → 32K Tail PPL	32K 8-N	`norm9` / 64K `Hotpot`
(a) 12.5**B MoE**	Tokamax Mosaic KDA (C=64, clip)	512.1 ± 6.6 (1.00×)	4,857 (10.2B)	24.78 / 20.05	19.52 ± 0.47	ref	12.75 → 13.98	34.8%	46.39% / 54.2%
`7L:1A`, D=1024	**Fused Iso-Param** `SLA-R4` (μ=0)	$\mathbf{455.2 \pm 6.0}$ (**1.125**×)	5,503 (11.5B)	25.04 / 20.16	$\mathbf{19.15 \pm 0.44}$	$\mathbf{-0.37 \pm 0.11}$ (-6.02, .027)	**12.50 → 12.16**	54.3%	**47.31% / 68.0%**
42.67 chip-h (n=3)	**Fused** `SLA-R4 + GDN-2` (μ=0)	479.2 ± 3.7 (**1.069**×)	5,230 (11.0B)	**24.59 / 19.90**	$\mathbf{19.11 \pm 0.33}$	$\mathbf{-0.41 \pm 0.35}$ (-2.01, .182)	12.36 → 12.15†	**61.7%**†	47.10%† / **69.5%**†
(b) 63.1**B MoE**	Tokamax Mosaic KDA (C=64, clip)	$1{,}697.5 \pm 6.5$ (1.00×)	1,243 (2.61B)	36.69 / 26.73	26.56 ± 0.32	ref	16.92 → 19.45	20.0%	40.98%‡ / 42.6%‡
`4L:1A`, D=1792	**Fused Iso-Param** `SLA-R4` (μ=0)	$\mathbf{1{,}508.3 \pm 6.8}$ (**1.125**×)	1,403 (2.94B)	36.73 / **25.91**	$\mathbf{25.08 \pm 0.44}$	$\mathbf{-1.48 \pm 0.21}$ (-12.1, .0068)	**16.01 → 16.10**	**36.2%**	**42.02%**‡ / –
37.33 chip-h (n=3)	**Fused** `SLA-R4 + GDN-2` (μ=0)	$1{,}548.7 \pm 2.5$ (**1.096**×)	1,357 (2.85B)	**36.13 / 25.78**	$\mathbf{25.14 \pm 0.21}$	$\mathbf{-1.41 \pm 0.16}$ (-15.0, .0044)	16.14 → 16.37‡	28.9%‡	–

*Iso-token checkpoints are steps 2,000/4,000 (4.19B/8.39B tok) at 12.5B and steps 600/1,200 (1.26B/2.52B tok) at 63.1B. Parameter deltas vs. Tokamax KDA (12.53B and 63.13B): `SLA-R4` has −3.0M (12.5B, shared expert 1024 → 1280) and −14.4M (63.1B, shared expert 3584 → 4352) fewer parameters; `SLA-R4 + GDN-2` has +1.7M (12.5B, shared expert 1024) and −33.5M (63.1B, shared expert 3584). Tokamax KDA (`tokamax` 0.0.14) requires $g \geq -3.5$ ($f \geq 0.030$) to avoid float32 `exp2` backward NaNs (0 non-finite steps across all 18 runs), which bounds sub-chunk fast forgetting on the baseline. †At 12.5B, all 9 runs completed 100% of training and all iso-token/final 11-set PPL evals ($n = 3$); for `SLA-R4` and Tokamax KDA, post-train OOD/downstream/N=512 64K QA are $n = 3$ (`SQuAD` 66.0% vs. 53.3%, $p = 0.067$), whereas `SLA-R4 + GDN-2` post-train OOD/downstream reflects seed `s1` (vs. Tokamax `s1`: 12.62 → 14.58 Tail, 54.7% 8-N, 46.85% `norm9`, 45.1% `HotpotQA`). ‡At 63.1B, all 9 runs completed 100% of training and all iso-token/final 11-set PPL evals ($n = 3$); `SLA-R4` and Tokamax KDA completed $n = 3$ 8K/32K Dolma & 32K NIAH (`SLA-R4 + GDN-2`: $n = 2$, `s1,s2`, vs. Tokamax 16.94 → 19.89 and 22.7%), and matched seed `s0` `norm9` is shown before reservation cutoff (all completed Tokamax `s0..s2` `norm9` [41.22% mean] and `s0,s1` 64K QA scores are in Appendix Table 39b).

Length extrapolation. We test zero-shot extrapolation from 8K training to 32K–64K on RULER [32] needle retrieval and 4-choice `SQuAD` [54] and `HotpotQA` [75] (Figure 7b–c). Unclamped

KDA degrades beyond the training length, severely at 63.1B, because about 1% of its channels learn retention horizons ($\tau = -1/\ln f$ tokens for per-token decay f) far beyond 8K (Table 24). SLA extrapolates without clamping: at 12.5B its 32K tail perplexity stays flat and 64K `HotpotQA` reaches 74%–78%. The advantage is limited, however. Clamping KDA's horizons to 8K, either at test time or during training (`TrainMaxTau=8K`), gives similar results at 12.5B. At 63.1B, SLA also needs evaluation-time `SSMax` [45] on the softmax layers (no retraining; 32K tail perplexity 22.78 → 8.41 for `SLA-R4`, Appendix Table 24), whereas clamped KDA reaches 7.70 without it. With `SSMax` on both (scale $\alpha = 0.5$ for SLA and 0.25 for KDA; one seed each), clamped KDA beats `SLA-R4` on 64K `SQuAD/HotpotQA` (90.6%/86.5% vs. 75.6%/79.1%), whereas SLA leads on 32K 8-Needle retrieval (83.2%–84.4% with zero gate clamping vs. 74.2% clamped KDA and 15.6% unclamped; Table 7). With 128–512 questions per QA set (standard error 2–4 points), many differences in Tables 7–8 are within noise. Retention-horizon dispersion and clamping audits across scales are in Appendix F (Table 24). Auditing the public `Kimi-Linear-48B-A3B` checkpoints confirms this slow-channel tail is intrinsic to per-channel KDA rather than specific to our baseline training recipe (Appendix I, Table 36): across all 81,920 KDA channels (20 KDA layers), Monte Carlo weight projections already show 0.51%–0.61% of channels with $\tau >$ 8K (0.04%–0.05% $>$ 1M), and live 27-layer forward passes of `Kimi-Linear-48B-A3B-Instruct` on FineWeb-Edu show 12.1% of channels above 8K (4.2% $>$ 64K, 0.34% $>$ 1M) at 16K. Although long-context pre-training keeps the public checkpoint finite, zero-retraining clamping to $\tau_{\max} = 4{,}096$ still lowers its 16K perplexity from 9.50 → 9.24 (11.34 → 10.90 tail-2K, −0.44 PPL).

Decay granularity and why $G = 4$ is the quality-speed Pareto knee. Table 10 sweeps $G \in \{1, 4, 16, 64, 256\}$ (tile width $b = 256/G$) on the identical $[4, 256, 512]$ `Exact Inv.` FP32 kernel ($\mu = 0$), expanding the direct $G = 1$ vs. $G = 4$ comparison to $n = 6$ paired seeds (`s0..s5`; per-seed breakdown in Appendix Table 39c). On matched `s0,s1` seeds across all five granularities, 8K 11-set perplexity favors $G \geq 4$ without a strict monotone trend (17.946, 17.749, 17.873, 17.748, 17.694 for $G = 1, 4, 16, 64, 256$; all $G \geq 4$ within 0.18 PPL). Across all $n = 6$ paired 12.5B MoE seeds, $G = 4$ improves over scalar $G = 1$ on 6/6 seeds (ΔPPL $= -0.34, -0.06, -0.06, -0.66, -0.07, -0.21$; mean -0.232 ± 0.235, $t = -2.42$, two-sided $p = 0.060$, exact two-sided sign test $p = 0.031$), improves 8K Dolma tail PPL on 5/6 seeds (11.57 vs. 11.64, -0.08 ± 0.08, $t = -2.48$, $p = 0.056$), and edges 9-task `lm-eval` on 4/6 seeds (48.56% vs. 48.08%, $+0.49 \pm 0.94$ pts) while matching per-channel KDA ($G = d_k$) within 0.055 PPL on `s0,s1`. On hardware, $G = 1$ (scalar GDN) runs at the same single-kernel VMEM speed as $G = 4$ (2.62 vs. 2.85 ms at 8K and 14.16 vs. 14.08 ms at $D = 1792$, 32K in Table 17) because 1 or 4 gate vectors both fit in VMEM vector registers without $[B, H_k, T/C, C, C, d_k]$ HBM materialization, whereas per-channel KDA ($G = d_k$) incurs a 1.62×–2.45× XLA slowdown (and 1.04×–1.12× in Pallas chunk kernels; Table 17). On synthetic zero-shot length extrapolation without clamping, $G = 1$ ($b = 256$, $\tau_{\max} \leq$ 56K) reaches higher 32K 4/8-Needle across $n = 6$ seeds (70.3%/59.6% vs. 46.3%/35.4%, $p = 0.017/0.048$, because unclamped $G = 4$ develops a slow $>$ 100K tile on seeds `s2, s4, s5` while reaching 55.9%–66.0% / 45.7%–61.7% on `s0, s1, s3`) and 3–4 points higher unclamped 64K QA (68.2%/71.0% vs. 64.0%/68.1%, $p = 0.20/0.22$; matched on `s0,s1` at 67.5%/71.2%), while finer groups ($G \geq 16$, $b \leq 16$, $\tau_{\max} =$ 131K) drop sharply to 19.0%–27.1% 8-Needle and 50.8%–61.9% QA. Complementing the $n = 6$ MoE runs without routing noise, our $n = 6$ paired-seed dense LM control (274.7M params, 104.9M FineWeb-Edu tokens, identical fused `SLA-R4` kernel, $\mu = 0$; Appendix I, Table 37) separates $G = 4$ from $G = 1$ on 6/6 seeds across every metric (-1.15 ± 0.49 Val PPL, $t = -5.78$; -5.28 ± 2.61 OOD PPL, $t = -4.96$; -1.75 ± 0.61 16K tail PPL, $t = -7.01$), captures 78.5% of the $G = 1 \to G = 256$ in-domain gain, and beats unclamped $G = 256$ at 16K (46.08 vs. 46.47 full, 46.32 vs. 47.24 tail). On a synthetic 2-layer multi-timescale associative recall diagnostic ($n = 5$ seeds), per-channel $G = 256$ solves the task reliably (84.1% joint) while $G \in \{1, 2, 4\}$ remain bimodal or underfit, so we treat $G \in \{1, 4\}$ as a zero-overhead architectural knob ($G = 4$ for language-modeling perplexity, $G = 1$ for uniform unclamped synthetic passkey retrieval).

FP8 state. SLA states quantize well: zero-retrain row-scaled E4M3 state quantization changes 8K/32K loss by less than 10^{-4} nats (Section 3.4), and FP8 output QAT with `Exact Inv. (A6)` is within 0.09 perplexity of BF16 on `s0` (18.006 vs. 17.913).

Optional low-precision key whitening (`Polar-2+`). In hybrids with sparser softmax refresh such as `15L:1A`, middle linear layers hold 10×–35× more active key-value associations and develop poorly conditioned chunk key Gram matrices. Appendix A introduces `SPD-SLA-W`, an optional extension

Table 10: Unconfounded 5-point $[4, 256, 512]$ decay-granularity sweep on the identical `chunk_kda Exact Inv.` FP32 kernel ($\mu = 0$, 12.5B MoE, `7L:1A`, 15.10B tok), reporting matched `s0,s1` 11-set PPL (valid across all 5 granularities) alongside multi-seed means ($n = 6$ paired seeds `s0..s5` for $G \in \{1, 4\}$; $n = 3$ `s0..s2` for $G \in \{16, 64\}$), paired ΔPPL vs. $G = 1$, 8K $\rightarrow$ 32K Tail PPL, 32K 4/8-N, `lm-eval norm9`, and $N = 512$ (1 SE = 2.0 pts) 64K `RULER` QA.

$[4, 256, 512]$ `Exact Inv.` FP32 Arm ($\mu = 0$)	G (b)	Within-Tile Std	`s0,s1` Mean PPL (n-Seed Mean)*	Paired ΔPPL vs. $G = 1$	8K $\rightarrow$ 32K Tail	32K 4-N / 8-N	`norm9`	$N = 512$ 64K `SQuAD / Hotpot`
Full-Tile Scalar Decay ($n = 6$)	**1** (256)	**0.00**	17.946 (17.822 $\pm$ 0.606)	0.000 (ref)	11.64 $\rightarrow$ **11.34**	**70.3%/59.6%**	48.08%	**68.2 $\pm$ 2.5%/71.0 $\pm$ 1.3%**
Systolic Sub-Tile Default ($n = 6$)	**4** (64)	**0.00**	17.749 (**17.590 $\pm$ 0.391**)	**-0.232 ± 0.235** (6/6, $t = -2.42$)	**11.57** $\rightarrow$ 11.43	46.3%/35.4%*	**48.56%**	64.0 $\pm$ 4.6%/68.1 $\pm$ 4.3%*
Intermediate Sub-Tile ($n = 3$)	16 (16)	4.34	17.873 (17.626)	-0.064 ($-0.14, -0.00, -0.05$)	11.75 $\rightarrow$ 12.12 $\pm$ 0.72	41.4%/27.1%	48.35%	60.3 $\pm$ 4.2%/61.9 $\pm$ 5.4%
Fine Sub-Tile Decay ($n = 3$)	64 (4)	5.05	17.748 (17.523)	-0.167 ($-0.36, -0.04, -0.11$)	11.53 $\rightarrow$ 12.94 $\pm$ 0.23	27.7%/19.0%	48.61%	53.5 $\pm$ 2.2%/50.8 $\pm$ 1.8%
Per-Channel Wide-Key ($n = 2$ PPL, 3 OOD)	256 (1)	4.78	**17.694** (`s0,s1`)	**-0.251** ($-0.36, -0.14$)	11.53 $\rightarrow$ 12.86 $\pm$ 1.13	41.0%/25.1%	48.54%	56.8%/58.9%†

*Per-seed 11-set PPL across `s0..s5` ($n = 6$, deterministic `OLMO_DATA_STEP_OFFSET` stream slices 0K..500K): $G = 1$ (18.210, 17.681, 17.178, 18.845, 17.445, 17.572; `s0..s2` mean 17.690), $G = 4$ (17.873, 17.625, 17.117, 18.189, 17.372, 17.363; `s0..s2` mean 17.538; paired ΔPPL $= -0.34, -0.06, -0.06, -0.66, -0.07, -0.21$, two-sided $p = 0.060$, sign test $p = 0.031$). For $G \in \{16, 64, 256\}$ on `s0..s2`: $G = 16$ (18.067, 17.679, 17.132), $G = 64$ (17.851, 17.644, 17.073), $G = 256$ (17.850, 17.538, `s2` 11-set not run). On `s0, s1, s3`, unclamped $G = 4$ reaches 55.9%–66.0% / 45.7%–61.7% 32K 4/8-N and 66.2%–69.0% / 71.1%–72.9% 64K QA, whereas slow tiles ($\tau_{\max} > 100$K) on `s2, s4, s5` pull unclamped 32K 4/8-N to 25.0%–36.3% / 15.2%–19.5% (full per-seed table in Appendix Table 39c). †For $G = 256$, $N = 128$ 64K QA is 57.3%/59.1%.

that preconditions chunk keys with `Polar-2+`, a five-matmul BF16-safe inverse-square-root operator (0.90 ms F+B on TPU v7x, zero NaNs across 524K tokens; Table 11). Because unwhitened SLA already matches or slightly improves on `SPD-SLA-W` in general language modeling by 60.4B tokens at lower step time, SLA remains unwhitened by default and reserves `SPD-SLA-W` for correlated key-overwrite workloads (60.3 dB BF16 recall SNR; Appendix B) and aggressive `15L:1A` hybrids (Table 25).

Shared-β_t isolation, open-source solver audits, and exact vs. approximate inversion. Four targeted TPU v7x control studies (Appendix H, Tables 28–33) verify our solver attribution and design choices. First, a direct Tokamax input-precision audit confirms that BF16 QKV inputs activate Tokamax's MXU Neumann–Horner path, running 1.68×–2.25× faster in core forward (2.41×–2.58× core F+B; 1.63×–1.71× full-layer F+B) than FP32 sequential substitution while matching outputs within 4.14×10^{-3}–4.19×10^{-3} relative L_2 error; Table 5 already supplied BF16 QKV inputs to Tokamax, so its ratios measure SLA against Tokamax's fast MXU path. Second, against FP64 on real Qwen3.5-0.8B/2B chunks and high-coherence $G = 4$ SLA chunks ($\beta_t \in (0, 2)$), FP32 block-doubling holds 1.4×10^{-7}–3.0×10^{-7} worst-case relative error from $C = 32$ to $C = 256$, whereas BF16 blocked Neumann–Horner reaches 10.6–12.7 relative error at $C = 128$ (16.7%–17.4% of Qwen3.5 chunks above 0.1) and unblocked Neumann products diverge on 21.5%–23.3% of Qwen3.5 chunks and 100% of high-coherence $C \geq 128$ chunks (Table 29). Third, holding the layer shape fixed at $H_k = 2, d_k = 256, H_v = 8, G = 4, C = 128$ (`Exact Inv.`), sharing β_t across the four packed value heads (2 solves vs. 8 solves per chunk) accelerates core forward by 2.45×–3.28×, core F+B by 2.25×–3.00×, and full-layer F+B by 1.87×–2.48× across $D \in \{1792, 2048\}$ and $T \in \{8\text{K}, 32\text{K}\}$ (with 0.0 output difference under equal gates), and speeds up 223M-parameter 3-seed pre-training by 1.68×–1.69× (72.9 vs. 122.4–123.4 ms/step; validation perplexity 74.24 ± 3.56 vs. 72.66 ± 1.45 and 72.89 ± 1.89, beating the per-head control on 2/3 seeds; Table 30). Fourth, once values are packed, exact inversion adds only 4.0%–7.7% to full-LM step time, gives the lowest mean perplexity and tightest seed spread at 204M (69.94 ± 1.15, consistent with the $+0.22$ PPL cost of 1-step Neumann at 12.5B), and compiles cleanly at $C = 256$ once the lowering avoids the Mosaic `kVectorPack` assertion (Table 31). For zero-shot inference on exact-trained Qwen3.5 weights, banded exact 16×16 solvers match exact $N = 1{,}425$ MMLU within 0.00–0.42 points (KL 0.0029–0.0034), whereas matmul-free 1-step operators require healing and still lose 4.6–16.8 MMLU points, keeping exact inversion as SLA's default (Appendix H).

4 Related Work

Decay Granularity in Linear Recurrences. Linear attentions replace the growing softmax KV cache with a fixed-size state updated once per token [14, 37]. Using the decay-group view of Section 2.1, most follow-up models differ in how finely their forget gate is resolved across key channels: RetNet [64], Mamba-2 [16], and GDN [74] use one decay rate per head, whereas GLA [72], RWKV-6/7 [47, 48], KDA [78], and GDN-2 [29] give every key channel its own rate. Mamba [26],

HGRN2 [51], xLSTM [7], and Mamba-3 [38] explore related state and gating structures. SLA takes a middle point with four decay rates per 256-channel head. This is coarse enough that the decay factors out of 64-channel MXU matmuls. In our sweeps it lowers perplexity relative to one rate per head across all six paired 12.5B MoE and dense-LM seeds while trailing on unclamped synthetic 32K needle retrieval (35.4% vs. 59.6%), so one rate per head remains a strong zero-overhead alternative (Section 3.5).

Delta Rules, Higher-Order Memories, and Momentum. DeltaNet [55, 73], test-time-trained memories [9, 10, 65], and the MIRAS online-optimization framework [11] view the state update as online regression with retention regularization, which improves associative recall over decay-only models [5, 34]. Later work adds multiple Householder reflections [58], Householder-diagonalized decay in HDLA [77], negative eigenvalues [24], decoupled per-channel erase and write gates [29], second-order preconditioning [68, 70], log-linear state hierarchies [27], and heavy-ball momentum [10, 46, 49]. Concurrent work, Muir [20], introduces sensitivity-allocated low-rank erase and write gates (reinvesting saved parameters in wider FFNs), a matched-filter readout gate that re-weights slow-decaying large-norm state rows, and a numerically stable large-chunk parallel update tailored to wide matrix units. Muir reports 1.8%–2.0% lower perplexity than GDN-2 across 166M–1.34B unpacked-head models and streaming extrapolation from 4K training up to 10M tokens; we do not run a head-to-head comparison against Muir in this work, and its gate allocation and readout selectivity are complementary to SLA's 4:1 head sharing, tile-grouped decay, and single-kernel VMEM fusion across 6.5B–63.1B MoE models. More broadly, these algorithmic extensions enrich what the state update can express, whereas SLA keeps GDN's single-reflection update and reshapes it around systolic hardware. The two directions are complementary: GDN-2 reaches lower per-token perplexity than SLA in our 15.1B-token MoE runs, and combining GDN-2's decoupled gates with SLA's systolic mapping inside one VMEM-fused kernel (`SLA-R4 + GDN-2`) beats open-source Tokamax KDA at equal chip-hours across $3/3$ paired seeds at both 12.5B and 63.1B scales (Table 9).

Sharing Key Heads and Write Gates Across Value Heads. GQA [2] shares key and value heads across query heads in softmax attention. For linear attention, Kimi Linear [78], Qwen3-Next [53], and `fla` [71] share query-key heads across value heads at $d_k = 128$, while Gated Attention [52] establishes sigmoid output gating; however, these layers keep independent decay and write gates (g, β_t) per value head and still execute H_v chunk solves per chunk. SLA widens keys to $d_k = 256$ and shares both key heads and β_t across R packed value heads (keeping per-value-head normalization), filling whole MXU tiles while cutting chunk solves by $R\times$.

Hardware-Aware Kernels and MXU Chunk Solvers. FlashAttention [15, 17, 76] re-architects softmax attention around the GPU memory hierarchy. For linear recurrences, `fla` [71] provides chunkwise WY kernels [12], TFLA [8] introduces a second level of parallelism inside each chunk so linear RNN and xLSTM kernels scale to large chunk sizes, and FlexLinearAttention [18] compiles a unified three-phase chunk abstraction on GPU tensor cores, while on TPUs and NPUs FLAT [44] fuses Mamba-2 and GLA in Pallas, Sobczyk et al. [59] analyze matmul-rich direct and iterative triangular inversion for delta-rule layers on NPUs, and open-source MaxText and Tokamax [1, 23] implement TPU GDN and KDA kernels that already move individual chunk solves onto the MXU (BF16 8-token Neumann plus Horner doubling with FP32 accumulation in Tokamax, validated at $C = 64$ for per-channel KDA alongside an FP32 sequential fallback and an inference op fusing 1D causal convolution with GDN/KDA state updates, and XLA block-doubling in MaxText's GDN path alongside our block-doubling Pallas KDA baseline). These systems optimize kernels for a fixed recurrence (with Tokamax's convolution fusion targeting single-step inference); SLA also co-designs the training recurrence itself ($R\times$ fewer solves via shared β_t, $G = 4$ tile-grouped decay, and full chunkwise WY forward-backward VMEM fusion with custom-VJP inverse reuse; distinct from works that abbreviate sparse-linear or softmax-linear hybrids as SLA).

Context Parallelism for Linear Recurrences. Head-sharded Ulysses [33] requires at least as many key heads as devices, ring-based methods [40, 50, 62] pass the state through $\mathcal{O}(P)$ sequential hops, and LASP-2 [63] gathers chunk states with one AllGather. Because every gated delta chunk transition with scalar β_t is affine, chunks can also be combined with a parallel prefix scan [13, 43]; we use this standard construction for both SLA and KDA (Section 3.3).

5 Conclusion

Summary. Recent gated-delta models were shaped around GPUs, where unpacked heads, per-channel gate rescaling, and per-value-head chunk solves (including sequential in-block fallbacks) are cheap relative to matmuls. On systolic accelerators, where matrix units outpace vector lanes by two orders of magnitude, kernel tuning or MXU chunk solves alone cannot hide an unpacked, vector-bound recurrence. Systolic Linear Attention (SLA) resolves this mismatch by co-designing the recurrence around the systolic array: sharing 256-wide keys and write gates β_t across value heads fills full tiles and cuts chunk solves by $R\times$, grouping forget gates into four tiles factors decay out of matmuls, and single-kernel VMEM fusion builds on open-source MXU block-doubling solvers [1, 23] to keep exact chunk inversion and its backward pass on-chip. Because value packing already makes exact inversion cheap (4%–8% of step time), exact FP32 block-doubling stays optimal for pre-training, whereas zero-shot serving on exact-trained weights should use exact or banded-exact inversion rather than unhealed 1-step approximations (Appendix H). On TPU v7x, a 2-stage `SLA-R4` layer with exact inversion runs $2.1\times$–$2.6\times$ faster than open-source MaxText KDA and $1.4\times$–$2.4\times$ faster than open-source Tokamax KDA (1-kernel fused `SLA-R4`: $2.5\times$–$2.8\times$ vs. MaxText and $1.5\times$–$2.4\times$ vs. Tokamax under single-chip protocols), with 49%–61% less temporary HBM; building a fused per-channel $G = d_k$ KDA kernel confirms that 4:1 head packing ($1.28\times$–$1.41\times$), VMEM fusion ($1.10\times$–$1.12\times$, which actually regresses to $0.44\times$–$0.69\times$ without packing), and $G = 4$ tile grouping ($1.07\times$–$1.20\times$) each contribute multiplicatively (Appendix I). Across 6.5B–63.1B MoE models on up to 512 chips per run, SLA trains 11%–24% faster than MaxText KDA ($1.08\times$ on the short 5B-token 31.5B `SLA-R2` run; $1.14\times$–$1.24\times$ for `SLA-R4`) and 4%–21% faster end-to-end than open-source Tokamax Mosaic KDA under the unified 4-scale protocol (Appendix Table 20b) at comparable per-token perplexity. Under controlled multi-seed equal-chip-hour pre-training on Dolma 3.5 across 64-chip TPU v7x slices ($n = 3$ paired seeds per arm; Table 9), 1-kernel fused iso-parameter `SLA-R4` ($\mu = 0$) and fused `SLA-R4 + GDN-2` beat safe-clipped open-source Tokamax Mosaic KDA on all 3/3 seeds at both 12.5B (-0.37 ± 0.11 PPL, $t = -6.02, p = 0.027$, and $+0.92 \pm 0.31$ `lm-eval` points, $p = 0.036$, for `SLA-R4`; -0.41 ± 0.35 PPL for `SLA-R4 + GDN-2`) and 63.1B (-1.48 ± 0.21 PPL, $t = -12.08, p = 0.0068$, and -1.41 ± 0.16 PPL, $t = -15.01, p = 0.0044$). Extended TPU v7x control studies (Appendix I) validate our analytical roofline out of sample (18 held-out shapes, zero-parameter Spearman $\rho = 0.893$, 13.9% packed-regime MAPE, +3.5% error on held-out $(R = 8, d_k = 512, C = 256, d_v = 256)$), confirm that public `Kimi-Linear-48B-A3B` checkpoints exhibit the same slow-channel decay tail ($>$ 1M-token channels) and improve under zero-retrain clamping (-0.44 16K tail PPL), evaluate $G = 4$ vs. $G = 1$ across $6/6$ paired seeds on both dense LMs ($t = -5.78$ Val / $t = -7.01$ 16K tail) and 12.5B MoE (-0.23 ± 0.24 PPL, $6/6$ seed wins, sign-test $p = 0.031$), and establish $\mu = 0$ as SLA's recommended default with an optional Schur-stable guard ($5 \to 0$ unstable grid points).

Limitations and Code Availability. SLA targets accelerators with wide systolic tiles and a large matrix-to-vector throughput gap; on GPUs, smaller Tensor Core tiles and faster vector lanes already saturate narrow states. Several empirical boundaries warrant transparency: (1) our analytical roofline achieves 13.9% calibrated MAPE inside the packed regime ($R \geq 4, d_k \geq 256$), whereas unpacked $R = 1$ shapes run 42%–53% slower than the single-η packed extrapolation due to unmodeled VPU and scheduling overhead; (2) at fixed token count on 12.5B MoE, BF16 routing variance ($\approx$0.18 PPL) yields per-token parity rather than a large per-token gap (though $G = 4$ beats $G = 1$ on $6/6$ paired 12.5B seeds by -0.23 ± 0.24 PPL while $G = 1$ leads synthetic 32K passkey retrieval by $+24.3$ points, and on synthetic 2-layer multi-timescale associative recall $G = 256$ outperforms $G \in \{1, 2, 4\}$; Tables 10 and 37), and the 120.8B-token 512-chip 63.1B horizon runs remain single-seed; (3) zero-retraining gate clamping closes much of KDA's 8K $\to$ 64K gap (and with `SSMax`, clamped 63.1B KDA beats `SLA-R4` on 64K QA while trailing on 32K 8-Needle); (4) in our multi-seed equal-chip-hour comparisons against open-source Tokamax Mosaic KDA (Tables 9, 38, and 39), Tokamax `0.0.14` required clipping $g \geq -3.5$ ($f \geq 0.030$) to prevent float32 `exp2` backward overflow, which bounds ultra-fast forgetting on the baseline, and two 12.5B `SLA-R4 + GDN-2` seeds plus late-stage 63.1B 64K RULER jobs exited after saving all in-loop validation checkpoints; and (5) chunk-boundary momentum introduces chunk-size dependence, doubles context-parallel transition size to 512×512, requires the 64 MiB scoped-VMEM cap at $D = 1792$, and showed no isolated perplexity gain over $\mu = 0$ on 3 paired dense-LM seeds, so $\mu = 0$ is our recommended default while historical 12.5B/63.1B MoE runs used $\mu = 0.25$. Pallas kernels, verification tests, and evaluation scripts will be released publicly upon publication to enable full reproduction on TPU v7x.

Future Work. Extending our multi-seed 12.5B and 63.1B equal-chip-hour `SLA-R4` and `SLA-R4 + GDN-2` results to full 100B$+$-token horizons and integrating concurrent sensitivity-allocated gating designs such as Muir [20] onto SLA's systolic head-packing and tile-grouped mapping remain direct next steps. Scaling multi-seed long-context evaluations to larger 64K–128K benchmarks and testing exact block-doubling at $C = 256$ in full MoE pre-training runs also remain open directions.

References

[1] AI-Hypercomputer Contributors. MaxText: A high performance, highly scalable, open-source LLM library written in pure Python/JAX. https://github.com/AI-Hypercomputer/maxtext, 2026.

[2] J. Ainslie, J. Lee-Thorp, M. de Jong, Y. Zemlyanskiy, F. Lebrón, and S. Sanghai. GQA: Training generalized multi-query transformer models from multi-head checkpoints. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2023. arXiv:2305.13245.

[3] N. Amsel, D. Persson, C. Musco, and R. M. Gower. The Polar Express: Optimal matrix sign methods and their application to the Muon algorithm. *arXiv preprint arXiv:2505.16932*, 2025.

[4] R. Anil, V. Gupta, T. Koren, K. Regan, and Y. Singer. Scalable second order optimization for deep learning. *arXiv preprint arXiv:2002.09018*, 2020.

[5] S. Arora, S. Eyuboglu, A. Timalsina, I. Johnson, M. Poli, J. Zou, A. Rudra, and C. Ré. Zoology: Measuring and improving recall in efficient language models. In *International Conference on Learning Representations (ICLR)*, 2024.

[6] J. Austin, S. Douglas, R. Frostig, A. Levskaya, C. Chen, S. Vikram, F. Lebron, P. Choy, V. Ramasesh, A. Webson, and R. Pope. How to scale your model. https://jax-ml.github.io/scaling-book/, 2025.

[7] M. Beck, K. Pöppel, M. Spanring, A. Auer, O. Prudnikova, M. Kopp, G. Klambauer, J. Brandstetter, and S. Hochreiter. xLSTM: Extended long short-term memory. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.

[8] M. Beck, K. Pöppel, P. Lippe, and S. Hochreiter. Tiled flash linear attention: More efficient linear RNN and xLSTM kernels. *arXiv preprint arXiv:2503.14376*, 2025.

[9] A. Behrouz, Z. Li, P. Kacham, M. Daliri, Y. Deng, P. Zhong, M. Razaviyayn, and V. Mirrokni. Atlas: Learning to optimally memorize the context at test time. *arXiv preprint arXiv:2505.23735*, 2025.

[10] A. Behrouz, P. Zhong, and V. Mirrokni. Titans: Learning to memorize at test time. *arXiv preprint arXiv:2501.00663*, 2025.

[11] A. Behrouz, M. Razaviyayn, P. Zhong, and V. Mirrokni. It's All Connected: A Journey Through Test-Time Memorization, Attentional Bias, Retention, and Online Optimization. In *International Conference on Learning Representations (ICLR)*, 2026. arXiv:2504.13173.

[12] C. Bischof and C. Van Loan. The WY representation for products of Householder matrices. *SIAM Journal on Scientific and Statistical Computing*, 8(1):s2–s13, 1987.

[13] G. E. Blelloch. Prefix sums and their applications. Technical Report CMU-CS-90-190, Carnegie Mellon University, 1990.

[14] K. Choromanski, V. Likhosherstov, D. Dohan, X. Song, A. Gane, T. Sarlos, P. Hawkins, J. Davis, A. Mohiuddin, L. Kaiser, D. Belanger, L. Colwell, and A. Weller. Rethinking attention with performers. In *International Conference on Learning Representations (ICLR)*, 2021.

[15] T. Dao. FlashAttention-2: Faster attention with better parallelism and work partitioning. In *International Conference on Learning Representations (ICLR)*, 2024.

[16] T. Dao and A. Gu. Transformers are SSMs: Generalized models and efficient algorithms through structured state space duality. In *International Conference on Machine Learning (ICML)*, 2024.

[17] T. Dao, D. Y. Fu, S. Ermon, A. Rudra, and C. Ré. FlashAttention: Fast and memory-efficient exact attention with IO-awareness. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

[18] H. Duanmu, S. Zheng, N. Zheng, J. Lu, X. Zheng, X. Zhang, L.-W. Chang, X. Liu, and D. Lin. FlexLinearAttention: Compiling a Unified Abstraction into Scalable Kernels for Linear Attention. In *International Conference on Learning Representations (ICLR)*, 2026.

[19] W. Fedus, B. Zoph, and N. Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research (JMLR)*, 23(120): 1–39, 2022. arXiv:2101.03961.

[20] T. Gangavarapu, G. Tabak, and C. J. S. Schaefer. Muir: Linear Attention with Sensitivity-Allocated Gating and Matched-Filter Readout. *arXiv preprint arXiv:2611.08376*, 2026.

[21] L. Gao, J. Tow, B. Abbasi, S. Biderman, S. Black, A. DiPofi, C. Foster, L. Golding, J. Hsu, A. L. Noac'h, H. Li, K. McDonell, N. Muennighoff, C. Ociepa, J. Phang, L. Reynolds, H. Schoelkopf, A. Skowron, L. Sutawika, E. Tang, A. Thite, B. Wang, K. Wang, and A. Zou. A framework for few-shot language model evaluation. Zenodo, https://doi.org/10.5281/zenodo.10256836, 2023.

[22] Google Cloud. TPU7x (Ironwood) architecture and performance optimizations. https://docs.cloud.google.com/tpu/docs/ironwood-performance, 2025.

[23] Google DeepMind. Tokamax: Open-source accelerator kernel library for JAX. https://github.com/openxla/tokamax, 2026.

[24] R. Grazzi, J. Siems, J. K. H. Franke, A. Zela, F. Hutter, and M. Pontil. Unlocking state-tracking in linear RNNs through negative eigenvalues. In *International Conference on Learning Representations (ICLR)*, 2025. arXiv:2411.12537.

[25] E. Grishina, M. Smirnov, and M. Rakhuba. Accelerating Newton-Schulz iteration for orthogonalization via Chebyshev-type polynomials. *arXiv preprint arXiv:2506.10935*, 2025.

[26] A. Gu and T. Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2023.

[27] H. Guo, S. Yang, T. Goel, E. P. Xing, T. Dao, and Y. Kim. Log-linear attention. *arXiv preprint arXiv:2506.04761*, 2025.

[28] V. Gupta, T. Koren, and Y. Singer. Shampoo: Preconditioned stochastic tensor optimization. In *International Conference on Machine Learning (ICML)*, 2018.

[29] A. Hatamizadeh, Y. Choi, and J. Kautz. Gated DeltaNet-2: Decoupling erase and write in linear attention. *arXiv preprint arXiv:2605.22791*, 2026.

[30] N. J. Higham. *Functions of Matrices: Theory and Computation*. Society for Industrial and Applied Mathematics (SIAM), 2008.

[31] J. Hoffmann, S. Borgeaud, A. Mensch, E. Buchatskaya, T. Cai, E. Rutherford, D. de Las Casas, L. A. Hendricks, J. Welbl, A. Clark, T. Hennigan, E. Noland, K. Millican, G. van den Driessche, B. Damoc, A. Guy, S. Osindero, K. Simonyan, E. Elsen, J. W. Rae, O. Vinyals, and L. Sifre. Training compute-optimal large language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. arXiv:2203.15556.

[32] C.-P. Hsieh, S. Sun, S. Kriman, S. Acharya, D. Rekesh, F. Jia, Y. Zhang, and B. Ginsburg. RULER: What's the real context size of your long-context language models? In *Conference on Language Modeling (COLM)*, 2024.

[33] S. A. Jacobs, M. Tanaka, C. Zhang, M. Zhang, S. L. Song, S. Rajbhandari, and Y. He. DeepSpeed Ulysses: System optimizations for enabling training of extreme long sequence transformer models. *arXiv preprint arXiv:2309.14509*, 2023.

[34] S. Jelassi, D. Brandfonbrener, S. M. Kakade, and E. Malach. Repeat after me: Transformers are better than state space models at copying. In *International Conference on Machine Learning (ICML)*, 2024.

[35] N. P. Jouppi, C. Young, N. Patil, D. Patterson, G. Agrawal, R. Bajwa, S. Bates, S. Bhatia, N. Boden, A. Borchers, R. Boyle, P.-L. Cantin, C. Chao, C. Clark, J. Coriell, M. Daley, M. Dau, J. Dean, B. Gelb, T. V. Ghaemmaghami, R. Gottipati, W. Gulland, R. Hagmann, C. R. Ho, D. Hogberg, J. Hu, R. Hundt, D. Hurt, J. Ibarz, A. Jaffey, A. Jaworski, A. Kaplan, H. Khaitan, A. Koch, N. Kumar, S. Lacy, J. Laudon, J. Law, D. Le, C. Leary, Z. Liu, K. Lucke,

A. Lundin, G. MacKean, A. Maggiore, M. Mahony, K. Miller, R. Nagarajan, R. Narayanaswami, R. Ni, K. Nix, T. Norrie, M. Omernick, N. Penukonda, A. Phelps, J. Ross, M. Ross, A. Salek, E. Samadiani, C. Severn, G. Sizikov, M. Snelham, J. Souter, D. Steinberg, A. Swing, M. Tan, G. Thorson, B. Tian, H. Toma, E. Tuttle, V. Vasudevan, R. Walter, W. Wang, E. Wilcox, and D. H. Yoon. In-datacenter performance analysis of a tensor processing unit. In *International Symposium on Computer Architecture (ISCA)*, 2017.

[36] N. P. Jouppi, G. Kurian, S. Li, P. Ma, R. Nagarajan, L. Nai, N. Patil, S. Subramanian, A. Swing, B. Towles, C. Young, X. Zhou, Z. Zhou, and D. Patterson. TPU v4: An optically reconfigurable supercomputer for machine learning with hardware support for embeddings. In *International Symposium on Computer Architecture (ISCA)*, 2023.

[37] A. Katharopoulos, A. Vyas, N. Pappas, and F. Fleuret. Transformers are RNNs: Fast autoregressive transformers with linear attention. In *International Conference on Machine Learning (ICML)*, 2020.

[38] A. Lahoti, K. Y. Li, B. Chen, C. Wang, A. Bick, J. Z. Kolter, T. Dao, and A. Gu. Mamba-3: Improved sequence modeling using state space principles. In *International Conference on Learning Representations (ICLR)*, 2026. arXiv:2603.15569.

[39] Y. Leviathan, M. Kalman, and Y. Matias. Fast inference from transformers via speculative decoding. In *International Conference on Machine Learning (ICML)*, 2023. arXiv:2211.17192.

[40] H. Liu, M. Zaharia, and P. Abbeel. Ring attention with blockwise transformers for near-infinite context. In *International Conference on Learning Representations (ICLR)*, 2024.

[41] J. Liu, J. Su, X. Yao, Z. Jiang, G. Lai, Y. Du, Y. Qin, W. Xu, E. Lu, J. Yan, Y. Chen, H. Zheng, Y. Liu, S. Liu, B. Yin, W. He, H. Zhu, Y. Wang, J. Wang, M. Dong, Z. Zhang, Y. Kang, H. Zhang, X. Xu, Y. Zhang, Y. Wu, X. Zhou, and Z. Yang. Muon is scalable for LLM training. *arXiv preprint arXiv:2502.16982*, 2025.

[42] I. Loshchilov and F. Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)*, 2019. arXiv:1711.05101.

[43] E. Martin and C. Cundy. Parallelizing linear recurrent neural nets over sequence length. In *International Conference on Learning Representations (ICLR)*, 2018.

[44] M. Mishra, Z. Han, P. Liang, I. Stoica, T. Dao, R. Iyer, and J. E. Gonzalez. FLAT: Flash Linear Attention on TPUs. https://trainstation-project.github.io/blog/flat/, Oct. 2026.

[45] K. M. Nakanishi. Scalable-softmax is superior for attention. *arXiv preprint arXiv:2501.19399*, 2025.

[46] T. M. Nguyen, R. G. Baraniuk, R. M. Kirby, S. J. Osher, and B. Wang. Momentum transformer: Closing the performance gap between self-attention and its linearization. *arXiv preprint arXiv:2208.00579*, 2022.

[47] B. Peng, D. Goldstein, Q. Anthony, A. Albalak, E. Alcaide, S. Biderman, E. Cheah, X. Du, T. Ferdinan, H. Hou, P. Kazienko, K. K. GV, J. Kocoń, B. Koptyra, S. Krishna, R. McClelland, Jr., J. Lin, N. Muennighoff, F. Obeid, A. Saito, G. Song, H. Tu, C. Wirawan, S. Woźniak, R. Zhang, B. Zhao, Q. Zhao, P. Zhou, J. Zhu, and R.-J. Zhu. Eagle and finch: RWKV with matrix-valued states and dynamic recurrence. In *Conference on Language Modeling (COLM)*, 2024.

[48] B. Peng, R. Zhang, D. Goldstein, E. Alcaide, X. Du, H. Hou, J. Lin, J. Liu, J. Lu, W. Merrill, G. Song, K. Tan, S. Utpala, N. Wilce, J. S. Wind, T. Wu, D. Wuttke, and C. Zhou-Zheng. RWKV-7 "goose" with expressive dynamic state evolution. *arXiv preprint arXiv:2503.14456*, 2025.

[49] B. T. Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5):1–17, 1964.

[50] Z. Qin, W. Sun, D. Li, X. Shen, W. Sun, and Y. Zhong. Various lengths, constant speed: Efficient language modeling with lightning attention. In *International Conference on Machine Learning (ICML)*, 2024. arXiv:2405.17381.

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

[52] Z. Qiu, Z. Wang, B. Zheng, Z. Huang, K. Wen, S. Yang, R. Men, L. Yu, F. Huang, S. Huang, D. Liu, J. Zhou, and J. Lin. Gated attention for large language models: Non-linearity, sparsity, and attention-sink-free. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. arXiv:2505.06708.

[53] Qwen Team. Qwen3-Next: Towards ultimate training and inference efficiency. https://qwenlm.github.io/blog/qwen3-next/, 2025.

[54] P. Rajpurkar, J. Zhang, K. Lopyrev, and P. Liang. SQuAD: 100,000+ questions for machine comprehension of text. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2016. arXiv:1606.05250.

[55] I. Schlag, K. Irie, and J. Schmidhuber. Linear transformers are secretly fast weight programmers. In *International Conference on Machine Learning (ICML)*, 2021.

[56] N. Shazeer. GLU variants improve transformer. *arXiv preprint arXiv:2002.05202*, 2020.

[57] N. Shazeer, A. Mirhoseini, K. Maziarz, A. Davis, Q. Le, G. Hinton, and J. Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *International Conference on Learning Representations (ICLR)*, 2017. arXiv:1701.03538.

[58] J. Siems, T. Carstensen, A. Zela, F. Hutter, M. Pontil, and R. Grazzi. DeltaProduct: Improving state-tracking in linear RNNs via Householder products. In *International Conference on Learning Representations (ICLR)*, 2026. arXiv:2502.10297.

[59] A. Sobczyk, G. Gottardo, C. K. Matzoros, M. D. Vita, F. Skogh, A. Zouzias, and J. Zhuang. Fast and stable triangular inversion for delta-rule linear transformers. *arXiv preprint arXiv:2605.21325*, 2026.

[60] L. Soldaini, R. Kinney, A. Bhagia, D. Schwenk, D. Atkinson, R. Authur, B. Bogin, K. Chandu, J. Dumas, Y. Elazar, V. Hofmann, A. H. Jha, S. Kumar, L. Lucy, X. Lyu, N. Lambert, I. Magnusson, J. Morrison, N. Muennighoff, A. Naik, C. Nam, M. E. Peters, A. Ravichander, K. Richardson, Z. Shen, E. Strubell, N. Subramani, O. Tafjord, P. Walsh, L. Zettlemoyer, N. A. Smith, H. Hajishirzi, I. Beltagy, D. Groeneveld, J. Dodge, and K. Lo. Dolma: an open corpus of three trillion tokens for language model pretraining research. In *Association for Computational Linguistics (ACL)*, 2024.

[61] J. Su, M. Ahmed, Y. Lu, S. Pan, W. Bo, and Y. Liu. RoFormer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024. arXiv:2104.09864.

[62] W. Sun, Z. Qin, D. Li, X. Shen, Y. Qiao, and Y. Zhong. Linear attention sequence parallelism. *arXiv preprint arXiv:2404.02882*, 2024.

[63] W. Sun, Z. Qin, D. Li, X. Shen, Y. Qiao, and Y. Zhong. LASP-2: Rethinking sequence parallelism for linear attention and its hybrid. *arXiv preprint arXiv:2502.07563*, 2025.

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

[65] Y. Sun, X. Li, K. Dalal, J. Xu, A. Vikram, G. Zhang, Y. Dubois, X. Chen, X. Wang, S. Koyejo, T. Hashimoto, and C. Guestrin. Learning to (learn at test time): RNNs with expressive hidden states. In *International Conference on Machine Learning (ICML)*, 2024. arXiv:2407.04620.

[66] Team Olmo. Olmo 3. *arXiv preprint arXiv:2512.13961*, 2025.

[67] Team OLMo, P. Walsh, L. Soldaini, D. Groeneveld, K. Lo, S. Arora, A. Bhagia, Y. Gu, S. Huang, M. Jordan, N. Lambert, D. Schwenk, O. Tafjord, T. Anderson, D. Atkinson, F. Brahman, C. Clark, P. Dasigi, N. Dziri, A. Ettinger, M. Guerquin, D. Heineman, H. Ivison, P. W. Koh, J. Liu, S. Malik, W. Merrill, L. J. V. Miranda, J. Morrison, T. Murray, C. Nam, J. Poznanski, V. Pyatkin, A. Rangapur, M. Schmitz, S. Skjonsberg, D. Wadden, C. Wilhelm, M. Wilson, L. Zettlemoyer, A. Farhadi, N. A. Smith, and H. Hajishirzi. 2 OLMo 2 furious. *arXiv preprint arXiv:2501.00656*, 2025.

[68] N. Tumma, N. Loo, and D. Rus. Preconditioned DeltaNet: Curvature-aware sequence modeling for linear recurrences. *arXiv preprint arXiv:2604.21100*, 2026.

[69] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

[70] J. von Oswald, N. Scherrer, S. Kobayashi, L. Versari, S. Yang, S. Mittal, M. Schlegel, K. Maile, Y. Schimpf, O. Sieberling, A. Meulemans, R. A. Saurous, G. Lajoie, C. Frenkel, R. Pascanu, B. A. y Arcas, and J. Sacramento. MesaNet: Sequence modeling by locally optimal test-time training. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. arXiv:2506.05233.

[71] S. Yang and Y. Zhang. FLA: A Triton-based library for hardware-efficient implementations of linear attention mechanism. https://github.com/fla-org/flash-linear-attention, 2024.

[72] S. Yang, B. Wang, Y. Shen, R. Panda, and Y. Kim. Gated linear attention transformers with hardware-efficient training. In *International Conference on Machine Learning (ICML)*, 2024.

[73] S. Yang, B. Wang, Y. Zhang, Y. Shen, and Y. Kim. Parallelizing linear transformers with the delta rule over sequence length. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.

[74] S. Yang, J. Kautz, and A. Hatamizadeh. Gated Delta networks: Improving Mamba2 with Delta rule. In *International Conference on Learning Representations (ICLR)*, 2025. arXiv:2412.06464.

[75] Z. Yang, P. Qi, S. Zhang, Y. Bengio, W. W. Cohen, R. Salakhutdinov, and C. D. Manning. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2018. arXiv:1809.09600.

[76] T. Zadouri, M. Hoehnerbach, J. Shah, T. Liu, V. Thakkar, and T. Dao. FlashAttention-4: Algorithm and kernel pipelining co-design for asymmetric hardware scaling. *arXiv preprint arXiv:2603.05451*, 2026.

[77] J. Zhang, Z. Qin, J. Tong, S. Mei, J. Hu, Y. Pan, A. Hu, M. Yao, B. Xu, and G. Li. Householder-Diagonalized Linear Attention (HDLA): Utilizing Enhanced Decay Mechanism for Efficient Sequence Modeling. In *International Conference on Learning Representations (ICLR)*, 2026.

[78] Y. Zhang, Z. Lin, X. Yao, J. Hu, F. Meng, C. Liu, X. Men, S. Yang, Z. Li, Z. Yang, Y. Du, and Kimi Team. Kimi Linear: An expressive, efficient attention architecture. *arXiv preprint arXiv:2510.26692*, 2025.

A Conditioning Recurrent State in Low Precision: Oblique Whitening and `Polar-2+`

This appendix details `SPD-SLA-W`, an optional key-whitening extension for aggressive `15L:1A` hybrids summarized at the end of Section 3.5 (not used in the core SLA benchmarks of Section 3). Reducing softmax frequency from `7L:1A` (12.5%) to `15L:1A` (6.25%) cuts 131K cache+state HBM to 16.47 GiB (Table 6), which forces middle linear layers (`L5, L8, L11`) to hold 10×–35× more active key-value associations N_{eff}: the memory load $\rho_{\text{eff}} = N_{\text{eff}}/d_k$ rises to $[0.101, 0.152]$ versus ≤ 0.015 at `4L:1A`, while chunk key Gram matrices become poorly conditioned ($\kappa_2(I_C + K_c K_c^\top) \in [20.6, 62.9]$; Figure 8c). Whereas GPU second-order rules rely on 30-step Conjugate Gradient (CG) [70] or axis-aligned diagonal scaling [68], a 256×256 MXU executes five $[128, 128]$ matmuls across 64 chunks in 0.62 ms forward (0.90 ms F+B). Below we analyze when key whitening improves BF16 recall and introduce **Polar-2+**, a five-matmul BF16-safe inverse-square-root primitive.

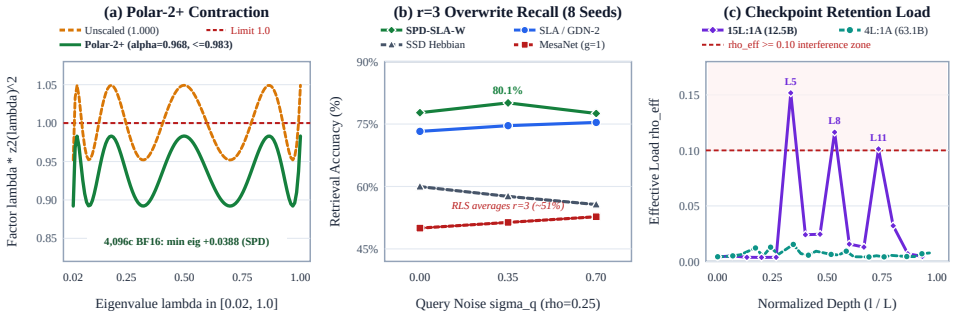

Figure 8: **Polar-2+ minimax contraction, overwrite recall, and checkpoint retention audit. (a)** Scalar inverse-square-root contraction $\lambda z_2(\lambda)^2$ on normalized eigenvalue $\lambda \in [0.02, 1.0]$ (Proposition A.2): unscaled polynomial factor $\alpha = 1.000$ hits $1.0492 > 1$, whereas `Polar-2+` (safety factor $\alpha_{\text{safe}} = 0.968$) bounds the peak to $0.9831 < 1$, keeping minimum eigenvalue $\lambda_{\min}(\Sigma_c) \geq 0.0388$ of key preconditioner Σ_c across 4,096 BF16 chunks (524K tok). **(b)** 8-seed Hadamard overwrite accuracy (overwrite multiplicity $r = 3$) at memory load $\rho_{\text{eff}} = N/d_k = 0.25$ ($N = 16$ stored keys, $d_k = 64$). **(c)** Trained 63.1B (`4L:1A`) and 12.5B (`15L:1A`) checkpoint load ρ_{eff}.

Proposition A.1 (Oblique Whitening vs. Euclidean & RLS)**.** *Let $K \in \mathbb{R}^{N \times d_k}$ stack $N \leq d_k$ unit keys $k_1, \ldots, k_N \in \mathbb{R}^{d_k}$ with mutual coherence $\nu_{coh} = \max_{i \neq j} |\langle k_i, k_j \rangle|$ and Gram matrix $\mathcal{G} = KK^\top$ satisfying $\kappa_2(\mathcal{G}) > 1$ (proofs in Section A.1).*

1. [Prior Delta [68, 73]] *For keys with coherence $\nu_{coh} > 0$, single-pass Euclidean delta lets each later write perturb the readout of earlier keys in proportion to their key overlap, and diagonal column-energy preconditioning $\text{diag}(p(K))^{-1/2} k_t$ [68] (where $p(K) \in \mathbb{R}^{d_k}$ has m-th entry $p_m(K) = \sum_{i=1}^{N} K_{i,m}^2$) fails under orthogonal key rotation $K \mapsto KQ$ ($Q \in \text{O}(d_k)$).*

2. [Prior RLS [70]] *MesaNet (with recursive least-squares key covariance $H_t = \gamma H_{t-1} + k_t k_t^\top$ and scalar forget factor $\gamma \in (0, 1]$) averages $r \geq 2$ overwrites $v^{(1)}, \ldots, v^{(r)} \in \mathbb{R}^{d_v}$ of the same key as $\frac{1}{r} \sum_{m=1}^{r} v^{(m)}$ at $\gamma = 1$, and suppresses stale values at $\gamma < 1$ only by decaying write-once keys stored $t - j$ steps earlier (at step $j < t$) at rate γ^{t-j}.*

3. ***[Ours]** Oblique whitened keys $\tilde{K}_c \in \mathbb{R}^{C \times d_k}$ and causal oblique keys $\tilde{k}_t \in \mathbb{R}^{d_k}$ (Eq. (14)) inside Eq. (4) decorrelate intra-chunk writes while erasing selectively along $\tilde{k}_t$.*

Non-circular normalization and `Polar-2+` formulation. Matrix-sign polynomials (Polar Express [3], CANS [25]) collapse singular values to 1, while coupled `NS-5` [4, 28, 30, 41] uses 10 $[C, C]$ matmuls [Prior Polynomials]. Here $K_c \in \mathbb{R}^{C \times d_k}$ stacks the chunk's keys, $\Sigma_{c-1} \succ 0$ is the $d_k \times d_k$ key preconditioner carried across chunks (Eq. (15); initialized at $\Sigma_0 = I_{d_k}$ and relaxed toward I_{d_k} at rate $\rho_p \in [0, 1]$), $\eta_c \in (0, 1]$ is the whitening step size, and $\varepsilon_\beta > 0$ is the erasure safety margin (Section 2.3). To avoid circular regularization, we form the raw chunk Gram matrix $\mathcal{G}_0 \in \mathbb{R}^{C \times C}$, its

Frobenius normalization $\bar{\mathcal{G}}_0$, spectral upper bound $\hat{u}_0$, and adaptive ridge ε_{reg}:

$$\mathcal{G}_0 = I_C + \eta_c K_c \Sigma_{c-1} K_c^\top, \qquad \bar{\mathcal{G}}_0 = \frac{\mathcal{G}_0}{\|\mathcal{G}_0\|_F}, \qquad \hat{u}_0 = \max\big(\|\mathcal{G}_0\|_F \left\|\bar{\mathcal{G}}_0^2\right\|_F^{1/2}, 1\big) \geq \lambda_{\max}(\mathcal{G}_0),$$

and

$$\varepsilon_{\text{reg}} = \max\left(0.02,\ \frac{0.0205}{0.9795}\hat{u}_0 - 1\right).$$

Setting regularized Gram $\mathcal{G}_{\text{reg}} = \mathcal{G}_0 + \varepsilon_{\text{reg}} I_C$ and scale $u_0 = \hat{u}_0 + \varepsilon_{\text{reg}}$ maps the eigenvalues of $T_0 = \mathcal{G}_{\text{reg}}/u_0$ into $[0.02, 1.0]$ while reusing $\bar{\mathcal{G}}_0^2$ (T_0^2 adds zero matmuls; 5 matmuls total), with step-size clamp

$$\beta_t \leq \frac{2 - \varepsilon_\beta}{\max(k_t^\top \tilde{k}_t,\ \varepsilon_\beta)}.$$

Letting $K_{<t}$ stack the chunk keys strictly preceding token t, $\mathcal{M}_\ell \in \mathbb{R}^{C\times C}$ denote the degree-2 polynomial factors for stages $\ell \in \{0, 1\}$ with $T_1 = \mathcal{M}_0 T_0 \mathcal{M}_0$ and minimax coefficients $(a_0, b_0, e_0) = (4.92864, -11.29513, 7.70099)$ and $(a_1, b_1, e_1) = (2.00410, -1.38700, 0.38894)$, and $\mathcal{W}_{\text{isq}} \in \mathbb{R}^{C\times C}$ denote the safety-scaled ($\alpha_{\text{safe}} = 0.968$) inverse-square-root factor, **Polar-2+** computes:

Low-Precision Conditioning: Oblique Whitening & BF16-Safe `Polar-2+`

$$\tilde{K}_c = \sqrt{\eta_c}\,\mathcal{W}_{\text{isq}} K_c \Sigma_{c-1}, \qquad \tilde{k}_t \propto \big(I_{d_k} - \eta_c \Sigma_{c-1} K_{<t}^\top K_{<t}\big)\Sigma_{c-1} k_t, \tag{14}$$

$$\Sigma_c = (1-\rho_p)\big(\Sigma_{c-1} - \tilde{K}_c^\top \tilde{K}_c\big) + \rho_p I_{d_k}, \tag{15}$$

$$\mathcal{M}_\ell = a_\ell I_C + b_\ell T_\ell + e_\ell T_\ell^2, \qquad \mathcal{W}_{\text{isq}} = \alpha_{\text{safe}} u_0^{-1/2} \mathcal{M}_1 \mathcal{M}_0. \tag{16}$$

The symmetric factor $\tilde{K}_c$ runs in the standalone chunk-whitening benchmarks (Table 11), the 4,096-chunk ($\rho_p = 0$) BF16 Woodbury stress test, and the lagged block-32 causal variant (`polar2_causal_blk32`, Appendix C). Because dense $\mathcal{W}_{\text{isq}}$ mixes tokens within a chunk, packed autoregressive LM training instead uses the strictly causal Schur-oblique key $\tilde{k}_t$ in Eq. (14) (normalized to $k_t^\top \tilde{k}_t = 1$) inside the chunk and runs `Polar-2+` in the chunk-boundary Woodbury tail. Setting $\rho_p = 1$ resets $\Sigma_c = I_{d_k}$ for compatibility with `SLA Tree-CP` (unscaled $\alpha = 1.000$ still diverges at step 299 through the tail gradients); carrying Σ_c across chunks ($\rho_p = 0.02$) replaces the additive update in Eq. (15), which hits `NaN` at step 282 in BF16, with the multiplicative Sylvester congruence $\mathcal{B}_c \Sigma_{c-1} \mathcal{B}_c^\top$ ($\mathcal{B}_c \in \mathbb{R}^{d_k \times d_k}$; Section A.1).

Proposition A.2 (Exact Equioscillation and Safety Margin of Polar-2+)**.** *Writing $m_\ell(x) = a_\ell + b_\ell x + e_\ell x^2$, let*

$$z_2(\lambda) = m_1\big(\lambda\, m_0(\lambda)^2\big)\, m_0(\lambda)$$

be the scalar map that Eq. (16) applies to an eigenvalue λ of T_0 (since $T_1 = \mathcal{M}_0 T_0 \mathcal{M}_0$ has eigenvalue $\lambda\, m_0(\lambda)^2$), and let $\kappa_{post} = \kappa_2(\mathcal{W}_{isq}\mathcal{G}_{reg}\mathcal{W}_{isq})$ be the post-whitening condition number. On $\lambda \in [0.02, 1.0]$, z_2 equioscillates with $\lambda z_2(\lambda)^2 \in [0.9520, 1.0492]$, exceeding 1 for unscaled $\alpha = 1$ and driving recursive Woodbury indefinite. Scaling by $\alpha_{safe} = 0.968$ bounds exact $\lambda(\alpha_{safe} z_2(\lambda))^2 \in [0.8920, 0.9831]$ (1.69% BF16 contractive margin below 1; exact $\kappa_{post} \leq 1.102$, BF16 5-matmul $\kappa_{post} \leq 1.123$).

When whitening helps vs. unwhitened SLA. `Polar-2+` runs 2.66× faster than Panel-16 Cholesky and 12.4×–30.8× faster than MesaNet `CG-10/30` (0.897 ms F+B; Table 11), keeping 4,096-chunk (524K-token) BF16 Woodbury SPD ($\lambda_{\min} \geq 0.0388$). It trains 2-seed 12.5B `15L:1A` with zero NaNs and matches 10-matmul `NS-5` (`s1`: 17.936 vs. 17.970 PPL), whereas $\alpha = 1.000$ diverges at step 299 and BF16 Cholesky hits NaN at step 54 (Appendix C). Because its language-modeling benefit shrinks to 0.004 nats by 15B tokens and unwhitened $G = 4$ SLA already saturates PPL at 60.4B `15L:1A+MS` (15.443 vs. 15.624 at +3.7% step time), unwhitened SLA is our default mixer.

`SPD-SLA-W` serves three high-collision regimes:

Table 11: Measured TPU v7x chunk-whitening latency (64 chunks = 8,192 tokens, $C = 128$, $H_k = 4, d_k = 256$; MesaNet `CG-10`: 1.74/11.16 ms), exact conditioned Gram κ_{post} (≤ 1.123 in BF16), and 4,096-chunk (524K-token, $\rho_p = 0$) BF16 Woodbury stress-test $\lambda_{\min}$.

Whitening / Second-Order Solver	Ops	Fwd (ms)	F+B (ms)	κ_{post}	$\lambda_{\min}$ (4,096c BF16)
MesaNet Gated `CG-30`	VPU+MXU	3.050	27.640	1.02	– (CG)
Panel-16 Cholesky (VPU)	Tri. Solve	1.412	2.385	1.00	+0.0148 (SPD)
Coupled Newton-Schulz (`NS-5`)	10 MM	0.708	1.024	1.04	+0.0294 (SPD)
Polar-2 Unscaled ($\alpha = 1.000$)	5 MM	0.615	0.895	1.102	−0.0412 (Indef.)
Polar-2+ Safe ($\alpha = 0.968$)	**5** MM	0.617	0.897	1.102	+0.0388 **(SPD)**

1. correlated $r \geq 2$ overwrite ($\kappa_2(\mathcal{G}) \approx 48$; Figure 8b, Table 12), reaching 60.3 dB BF16 SNR and 80.1% noisy $r = 3$ recall vs. 2.67 dB and 50.0% for MesaNet;
2. uniform-decay 12.5B `15L:1A` (15.1B tok, Table 25), edging `15L SLA-R2` on 32K 8-N (39.8% vs. 34.8%) and $N = 512$ 64K `HotpotQA` (69.5% vs. 68.2%); and
3. 63.1B `15L:1A`, reaching 12.42 vs. 13.01 32K Tail PPL at 5.0B tok (Table 25) and 11.82 Tail PPL with 57.8%/64.8% 64K `SQuAD/HotpotQA` at 30.2B ($+12.5/+25.0$ pts over `0L:1A`; Table 26).

A.1 Proofs of Propositions A.1 and A.2

Part 1 (Euclidean Interference and Diagonal Non-Invariance). Let unit keys $k_1, \dots, k_N \in \mathbb{R}^{d_k}$ satisfy $|\langle k_i, k_j \rangle| = \nu_{\text{coh}}$ for $i \neq j$ (coherence ν_{coh}; e.g., shared-subspace correlated dictionary) and draw independent zero-mean values $v_j \in \mathbb{R}^{d_v}$ with covariance $\mathbb{E}[v_j v_j^\top] = \frac{s_v^2}{d_v} I_{d_v}$ (for $d_v \times d_v$ identity I_{d_v}), so $s_v^2 = \operatorname{tr}\big(\frac{s_v^2}{d_v} I_{d_v}\big) = \mathbb{E}\,\|v_j\|_2^2$. Under single-pass Euclidean delta updates with $\beta_t = 1$ from $S_0 = 0$, updating $S_t = (I_{d_k} - k_t k_t^\top) S_{t-1} + k_t v_t^\top$ writes later values $j > i$ with projection $\langle k_j, k_i \rangle v_j$ along k_i: for $i = N - 1$,

$$S_N^\top k_{N-1} - v_{N-1} = \langle k_N, k_{N-1} \rangle \big(v_N - S_{N-1}^\top k_N\big),$$

so, because v_N has zero mean and is independent of S_{N-1},

$$\mathbb{E}\,\big\|S_N^\top k_{N-1} - v_{N-1}\big\|_2^2 = \nu_{\text{coh}}^2 \big(s_v^2 + \mathbb{E}\,\big\|S_{N-1}^\top k_N\big\|_2^2\big) \geq \nu_{\text{coh}}^2 s_v^2.$$

Every later write with nonzero coherence therefore perturbs the readout of earlier keys, although the total error need not grow linearly in N.

For the diagonal column-energy preconditioner $K \operatorname{diag}(p(K))^{-1/2}$ where $p_m(K) = \sum_{i=1}^N K_{i,m}^2$, rotating $K \mapsto KQ$ with $Q = \operatorname{diag}(\xi) Q_H \in \mathrm{O}(d_k)$ for random Rademacher signs $\xi \in \{\pm 1\}^{d_k}$ followed by a normalized Hadamard matrix $Q_H \in \{\pm d_k^{-1/2}\}^{d_k \times d_k}$ equalizes all column energies in expectation,

$$\mathbb{E}_\xi \big[p_m(K \operatorname{diag}(\xi) Q_H)\big] = \frac{1}{d_k} \|K\|_F^2$$

(and by Schur-Horn some rotation $Q \in \mathrm{O}(d_k)$ makes $Q^\top K^\top K Q$ exactly constant-diagonal), reducing $\operatorname{diag}(p(KQ))^{-1/2}$ to a scalar multiple of I_{d_k} that leaves $\kappa_2((KQ)(KQ)^\top) = \kappa_2(KK^\top)$ unchanged; by contrast, $(I_C + \eta_c K_c K_c^\top)^{-1/2} K_c$ transforms equivariantly ($K_c \mapsto K_c Q \implies \tilde{K}_c \mapsto \tilde{K}_c Q$) for all $Q \in \mathrm{O}(d_k)$.

Part 2 (Gated RLS Overwrite-Retention Trade-off). Under gated RLS [70], with key covariance $H_t \in \mathbb{R}^{d_k \times d_k}$ and key-value moment $B_t \in \mathbb{R}^{d_k \times d_v}$ (local to this paragraph), $H_t = \gamma H_{t-1} + k_t k_t^\top$ and $B_t = \gamma B_{t-1} + k_t v_t^\top$ with scalar forget factor $\gamma \in (0, 1]$ and ridge ϵI_{d_k} ($\epsilon > 0$); readout on query $k \in \mathbb{R}^{d_k}$ is $B_t^\top H_t^{-1} k$. If a unit key k is written $r \geq 2$ times, with value $v^{(m)}$ at step t_m ($t_1 < \cdots < t_r = t$) and weights $w_m = \gamma^{t - t_m}$, Sherman-Morrison on $H_t = \epsilon I_{d_k} + (\sum_{m=1}^r w_m) k k^\top$ gives

$$\lim_{\epsilon \to 0^+} B_t^\top H_t^{-1} k = \frac{\sum_{m=1}^r w_m v^{(m)}}{\sum_{m=1}^r w_m}.$$

Letting $\bar{v}_{1:r-1} = \frac{1}{r-1}\sum_{m=1}^{r-1} v^{(m)}$ denote the mean of the first $r-1$ values, readout at $\gamma = 1$ ($w_m \equiv 1$) equals $\frac{1}{r}\sum_{m=1}^{r} v^{(m)}$ with error $\frac{r-1}{r}\left\|v^{(r)} - \bar{v}_{1:r-1}\right\|_2$ relative to $v^{(r)}$. Already for $r = 2$ consecutive writes and any target relative tolerance $\varepsilon \in (0, 1/2)$, the error is $\frac{\gamma}{1+\gamma}\left\|v^{(2)} - v^{(1)}\right\|_2$, so making it at most $\varepsilon\left\|v^{(2)} - v^{(1)}\right\|_2$ requires $\gamma \le \varepsilon/(1-\varepsilon)$, which simultaneously decays any non-overwritten key stored $t-j$ steps earlier (at step $j < t$) by γ^{t-j} relative to the regularizer ϵI_{d_k}, causing rapid forgetting in finite precision (verified in Table 12). Under oblique whitened delta, intra-chunk Gram whitening orthogonalizes concurrent keys while the oblique erasure $(I_{d_k} - \beta_t \tilde{k}_t k_t^\top)$ subtracts the current state projection $\hat{v}_t = (F_t \tilde{S}_{t-1})^\top k_t$ selectively along $\tilde{k}_t$ without decaying orthogonal associations.

Part 3 (Polar-2+ Non-Circular Normalization, Equioscillation, and BF16 Safety). Let $\Sigma_{c-1} \succ 0$ (initialized from $\Sigma_0 = I_{d_k}$) and let $Y_c = K_c \Sigma_{c-1}^{1/2} \in \mathbb{R}^{C \times d_k}$ have singular values $\sigma_i \ge 0$ (unrelated to the sigmoid σ), so the unregularized Gram $\mathcal{G}_0 = I_C + \eta_c Y_c Y_c^\top$ ($\eta_c \in (0, 1]$) has eigenvalues $\omega_{0,i} = 1 + \eta_c \sigma_i^2$. Computing $\bar{\mathcal{G}}_0 = \mathcal{G}_0 / \left\|\mathcal{G}_0\right\|_F$ and

$$\hat{u}_0 = \max\left(\left\|\mathcal{G}_0\right\|_F \left\|\bar{\mathcal{G}}_0^2\right\|_F^{1/2},\, 1\right) = \max\left(\left(\sum_j \omega_{0,j}^4\right)^{1/4},\, 1\right) \ge \max_i \omega_{0,i}$$

yields a non-circular upper bound on $\lambda_{\max}(\mathcal{G}_0)$ from one $[C, C]$ matmul $\bar{\mathcal{G}}_0^2$. Setting adaptive ridge $\varepsilon_{\text{reg}} = \max\left(0.02,\, \frac{0.0205}{0.9795}\hat{u}_0 - 1\right)$, $\mathcal{G}_{\text{reg}} = \mathcal{G}_0 + \varepsilon_{\text{reg}} I_C$ with shifted eigenvalues $\omega_i = \omega_{0,i} + \varepsilon_{\text{reg}}$, and $u_0 = \hat{u}_0 + \varepsilon_{\text{reg}} \ge \max_i \omega_i$ guarantees normalized eigenvalues $\lambda_i = \omega_i / u_0 \in [0.02, 1.0]$ for all $\hat{u}_0 \ge 1$ while expanding

$$T_0^2 = \frac{(\mathcal{G}_0 + \varepsilon_{\text{reg}} I_C)^2}{u_0^2} = \frac{\left\|\mathcal{G}_0\right\|_F^2 \bar{\mathcal{G}}_0^2 + 2\varepsilon_{\text{reg}} \mathcal{G}_0 + \varepsilon_{\text{reg}}^2 I_C}{u_0^2}$$

with zero additional matmuls (5 matmuls total across stages $\ell \in \{0, 1\}$, where $\mathcal{M}_\ell = m_\ell(T_\ell)$ evaluates $m_0(\lambda) = a_0 + b_0\lambda + e_0\lambda^2$ and $m_1(x) = a_1 + b_1 x + e_1 x^2$).

Clamping $\beta_t \le (2 - \varepsilon_\beta)/\max(k_t^\top \tilde{k}_t, \varepsilon_\beta)$ keeps the nontrivial eigenvalue $1 - \beta_t k_t^\top \tilde{k}_t$ of $I_{d_k} - \beta_t \tilde{k}_t k_t^\top$ in $[-1 + \varepsilon_\beta, 1]$ whenever $k_t^\top \tilde{k}_t > 0$ (in packed LM training, intra-chunk causal Schur-oblique projection sets

$$\tilde{k}_t = \frac{(I_{d_k} - \eta_c \Sigma_{c-1} K_{<t}^\top K_{<t}) \Sigma_{c-1} k_t}{k_t^\top (I_{d_k} - \eta_c \Sigma_{c-1} K_{<t}^\top K_{<t}) \Sigma_{c-1} k_t}$$

so $k_t^\top \tilde{k}_t = 1$ depends strictly on rows $< t$, while symmetric `Polar-2+` governs the chunk-boundary Woodbury tail, two-sided chunk-whitening kernels, and lagged block-32 causal whitening). Substituting $\mathcal{W}_{\text{isq}} = \alpha_{\text{safe}} u_0^{-1/2} \mathcal{M}_1 \mathcal{M}_0 = \alpha_{\text{safe}} u_0^{-1/2} z_2(T_0)$ and $\tilde{K}_c = \sqrt{\eta_c}\, \mathcal{W}_{\text{isq}} K_c \Sigma_{c-1}$ (Eq. (14)) into the unrelaxed Woodbury update $\tilde{\Sigma}_c = \Sigma_{c-1} - \tilde{K}_c^\top \tilde{K}_c$ (Eq. (15) at $\rho_p = 0$) gives $\tilde{\Sigma}_c = \Sigma_{c-1}^{1/2}(I_{d_k} - \eta_c Y_c^\top \mathcal{W}_{\text{isq}}^2 Y_c)\Sigma_{c-1}^{1/2}$, whose i-th singular-direction multiplier is

$$1 - \eta_c \sigma_i^2 u_0^{-1} \alpha_{\text{safe}}^2 z_2(\lambda_i)^2 = 1 - \frac{\eta_c \sigma_i^2}{1 + \varepsilon_{\text{reg}} + \eta_c \sigma_i^2} \lambda_i \left(\alpha_{\text{safe}} z_2(\lambda_i)\right)^2. \tag{17}$$

Direct evaluation of the scalar map z_2 (Proposition A.2) on $\lambda \in [0.02, 1.0]$ gives minimax equioscillation $\sqrt{\lambda} z_2(\lambda) \in [0.9757, 1.0243]$, hence $\lambda z_2(\lambda)^2 \in [0.9520, 1.0492]$ (and < 0.952 on $(0, 0.02)$ since $z_2(0) = a_1 a_0 \approx 9.877$). Thus unscaled $\alpha = 1$ exceeds 1 by $+4.92\%$ near equioscillation peaks, driving Eq. (17) negative when $\eta_c \sigma_i^2 \gg 1$. Setting $\alpha_{\text{safe}} = 0.968$ ($\alpha_{\text{safe}}^2 = 0.937024$) maps exact $\lambda(\alpha_{\text{safe}} z_2(\lambda))^2$ onto $[0.8920, 0.9831]$ across $[0.02, 1.0]$, giving exact $\kappa_2(\mathcal{W}_{\text{isq}} \mathcal{G}_{\text{reg}} \mathcal{W}_{\text{isq}}) = 0.9831/0.8920 = 1.102$ (≤ 1.123 in 5-matmul BF16 hardware execution). At $\rho_p = 0$, 4,096 correlated BF16 chunks (524,288 tokens) maintain $\lambda_{\min}(\Sigma_c) \ge 0.0388 > 0$; when evolving Σ_c across thousands of packed BF16 LM gradient steps with $\rho_p < 1$, replacing additive subtraction with multiplicative Sylvester congruence $\mathcal{B}_c \Sigma_{c-1} \mathcal{B}_c^\top$ (with contraction matrix $\mathcal{B}_c = I_{d_k} - \frac{\eta_c}{2}\Sigma_{c-1} K_c^\top \mathcal{W}_{\text{isq}}^2 K_c$, or $\mathcal{B}_c = I_{d_k} - \hat{C}_c$ using row-sum–normalized key covariance $\hat{C}_c \propto \eta_c K_c^\top K_c$ in our kernel) enforces PSD preservation unconditionally (Appendix C). □

B Associative Recall, Gated RLS Trade-Off, and Gate Conditioning

Table 12 reports synthetic associative-recall SNR (dB) in both BF16 and FP32 alongside Mod-2 parity accuracy and the chunk forget-gate condition number $\kappa_{\text{forget}} = \kappa_2(F_c)$ ($d_k = 256, C = 128$, 8 seeds, correlated keys with $\kappa_2(\mathcal{G}) \approx 48$). Separating BF16 from FP32 clarifies precision ceilings: exact inverse-square-root whitening reaches > 140 dB ($r = 2, 3$ overwrite) and 51.9 dB (write-once) in FP32, whereas BF16 readout caps SNR near 59–60 dB due to 7-bit mantissa rounding ($20 \log_{10}(2^7\sqrt{12}) \approx 53$–$60$ dB). Unwhitened Euclidean delta rules reach 58.40–60.97 dB on the final $r = 3$ overwrite via recency-weighted last-write erasure, whereas at $r = 2$ KDA's channel decay ($\kappa_{\text{forget}} = 1.1 \times 10^7$) collapses to 0.07 dB (and write-once -0.84 dB) and unwhitened SLA ($\kappa_{\text{forget}} = 1.00$) drops to 24.10 dB vs. 60.33 dB for `SPD-SLA-W`. Sweeping gated MesaNet's forget factor $\gamma \in \{1.00, 0.98, 0.80\}$ directly confirms Proposition A.1(2): $\gamma = 1.00$ and $\gamma = 0.98$ suffer severe stale-value averaging (1.90 and 4.13 dB at $r = 3$), whereas $\gamma = 0.80$ improves overwrite SNR (29.53 dB) only by degrading BF16 write-once retention from 10.63 dB down to 5.78 dB.

Table 12: Measured associative recall SNR (dB, reported as **BF16 [FP32]**) and Mod-2 parity accuracy across linear attention update rules ($d_k = 256, C = 128$, 8 seeds, $\kappa_2(\mathcal{G}) \approx 48$); $\kappa_{\text{forget}} = \kappa_2(F_c)$ is the condition number of the chunk forget gate.

Mechanism / Update Rule	Axis-Aligned (FP32)	Hadamard-Rot. (FP32)	Write-Once BF16 [FP32]	$r = 2$ Overwrite BF16 [FP32]	$r = 3$ Overwrite BF16 [FP32]	Mod-2 Parity (κ_{forget})
Mamba-2 / SSD (Scalar Hebbian)	2.14 dB	2.11 dB	1.85 [2.11]	1.72 [1.89]	1.40 [1.52]	50.2% (1.00)
KDA Real Channel Decay ($\beta \in (0, 1)$)	5.31 dB	5.29 dB	-0.84 [-0.91]	0.07 [0.26]	60.97 [> 140]	51.2% (1.1×10^7)
Diagonal Preconditioned Delta [68]	22.41 dB	6.12 dB	4.92 [6.12]	16.80 [19.05]	18.10 [21.40]	51.4% (1.1×10^7)
MesaNet Unforgetted ($\gamma = 1.00$) [70]	95.08 dB	95.06 dB	**10.63** [22.79]	2.67 [3.23]	1.90 [1.73]	50.8% (1.00)
MesaNet Gated ($\gamma = 0.98$)	91.40 dB	91.38 dB	9.54 [22.56]	7.32 [8.35]	4.13 [4.07]	51.0% (1.00)
MesaNet Aggressively Gated ($\gamma = 0.80$)	74.20 dB	74.18 dB	5.78 [26.09]	47.68 [89.36]	29.53 [30.22]	51.1% (1.00)
SLA ($G = 4$ + SO(2) + $\beta \in (0, 2)$)	7.84 dB	7.82 dB	6.12 [7.82]	24.10 [26.74]	58.40 [> 140]	**82.2%** (**1.00**)
SPD-SLA-W (Polar-2+ Oblique Delta)	**95.05** dB	**95.05** dB	7.01 [**51.93**]	**60.33** [> 140]	**59.26** [> 140]	**82.0%** (**1.00**)

C Negative Results and Reproducibility Details

We record three negative results observed during hardware-algorithm co-design alongside baseline software versions:

1. **Tier-1 VMEM Fusion Without Shape Co-Design** ($1.03\times$)**.** Fusing Ops 2–7a on GPU-shaped $8 \times 128 \times 256$ KDA yields only 3% forward-backward speedup on TPU v7x ($7.108 \to 6.912$ ms; Table 4) because per-channel VPU sub-block rescaling, $H_v = 8$ per-value-head chunk solves, and 50%-padded FP32 state scans remain VPU- and compute-bound inside VMEM.

2. **Unscaled Minimax Polynomials, BF16 Cholesky, and Additive Woodbury Drift in Live LM Training.** Using exact-arithmetic minimax coefficients ($\alpha = 1.000$) causes $\lambda z_2(\lambda)^2$ to reach $1.0492 > 1$, driving BF16 Woodbury covariance Σ_c indefinite ($\lambda_{\min} = -0.0412$; Table 11) and diverging 12.5B MoE pre-training to 258.5 PPL by step 299, while direct BF16 Cholesky/Woodbury hits `NaN` at step 54. Scaling by $\alpha_{\text{safe}} = 0.968$ bounds the exact peak to 0.9831, preserves $\lambda_{\min} \geq 0.0388$ across 4,096 chunks (524K tokens at $\rho_p = 0$), and trains cleanly for 7,200 steps across 2 seeds at $\rho_p = 1.0$ (0 NaNs; seed-matched `s1`: 17.936 vs. 17.970 PPL for `NS-5`, 2-seed `Polar-2+` mean 17.970 on the `+MS` arms [`s0` 18.004]; the uniform-decay arm reaches 17.575 in Table 25). When evolving Σ_c across thousands of packed BF16 LM gradient steps with $\rho_p = 0.02$, additive subtraction $\Sigma_{c-1} - \tilde{K}_c^\top \tilde{K}_c$ still accumulates BF16 rounding drift (`NaN` at step 282), whereas multiplicative Sylvester congruence $\mathcal{B}_c \Sigma_{c-1} \mathcal{B}_c^\top$ preserves PSD unconditionally (and strictly causal block-32 lagged two-sided `Polar-2+`, `polar2_causal_blk32`, completes all 7,200 steps [15.10B tokens] on 12.5B `15L:1A` with 0 NaNs, compressing $\kappa_2(I_C + M_c)$ from 11.9–79.3 to 1.12–1.40 across all 15 linear layers and reaching 17.684 11-set PPL, 11.49 32K Tail PPL, 49.0% `norm9` [49.1% FP8], and 68.8% 64K `HotpotQA`).

3. **Global First-Order Neumann Inversion at $C \geq 128$, Write-Budget Horizon, and Roofline-Capacity Proxy Validation.** On a cluster of b aligned keys ($H = \mathbf{1}_b \mathbf{1}_b^\top$), one stationary Jacobi sweep $I_b - \beta H$ has spectral radius $|b\beta - 1| > 1$ whenever $b\beta > 2$, causing unscaled $(I_C + M_c)^{-1} \approx I_C - M_c$ to incur 62.2% intra-chunk WY error on coherent bursts ($+0.014$ nats validation loss on Dolma 3). Our row-scaled variant $A_{\text{N1}}^{\text{row}}$ in `A3/A5` enforces Gershgorin row damping ($r_i = \max(1, 1.25 \, \|(M_c)_{i,:}\|_1)$), reducing coherent-burst WY error

to 38.1% at $+0.22$ pre-training perplexity (A2→A3, Table 8b), whereas inverting 16×16 diagonal blocks on the MXU (A_{N1}, `neumann_order=16`) cuts WY error $32\times$ to 1.18% (0.0 14-layer tail logit KL). Under isotropic unit keys, write erasure attenuates stored associations as $\prod_j (1 - \beta_j/d_k) \approx \exp(-\sum_j \beta_j/d_k)$ (and directionally as $\exp(-\beta c_i T)$ along key-covariance eigenvalues c_i, motivating per-tile whitening in Appendix A); across the 14 linear layers of our 12.5B `SLA-R4` checkpoint ($\bar{\beta}_\ell \in [0.251, 0.348]$), 8K-window write mass $\sum \beta_t \in [2{,}056, 2{,}851] = 8.0$–$11.1 d_k$ sets an associative horizon $d_k/\bar{\beta}_\ell \in [736, 1{,}020]$ tokens (mean 825). Scoring our 10 trained 12.5B arms (`KDA`, `A1`–`A6`, `GDN-2`, `SLA-R2`, `SLA-R4`) with the analytical systolic roofline $T(\theta)$ and multi-horizon capacity proxy achieves Spearman rank correlation $r_s = 0.988$ ($p = 9.3 \times 10^{-8}$) against measured step time ($r_s = 0.833, p = 0.010$ on isolated 32K layer latency) and $r_s = 0.769$ ($p = 0.0093$) against negative validation perplexity.

D Notation Summary, Momentum Ablation, and Multi-Scale Architecture Specifications

Tables 13, 14, and 15 provide the paper-wide symbol glossary (Section 2.3), the single-seed 128-chip TPU v4 chunk-boundary momentum ablation ($\mu \in \{0, 0.25\}$), and the exact architectural hyperparameters across all four hybrid `LatentMoE` pre-training rungs (Section 3.4).

Table 13: Summary of notation used across the paper.

Symbol	Meaning
$T,\ C,\ c$	Sequence length; chunk length ($C = 128$); chunk index $c \in \{1, \dots, T/C\}$.
t, j, s	Token indices inside a chunk, $1 \le j \le t \le C$ (with chunk-entry boundary 0 in $P_{t\leftarrow 0}$).
$D,\ H_k,\ H_v,\ R$	Model width; number of query-key heads; number of value heads; sharing ratio $R = H_v/H_k \in \{2, 4\}$.
$d_k,\ d_v,\ D_{v,\text{pack}},\ D_v$	Key width ($d_k = 256$); per-value-head width; packed value width per key head $D_{v,\text{pack}} = Rd_v$; total value width $D_v = H_v d_v$.
$G,\ b,\ g$	Number of decay tiles ($G = 4$), tile width $b = d_k/G = 64$, tile index $g \in \{1, \dots, G\}$.
$q_t, k_t \in \mathbb{R}^{d_k}, v_t \in \mathbb{R}^{D_{v,\text{pack}}}$	L_2-normalized query and key of one key head; packed values of its R value heads.
$\beta_t \in (0, s_\beta)$	Write gate $s_\beta \sigma(\tilde{\beta}_t)$ per token and key head, shared across the R packed value heads ($s_\beta = 2$ in SLA; $s_\beta = 1$ in GDN/KDA).
$f_{t,i},\ \gamma_{t,g} \in (0, 1]$	Per-channel decay of channel i (KDA, $G = d_k$); per-tile decay $\gamma_{t,g} = e^{-\operatorname{softplus}(\tilde{f}_{t,g})}$ of tile g (SLA, $f_{t,i} = \gamma_{t,g}$ on tile g).
$F_t,\ P_{t\leftarrow j},\ F_c$	Diagonal forget gate $F_t = \operatorname{diag}(f_t)$; cumulative decay $P_{t\leftarrow j} = \prod_{s=j+1}^{t} F_s$ with $P_{t\leftarrow t} = I_{d_k}$; chunk decay $F_c = P_{C\leftarrow 0}$.
$\Lambda_{t,i}$	Cumulative log-decay of channel i inside a chunk, $\Lambda_{t,i} = \sum_{s=1}^{t} \log f_{s,i}$.
$S_c,\ \tilde{S}_t,\ \Delta S_c$	Chunk-boundary state $S_c \in \mathbb{R}^{d_k \times D_{v,\text{pack}}}$; intra-chunk state after local token t ($\tilde{S}_0 = S_{c-1}$, $\tilde{S}_C$ written $\tilde{S}_c$ in Def. 2.1(ii)); increment $\Delta S_c = \tilde{S}_C - F_c S_{c-1}$.
$Z_c,\ \mu,\ \kappa_c$	Chunk-boundary momentum (velocity) buffer; momentum coefficient ($\mu = 0$ by default; $\mu = 0.25$ in earlier runs); reset indicator $\kappa_c \in \{0, 1\}$, equal to 0 whenever any document reset occurs inside chunk c (so Z_{c-1} and cross-document carry are discarded).
$\Psi_g,\ \Gamma_g \in \mathbb{R}^{C\times C}$	Tile-g causal decay masks: $(\Psi_g)_{t,j} = \mathbf{1}[t \ge j] \prod_{s=j+1}^{t} \gamma_{s,g}$ and $(\Gamma_g)_{t,j} = \mathbf{1}[t > j]\, \beta_j (\Psi_g)_{t,j}$.
$M_c,\ L_c,\ A_c \in \mathbb{R}^{C\times C}$	Strictly lower-triangular Gram matrix M_c (Op 3a), lower-triangular query-key score matrix L_c (Op 3b), and WY inverse $A_c = (I_C + M_c)^{-1}$ (Op 4).
$\bar{Q}_c, \bar{K}_c, \hat{K}_c \in \mathbb{R}^{C\times d_k}$	Row-stacked decayed queries $(P_{t\leftarrow 0} q_t)^\top$, decayed keys $(P_{t\leftarrow 0} k_t)^\top$, and carried keys $(P_{C\leftarrow t} k_t)^\top$.
$U_c,\ W_c,\ E_c$	WY factors $U_c = \operatorname{diag}(\beta_c) A_c V_c$, $W_c = \operatorname{diag}(\beta_c) A_c \bar{K}_c$, and chunk writes $E_c = U_c - W_c S_{c-1}$ (Ops 5–6a).
$P,\ p,\ L_{\text{loc}}$	Number of context-parallel shards (devices); shard index; chunks per shard.
$\sigma(\cdot),\ \kappa_2(\cdot)$	Logistic sigmoid; spectral condition number (distinct from the reset indicator κ_c).

Table 14: Chunk-boundary momentum ablation: 20-step rolling training cross-entropy at step 7,199 (single pass over Dolma 3.5, 12.5B MoE, 128 TPU v4 chips, seed 42; lower is better). †Trained on an 8,200-step cosine schedule and evaluated at step 7,199.

Configuration	μ	Loss @ 7,199
KDA (per-channel decay, FP32 state)	–	2.2925
SLA, $R = 2$ (two runs)	0	2.2912 / 2.2925
SLA, $R = 2$	0.25	2.2955
SLA, $R = 2$, input-dependent μ	adaptive	2.3216
SLA, $R = 4$ (two runs)	0	2.3103 / 2.3148
SLA, $R = 4^\dagger$	0.25	**2.2914**

Hybrid `LatentMoE` backbone and multi-scale parameterization. All four pre-training scales are instantiated from a single TPU-aligned hybrid linear-softmax `LatentMoE` decoder family implemented in open-source MaxText [1] (Table 15). Every weight matrix (embeddings, projections,

routers, and experts) uses fixed truncated-normal initialization (standard deviation $\sigma_{\text{init}} = 0.02$, clipped at $\pm 3\sigma_{\text{init}}$), and each decoder block applies four $\mathrm{RMSNorm}$ layers ($\epsilon = 10^{-6}$): pre- and post-sublayer norms around both the sequence mixer and the FFN, with post-norms applied to the sublayer output prior to residual addition, followed by a final stack $\mathrm{RMSNorm}$.

Layer 0 uses a dense SwiGLU FFN [56] of intermediate width F_{dense} with no routed experts to prevent early-layer routing collapse. Layers $1, \ldots, L-1$ pair one full-width shared SwiGLU expert of intermediate width F_{sh} with `LatentMoE`: a shared linear down-projection compresses the normalized residual stream from $\mathbb{R}^D$ to an MXU-aligned latent width $\ell < D$, $E = 256$ routed SwiGLU experts (intermediate width $F_{\text{moe}} = 2D$) execute inside $\mathbb{R}^\ell$, and a shared up-projection maps the weighted expert sum from $\mathbb{R}^\ell$ back to $\mathbb{R}^D$. Routing scores are still computed on the uncompressed residual in $\mathbb{R}^D$ in FP32, using top-k softmax weights rescaled to sum to k together with document-pool routing (restricting each document to a pool size in $[k, E]$). Along depth, each $(x{+}1)$-layer cycle places linear delta mixers (4-tap causal depthwise `Conv1D` with SiLU, L_2-normalized queries/keys, low-rank decay and sigmoid output-gate projections of rank r_{proj}, and per-group output $\mathrm{RMSNorm}$ at $\epsilon = 10^{-5}$) on the first x layers and gated Grouped-Query Attention (GQA) [2, 52] ($H_q = 8, H_{kv} = 4$) on every $(x{+}1)$-th layer, with QK-$\mathrm{RMSNorm}$, a $2d_{\text{head}}$-widened query projection (for per-head dimension d_{head}) split into attention queries and an elementwise sigmoid output gate, and either RoPE ($\theta_{\max} = 5 \times 10^5$) [61] or NoPE with Scalable Softmax [45].

Rather than choosing intermediate widths that pad 256×256 systolic tiles, we step across two 256-divisible width tiers ($D = 1024 = 4 \times 256$ for Rungs 1–2; $D = 1792 = 7 \times 256$ for Rungs 3–4, a $1.75\times$ width step) and scale depth L, active expert count k, and shared-expert width F_{sh} within each tier:

1. **Embeddings and untied LM head ($2VD$ linear scaling):** All four scales share a fixed vocabulary $V = 100{,}352$ with untied $V \times D$ input embedding and $D \times V$ output LM-head matrices. Input token embeddings are multiplied by $\sqrt{D}$ and passed through $\mathrm{RMSNorm}(D)$ before Layer 0 so initial residual-stream RMS remains invariant across D under fixed $\sigma_{\text{init}} = 0.02$ initialization. Because V is fixed, combined embedding and LM-head parameters $2VD$ grow strictly linearly with D (205.5M $\to$ 359.7M, $1.75\times$), whereas non-embedding capacity grows with both width and depth; accordingly, $2VD$ drops monotonically from 46.6% of active (2.71% of total) parameters at Rung 1 to 9.6% (0.57%) at Rung 4.
2. **Sublayer widths across tiers ($D : 1024 \to 1792$):** Stepping from Tier A to Tier B scales the `LatentMoE` bottleneck $1.5\times$ ($\ell : 512 \to 768$, keeping $\ell/D \in [0.43, 0.50]$ so expert All-to-All dispatch and grouped-matmul K-tiles stay compressed and 256-aligned) and routed expert width $1.75\times$ ($F_{\text{moe}} = 2D : 2048 \to 3584$). In parallel, global GQA head dimension ($d_{\text{head}} : 128 \to 256$), linear value width ($D_v : 2048 \to 4096$), low-rank gate bottleneck ($r_{\text{proj}} : 256 \to 512$), and output $\mathrm{RMSNorm}$ groups ($8 \to 16$, fixing 256 channels/group) all double ($2\times$). Baseline KDA doubles $d_k : 128 \to 256$ ($H_k = 8$), whereas `SLA-R2` ($H_k = 4$) and `SLA-R4` ($H_k = 2$) hold $d_k = 256$ fixed at every scale to preserve unpadded 256×256 MXU state updates while doubling packed value width Rd_v.
3. **Within-tier depth, sparsity, and iso-parameter compensation:** Within $D = 1024$ (Rung 1 $\to$ Rung 2), depth increases from $L = 10$ to 16, active routed experts double ($k : 4 \to 8$), and Layer 0 widens ($F_{\text{dense}} : 4D \to 8D$). Within $D = 1792$ (Rung 3 $\to$ Rung 4), depth doubles ($L : 15 \to 30$) and shared-expert width doubles ($F_{\text{sh}} : D \to 2D$). Finally, because 4:1 key-head sharing in `SLA-R4` reduces per-layer linear mixer parameters (7.63M $\to$ 6.31M at $D = 1024$; 27.05M $\to$ 20.48M at $D = 1792$), `Iso-Param SLA-R4` widens F_{sh} by $+256$ ($1024 \to 1280$ on Rungs 1–2; $1792 \to 2048$ on Rung 3) to match baseline parameter count.

Table 15: Exact model architecture and multi-scale parameterization across our four hybrid `LatentMoE` pre-training rungs ($V = 100{,}352$ untied vocabulary, $E = 256$ routed experts on layers $1, \ldots, L-1$, dense SwiGLU on layer 0).

Architecture / Scaling Parameter	**Rung 1**	**Rung 2**	**Rung 3**	**Rung 4**
Cohort Label (Active / Total)	$0.35\text{B}/6.5\text{B}^*$	$0.79\text{B}/12.5\text{B}$	$1.80\text{B}/31.5\text{B}^*$	$3.74\text{B}/63.1\text{B}$
Exact Total Params (KDA / `SLA-R2`)	$7.58\text{B}/7.57\text{B}$	$12.49\text{B}/12.49\text{B}$	$30.58\text{B}/30.52\text{B}$	$63.13\text{B}/63.02\text{B}$
Exact Active: incl. $2VD$ ($1VD$ FLOP)*	0.44B (0.34B)	0.79B (0.69B)	1.91B (1.73B)	3.74B (3.56B)
Decoder Depth L & Hybrid Cycle	10 (`4L`: 8+2)	16 (`7L`: 14+2; `15L`)	15 (`4L`: 12+3)	30 (`4L`: 24+6; `15L`)
Model Width D ($4 \times \text{RMSNorm}$/block)	1024 (4×256)	1024 (4×256)	1792 (7×256)	1792 (7×256)
Embeddings + Untied LM Head ($2VD$)	205.5M ($\sqrt{D}$+Norm)	205.5M ($\sqrt{D}$+Norm)	359.7M ($\sqrt{D}$+Norm)	359.7M ($\sqrt{D}$+Norm)
$2VD$ Share of Active / Total Params	46.6% / 2.71%	26.0% / 1.65%	18.8% / 1.18%	9.6% / 0.57%
Layer 0 Dense SwiGLU Width F_{dense}	4096 ($4D$)	8192 ($8D$)	14,336 ($8D$)	14,336 ($8D$)
`LatentMoE` Bottleneck ℓ (ℓ/D)	512 ($0.50D$)	512 ($0.50D$)	768 ($0.43D$)	768 ($0.43D$)
Routed Experts (E, Top-k, F_{moe})	256, k=4, 2048	256, k=8, 2048	256, k=8, 3584	256, k=8, 3584
Shared Expert F_{sh} (Base / `Iso-R4`)	1024 / 1280	1024 / 1280	1792 / 2048	3584 ($2D$)
Global Gated GQA (H_q/H_{kv}, d_{head})	8/4, d_{head}=128	8/4, d_{head}=128	8/4, d_{head}=256	8/4, d_{head}=256
Linear D_v, Gate Rank r_{proj}, Norm Grps	2048, r_{proj}=256, 8	2048, r_{proj}=256, 8	4096, r_{proj}=512, 16	4096, r_{proj}=512, 16
Linear Geometry ($H_k \times d_k \times Rd_v$)	KDA $8\times128\times256$; `R2` $4\times256\times512$; `R4` $2\times256\times1024$		KDA $8\times256\times512$; `R2` $4\times256\times1024$; `R4` $2\times256\times2048$	
Mixer Params / Layer (KDA vs. `R4`; GQA)	7.63 vs. 6.31M; 4.20M	7.63 vs. 6.31M; 4.20M	27.1 vs. 20.5M; 14.7M	27.1 vs. 20.5M; 14.7M
Pre-Training Horizons Studied	15.1B, 30.2B	15.1B, 60.4B (+32K)	5.03B, 36.2B	5.03B, 30.2B, 120.8B

*Cohort labels $0.35\text{B}/6.5\text{B}$ and $1.80\text{B}/31.5\text{B}$ reflect the 8-layer proxy (6.08B total, 0.38B active) and half-63.1B shorthand; extending Rung 1 to $L = 10$ ($2 \times$ `4L:1A` cycles) and Rung 3 to $L = 15$ ($3 \times$ `4L:1A` cycles, $F_{\text{sh}} = 1792$) ensures exact cycle divisibility and yields the parameter counts above. Active parameters are shown both with $2VD$ (matching the 0.79B and 3.74B convention) and as FLOP-bearing active parameters ($1VD$, excluding input token lookup; pure non-embedding/non-head active parameters are 0.24B, 0.59B, 1.55B, 3.38B).

E Extended Kernel, Multi-Generation Roofline, and 12.5B MoE Cube Serving Benchmarks

Table 16 reports isolated core-kernel latencies and iso-state / iso-parameter controls on TPU v7x normalized to both open-source MaxText Pallas KDA ($C = 128$) and Tokamax Mosaic KDA ($C = 64$), and Table 17 isolates each architectural and kernel factor at the identical `SLA-R4` full-layer shape, complementing Table 5. Table 18 tabulates the four-generation TPU forward latency and leave-one-out roofline predictions plotted in Figure 5 (Section 3.1).

Table 16: Isolated core-kernel latency (ms, Ops 3–7a, batch size $B = 1$) and F+B speedups vs. open-source `MaxText Pallas KDA` ($C = 128$) and `Tokamax Mosaic KDA` ($C = 64$) on TPU v7x across (a) $D = 1792$ (63.1B MoE) and (b) iso-controlled $D = 2048$ (12.5B). Full-layer latencies are in Table 5.

Regime	Core Kernel ($H_k \times d_k \times Rd_v$, Chunk C, State)	8K Fwd	8K Bwd	8K F+B	vs. Max.	vs. Tok.	32K F+B	vs. Max.	vs. Tok.
(a) Core $D = 1792$ (63.1B)	Open-Source MaxText Pallas KDA ($8 \times 256 \times 512, C = 128$, FP32)	3.32	2.00	5.32	1.00×	0.63×	20.43	1.00×	0.61×
	Open-Source Tokamax Mosaic KDA ($8 \times 256 \times 512, C = 64$, Base)	1.95	1.40	3.35	1.59×	1.00×	12.45	1.64×	1.00×
	SLA-R2 ($4 \times 256 \times 1024, C = 128$, BF16)	1.96	0.91	2.87	1.86×	1.17×	10.15	2.01×	1.23×
	SLA-R4 ($2 \times 256 \times 2048, C = 128$, BF16)	**1.37**	**0.67**	**2.04**	2.61×	1.64×	**6.18**	3.31×	2.02×
(b) Iso-Ctrl Core Kernel $D = 2048$	Open-Source MaxText Pallas KDA ($8 \times 128 \times 256, C = 128$, FP32)	2.67	1.58	4.25	1.00×	0.44×	16.46	1.00×	0.35×
	Open-Source Tokamax Mosaic KDA ($8 \times 128 \times 256, C = 64$, Base)	1.07	0.78	1.85	2.29×	1.00×	5.83	2.82×	1.00×
	SLA Iso-State ($4 \times 256 \times 256, C = 128$, equal state & scan FLOPs)	1.48	0.93	2.42	1.76×	0.77×	7.57	2.18×	0.77×
	SLA-R2 ($4 \times 256 \times 512, C = 128$, equal $D_v = 2048$; 2× KDA state)	1.57	1.01	2.58	1.65×	0.72×	8.24	2.00×	0.71×
	SLA Iso-Mixer-Param ($4 \times 256 \times 640, C = 128$, equal mixer params)	1.63	1.06	2.69	1.58×	0.69×	8.72	1.89×	0.67×
	SLA-R4 ($2 \times 256 \times 1024, C = 128$, equal 1.0 MiB bytes)	**1.05**	**0.65**	**1.70**	2.49×	1.09×	**5.00**	3.29×	1.17×

Iso-shape full-layer attribution ($H_k = 2, d_k = 256, R = 4$) and Table 5 reconciliation. To separate the throughput contributions of 4:1 head sharing, state precision, tile-grouped decay (G), chunk inversion, and single-kernel VMEM fusion, Table 17 benchmarks all three decay granularities ($G \in \{d_k, 4, 1\}$) at the identical `SLA-R4` shape ($H_k = 2, d_k = 256, R = 4, C = 128$, BF16 activations, $B = 1$ per TensorCore) across pure XLA, Pallas chunk kernels with an XLA wrapper, single-kernel fused Pallas layers, and open-source Tokamax Mosaic KDA (0.0.14) packed to $R = 4$. All 14 kernels match an FP32 token-by-token reference recurrence with forward and input/weight-gradient cosine similarity at least 0.99998. Each cell reports Measured mean ± SE over $n = 3$ independent process restarts (fresh compilation and allocation per restart) on one isolated TPU v7x chip under pre-allocated inputs, alongside protocol-anchor ratios matching Table 5's RNG-in-step harness. This table also directly reconciles the $D = 1792, 32$K entries between Table 5a (20.24 ms) and Table 17: Table 5a's 20.24 ms (20.237 ms) is the 2-stage chunk kernel with $\mu = 0.25$ (20.30 ± 0.23 ms in Table 17a), which is already 1.69× faster than single-chip Tokamax Mosaic KDA (34.31 ± 0.13 ms) and 2.09× faster under Table 5's 64-chip concurrent timing (42.38 ms), while the 1-kernel fused variant (14.08 ± 0.30 ms, $\mu = 0$) is 2.44× faster than Tokamax Mosaic KDA at $D = 1792$. Furthermore, Table 17a shows why $G = 4$ is the quality-speed Pareto knee: scalar $G = 1$ runs at the same single-kernel VMEM speed as $G = 4$ (2.62 vs. 2.85 ms at 8K and 14.16 vs. 14.08 ms at $D = 1792, 32$K) because 1 or 4 gate vectors both fit in VMEM registers without $[B, H_k, T/C, C, C, d_k]$ HBM materialization, but $G = 1$ loses multi-timescale retention quality and associative recall (Table 10), whereas $G = 4$ matches $G = d_k$ quality while avoiding $G = d_k$'s 1.62×–2.45× XLA and 4.47×–6.65× portable kernel slowdown.

End-to-end 12.5B sparse MoE 1M-token `SLA Tree-CP`, cube serving, speculative acceptance, and direct `Tokamax Mosaic KDA` step-time comparison. To complement the isolated 128-core `SLA Tree-CP` kernel and single-chip 32-layer dense-FFN serving benchmarks in Table 6, Table 21 benchmarks a full 16-layer 12.5B sparse MoE (0.79B active, $E = 256$ routed experts with $k = 8$ active per token plus one shared expert, $D = 1024, D_v = 2048$, vocabulary size $V = 100{,}352$) end-to-end across all 64 chips (128 TensorCores) of a TPU v7x `4x4x4` cube, and Table 19 directly compares end-to-end 12.5B (`7L:1A`) training step time against open-source Tokamax Mosaic KDA (`tokamax` 0.0.14):

1. **End-to-End `CP=EP=128` Training Steps up to 1.05M Tokens (Table 21a):** On the isolated context-parallel scan across 64 chips, `SLA Tree-CP` achieves both a 1.09× speedup over `KDA-CP` (22.58 ms vs. 24.64 ms) and a 7.65× speedup over dense `RingCP` (22.58 ms vs. 172.71 ms). End-to-end on the 64-chip 12.5B MoE at $T = 1{,}048{,}576$ ($T_{\text{loc}} = 8{,}192$ per core),

Table 17: **Measured iso-shape full-layer F+B attribution on one isolated TPU v7x chip** (ms, mean $\pm$ SE over $n = 3$ independent restarts; container `maxtext_e174258_20261008`, Tokamax 0.0.14, XLA default 16 MiB scoped-VMEM cap, pre-allocated inputs, $B = 1$/core). Non-reference rows hold shape fixed at `SLA-R4` ($H_k = 2, d_k = 256, R = 4$, $D_{v,\text{pack}} = 1024$ at $D = 2048$ and 2048 at $D = 1792$, $C = 128$). Panel (b) gives ratios of means $\pm$ delta-method SE.

(a) Full-Layer Implementation (Ops 1–7c)	Gate G	D=2048, 8K	D=2048, 32K	D=1792, 8K	D=1792, 32K
Open-Source MaxText Pallas KDA, native 8 heads, FP32 state (ref)	d_k	6.89 ± 0.02	26.82 ± 0.12	10.55 ± 0.11	40.49 ± 0.17
Open-Source Tokamax Mosaic KDA, native 8 heads, C=64 (ref)	d_k	4.37 ± 0.02	16.45 ± 0.13	8.73 ± 0.01	34.31 ± 0.13
Open-Source Tokamax Mosaic KDA, iso-shape R=4 packed values*	d_k	4.31 ± 0.01	15.86 ± 0.10	VMEM OOM	VMEM OOM
Open-Source Tokamax Mosaic KDA, shared keys repeated to 8 heads	d_k	4.79 ± 0.00	17.77 ± 0.12	7.32 ± 0.01	28.44 ± 0.06
XLA chunked WY (unfused, identical JAX code)	d_k	17.58 ± 0.12	72.67 ± 0.10	22.23 ± 0.03	85.27 ± 0.07
XLA chunked WY (unfused, identical JAX code)	4	8.05 ± 0.01	29.67 ± 0.07	13.72 ± 0.13	51.33 ± 0.15
XLA chunked WY (unfused, identical JAX code)	1	7.82 ± 0.03	27.95 ± 0.04	12.78 ± 0.04	48.84 ± 0.20
Pallas KDA chunk kernel + XLA wrapper, FP32 state	d_k	4.03 ± 0.01	15.26 ± 0.10	5.63 ± 0.01	21.86 ± 0.16
Pallas KDA chunk kernel + XLA wrapper, BF16 state	d_k	4.00 ± 0.00	15.30 ± 0.18	5.59 ± 0.02	21.62 ± 0.13
Pallas KDA chunk kernel, scalar gate broadcast	1	3.95 ± 0.02	15.00 ± 0.02	5.56 ± 0.01	21.30 ± 0.18
Pallas SLA chunk kernel + XLA wrapper ($\mu = 0$, BF16 state)	4	3.82 ± 0.01	13.71 ± 0.17	5.37 ± 0.01	19.69 ± 0.12
Pallas SLA chunk kernel, scalar gate broadcast ($\mu = 0$)	1	3.70 ± 0.02	13.85 ± 0.09	5.31 ± 0.02	20.02 ± 0.31
Pallas SLA chunk kernel + $\mu = 0.25$ (2-stage chunk kernel, `SLA-R4`)	4	3.86 ± 0.00	14.05 ± 0.09	5.47 ± 0.02	20.30 ± 0.23
Fused single-kernel layer (`Exact Inv.`, $\mu = 0$)	4	$\mathbf{2.85 \pm 0.01}$	$\mathbf{9.71 \pm 0.01}$	$\mathbf{3.92 \pm 0.01}$	$\mathbf{14.08 \pm 0.30}$
Fused single-kernel layer (`Exact Inv.`, $\mu = 0$, scalar gate)	1	2.62 ± 0.01	9.33 ± 0.02	3.71 ± 0.01	14.16 ± 0.06
Fused `SLA-R4` (`Exact Inv.`, $\mu = 0.25$)*	4	$\mathbf{2.90 \pm 0.01}$	$\mathbf{9.90 \pm 0.06}$	16M cap*	16M cap*
Fused `SLA-R4` (`1-Step Neumann`, $\mu = 0.25$)*	4	$\mathbf{2.48 \pm 0.01}$	$\mathbf{8.25 \pm 0.01}$	16M cap*	16M cap*
(b) Isolated Factor Attribution	Comparison	D=2048, 8K	D=2048, 32K	D=1792, 8K	D=1792, 32K
4:1 head sharing (MaxText KDA kernel, FP32)	native 8h $\to$ 2×256, R=4	$1.71 \times \pm 0.01$	$1.76 \times \pm 0.01$	$1.87 \times \pm 0.02$	$1.85 \times \pm 0.02$
4:1 head sharing alone on Tokamax (packed)	native 8h $\to$ Tok. R=4	$1.01 \times \pm 0.01$	$1.04 \times \pm 0.01$	OOM	OOM
BF16 vs. FP32 state (iso-shape KDA Pallas)	FP32 $\to$ BF16 state	$1.01 \times \pm 0.00$	$1.00 \times \pm 0.01$	$1.01 \times \pm 0.00$	$1.01 \times \pm 0.01$
Gate $G = d_k \to G = 4$ (XLA, identical code)	XLA $d_k \to$ XLA 4	$2.18 \times \pm 0.01$	$2.45 \times \pm 0.01$	$1.62 \times \pm 0.02$	$1.66 \times \pm 0.01$
Gate $G = d_k \to G = 4$ (Pallas chunk kernels)	KDA BF16 $\to$ SLA (μ=0)	$1.05 \times \pm 0.00$	$1.12 \times \pm 0.02$	$1.04 \times \pm 0.00$	$1.10 \times \pm 0.01$
Gate $G = 4 \to G = 1$ (fused single-kernel)	Fused 4 $\to$ Fused 1	$1.09 \times \pm 0.01$	$1.04 \times \pm 0.00$	$1.06 \times \pm 0.00$	$0.99 \times \pm 0.02$
Single-kernel VMEM fusion ($G = 4, \mu = 0$)	Unfused $\to$ Fused	$1.34 \times \pm 0.01$	$1.41 \times \pm 0.02$	$1.37 \times \pm 0.00$	$1.40 \times \pm 0.03$
1-Step Neumann vs. Exact Inv. (fused, $\mu = 0.25$)	Exact $\to$ 1-Step	$1.17 \times \pm 0.01$	$1.20 \times \pm 0.01$	–	–
Fused `SLA-R4` vs. Tokamax iso-shape R=4	Tok. R=4 $\to$ `SLA-R4`	$\mathbf{1.49 \times \pm 0.01}$	$\mathbf{1.60 \times \pm 0.01}$	Tok. OOM	Tok. OOM
Fused `SLA-R4` (μ=0) vs. best iso KDA Pallas	KDA BF16 $\to$ Fused G=4	$\mathbf{1.40 \times \pm 0.01}$	$\mathbf{1.58 \times \pm 0.02}$	$\mathbf{1.43 \times \pm 0.01}$	$\mathbf{1.54 \times \pm 0.03}$

*__VMEM budgets and Table 5 timing anchor:__ Tokamax overrides `vmem_limit_bytes` internally to the full 64 MiB hardware VMEM, so its $D = 1792$, $R = 4$ backward pass exceeds physical capacity (83.3 MiB $>$ 63.9 MiB from two 32 MiB `f32[32,64,2048]` windows). By contrast, this suite omits `-xla_tpu_scoped_vmem_limit_kib=65472` (keeping XLA's default 16 MiB scoped-VMEM cap): fused `SLA-R4` fits at $\mu = 0$ across all sizes and at $\mu = 0.25$ for $D = 2048$, while adding $\mu = 0.25$ velocity tiles at $D = 1792$ ($D_{v,\text{pack}} = 2048$) requires the 64 MiB scoped-VMEM flag used in Table 5 (5.57/20.24 ms) and 63.1B pre-training. Re-running on one chip under Table 5's RNG-inside-step timing ($n = 3$) gives MaxText 8.55/31.47 (D=2048) and 12.64/45.99 ms (D=1792), Tokamax native 5.86/22.78 and 10.42/42.79 ms (R=4 packed: 5.54/20.62 ms), KDA BF16 Pallas iso-shape 5.36/18.78 and 7.11/26.37 ms, unfused SLA (μ=0) 4.93/17.55 and 6.86/24.29 ms, and fused SLA 4.12/14.43 ms (μ=0.25 at D=2048; μ=0: 4.08/14.15 and 5.61/19.02 ms), yielding 1.34 × /1.43× vs. iso-shape Tokamax, 1.27×–1.39× vs. iso-shape KDA Pallas, and 2.08×–2.42× (vs. MaxText) and 1.42×–2.25× (vs. native Tokamax) across all four cells ($[H_k, Rd_v]$ vs. 8-head value layout changes time by < 0.03 ms).

Table 18: Measured forward core-kernel latency (ms, $B = 1, T = 8{,}192, D = 2048$) vs. portable JAX/Pallas KDA ($C = 64$) and out-of-fold LOO Roofline predictions across four TPU generations.

TPU Generation	Array Tile	MXU:VPU (BF16)	KDA Measured (ms)	SLA-R4 Measured (ms)	Measured Speedup	Out-of-Fold LOO Pred. (Error)
TPU v4	128 × 128	85.9×	14.84	3.32	4.47×	4.40× (−1.6%)
TPU v5p	128 × 128	85.0×	9.19	1.98	4.64×	4.71× (+1.5%)
TPU v6e	256 × 256	127.5×	11.02	1.88	5.86×	6.15× (+5.0%)
TPU v7x	256 × 256	128.2×	6.99	1.05	6.65×	6.38× (−4.1%)

`15L:1A SLA Tree-CP` runs in 4,870.6 ms (215.3K tok/s/cube) with momentum and 4,778.1 ms (219.5K tok/s/cube) without momentum (4,737.6 / 4,904.0 ms with vocab-chunked cross-entropy), achieving a 1.09× speedup (1.11×–1.13× without momentum) over `15L:1A KDA-CP` (5,307.1–5,337.6 ms) with 6.31 GiB/core lower peak HBM (25.50 vs. 31.81 GiB/core), a 7.50×–7.75× (7.65× without momentum) speedup over `0L:1A` Ring-CP (36,530.6–36,725.0 ms), and 83% lower mixer ICI traffic than `15L:1A Ring-LASP` (5,204.6 ms, 2,992 vs. 17,392 MiB/rank). At $T = 524{,}288$ ($T_{\text{loc}} = 4{,}096$/core), it reaches 4,170 tok/s/chip (1,965 ms uncoupled, 5.45× vs. `0L:1A` Ring-CP 10,712.1 ms; 2,003.1 ms with momentum, 5.35×).

2. **End-to-End 12.5B (`7L:1A`) Pre-Training Step Time vs. Tokamax Mosaic KDA and MaxText KDA (Table 19):** Across the full 16-layer 12.5B (`7L:1A`) architecture (14 linear layers + 2 global GQA layers + 16 `LatentMoE` layers with $k = 8$ active experts + AdamW), fused 1-kernel `SLA-R4` is 1.68× (23.61 ms, $\mu = 0.25$) to 1.73× (22.85 ms, $\mu = 0$) faster on the 14-layer linear mixer F+B than open-source Tokamax Mosaic KDA (39.62 ms; 2.65×–2.74× vs. MaxText KDA 62.50 ms), delivers a 1.24× (66.01 vs. 82.15 ms, $\mu = 0.25$) to 1.26× (65.24 ms, $\mu = 0$) and

Table 19: Measured end-to-end 12.5B (`7L:1A`, 16L $=$ 14 Linear $+$ 2 GQA $+$ 16 LatentMoE $k =$ $8 +$ AdamW) training step-time benchmark (ms, median [IQR]) comparing `MaxText Chunk-KDA`, open-source `Tokamax Mosaic KDA` (`tokamax` 0.0.14), `Scalar GDN` ($G = 1$), and `SLA-R4` ($G = 4$, 2-stage and 1-kernel fused) on 256×256 MXU TPU v6e (1-chip active-parameter and 8-chip `EP=8` all-to-all, $T = 4{,}096$) and 64-chip TPU v7x (1-microbatch and GA $= 4$, $T = 8{,}192$, `-xla_tpu_scoped_vmem_limit_kib=65472`).

Linear Mixer Arm (12.5B `7L:1A`, 16L MoE + AdamW)	Gate G	14L Mixer F+B (ms, vs. Tok.)	12.5B Active Step (ms, vs. Tok./Max.)	8c EP=8 Step (ms, vs. Tok.)	64c v7x 12.5B (1-MB/GA=4, vs. Tok.)
MaxText Chunk-KDA (FP32 / BF16 state)	d_k	62.50 (0.63×) / 60.92	108.06 (0.76× / 1.00×)	227.00 / 154.56 (0.87×)	766.8 / 2,225 (0.87× / 0.92×)
Open-Source Tokamax Mosaic KDA (`tokamax` 0.0.14)	d_k	39.62 [0.06] (1.00×)	82.15 [0.11] (1.00× / 1.32×)	134.06 [0.63] (1.00×)	668.4 / 2,049 (1.00×)
Scalar GDN ($G = 1$, 2-Stage / 1-Kernel, loses recall)	1	31.03 / 21.23 (1.87×)	73.54 / 63.98 (1.28× / 1.69×)	132.76 / 120.89 (1.11×)	468.8 / 1,796 (1.43× / 1.14×)
SLA-R4 (2-Stage Chunk Kernel, $\mu = 0.25$)	4	33.02 [0.05] (**1.20×**)	75.44 [0.09] (**1.09× / 1.43×**)	132.80 [0.45] (1.01×)	**469.9 / 1,800 (1.42× / 1.14×)**
Fused SLA-R4 (1-Kernel Exact Inv., $\mu = 0.25$)	**4**	**23.61** [**0.02**] (**1.68×**)	**66.01** [**0.09**] (**1.24× / 1.64×**)	**121.59** [**0.30**] (**1.10×**)	**390.1 / 1,738 (1.71× / 1.18×)**
Fused SLA-R4 (1-Kernel Exact Inv., $\mu = 0$)	**4**	**22.85** [**0.06**] (**1.73×**)	**65.24** [**0.09**] (**1.26× / 1.66×**)	**120.79** [**0.36**] (**1.11×**)	**385.1 / 1,729 (1.74× / 1.19×)**
Fused SLA-R4 (1-Kernel 1-Step Neumann)	**4**	**19.51** [**0.03**] (**2.03×**)	**61.83** [**0.08**] (**1.33× / 1.75×**)	**118.25** [**0.34**] (**1.13×**)	**361.9 / 1,692 (1.85× / 1.21×)**

Table 20: Unified single-protocol TPU v7x benchmark (3 independent process restarts per point, `-xla_tpu_scoped_vmem_limit_kib=65472` [64 MiB VMEM], pre-allocated inputs, `tokamax` 0.0.14) across (a) single-layer full Ops 1–7c F+B latency (ms, $B = 1$) at $D = 1792$ and $D = 2048$ (including the fused 1-kernel $G = 1$ scalar-gate control isolating $G = 4$'s 2.5%–4.2% mask cost) and (b) 4-scale end-to-end sparse MoE training step time (ms, $T_{\text{train}} = 8{,}192$) from 6.5B to 63.1B.

(a) Single-Layer Full Ops 1–7c ($B = 1$)	Gate G	D=1792, 8K	D=1792, 32K	vs. KDA / Tok.	D=2048, 8K	D=2048, 32K	vs. KDA / Tok.
Open-Source MaxText Pallas KDA ($C = 128$, FP32)	d_k	10.42 ± 0.03	39.81 ± 0.14	1.00× / 0.85×	6.87 ± 0.01	26.69 ± 0.09	1.00× / 0.61×
Open-Source MaxText Pallas KDA ($C = 128$, BF16)	d_k	10.09 ± 0.02	38.67 ± 0.12	1.03× / 0.88×	6.55 ± 0.01	25.49 ± 0.08	1.05× / 0.64×
Open-Source Tokamax Mosaic KDA ($C = 64$, Base)	d_k	8.70 ± 0.01	34.01 ± 0.11	1.17× / 1.00×	4.26 ± 0.01	16.28 ± 0.07	1.64× / 1.00×
2-Stage SLA-R4 (`Exact Inv.`, $\mu = 0$ / 0.25)	4	5.38 / 5.50	19.90 / 20.43	2.00× / 1.71×	3.72 / 3.86	13.86 / 14.23	1.93× / 1.17×
1-Kernel SLA-R4 (Scalar $G = 1$ Control, $\mu = 0$)	**1**	$\mathbf{3.73 \pm 0.01}$	$\mathbf{14.05 \pm 0.05}$	**2.83× / 2.42×**	$\mathbf{2.63 \pm 0.01}$	$\mathbf{9.20 \pm 0.02}$	**2.90× / 1.77×**
1-Kernel SLA-R4 (`Exact Inv.`, $\mu = 0$ / 0.25)	**4**	**3.84 / 3.89**	**14.06 / 14.41**	**2.76× / 2.36×**	**2.74 / 2.81**	**9.59 / 9.83**	**2.78× / 1.70×**
1-Kernel SLA-R4 (`1-Step Neumann`, $\mu = 0.25$)	**4**	$\mathbf{3.57 \pm 0.01}$	$\mathbf{12.93 \pm 0.04}$	**3.08× / 2.63×**	$\mathbf{2.39 \pm 0.01}$	$\mathbf{8.16 \pm 0.02}$	**3.27× / 2.00×**
(b) 4-Scale E2E Sparse MoE Step Time (ms)	Gate G	6.5B (10L)	12.5B (16L)	31.5B (15L)	63.1B (30L)	Spd. vs. KDA	Spd. vs. Tokamax
MaxText Pallas KDA ($C = 128$, FP32 state)	d_k	310.2	385.1	601.9	1,243.2	1.000×	0.865×–0.992×
Open-Source Tokamax Mosaic KDA ($C = 64$)	d_k	272.6	348.5	572.7	1,232.9	1.008×–1.138×	1.000×
2-Stage SLA-R4 (`Exact Inv.`, BF16)	4	269.5	371.9	511.6	1,050.8	1.035×–1.183×	0.937×–1.173×
1-Kernel SLA-R4 (`Exact Inv.`, BF16)	**4**	**262.2**	**330.4**	**495.2**	**1,015.7**	**1.166 ×–1.224×**	**1.040 ×–1.214× (+4.0%–21.4%)**
1-Kernel SLA-R4 (`1-Step Neumann`, BF16)	**4**	**260.4**	**325.4**	**490.8**	**1,006.8**	**1.184 ×–1.235×**	**1.047 ×–1.225× (+4.7%–22.5%)**

1.33× (61.83 ms, 1-step Neumann) end-to-end 12.5B step speedup over Tokamax Mosaic KDA on single-chip active-parameter steps (1.64×–1.75× vs. MaxText KDA 108.06 ms), delivers 1.10×–1.11× (121.59/120.79 ms vs. 134.06 ms) on 8-chip `EP=8` with MoE all-to-all (1.87× vs. MaxText FP32 227.00 ms), and achieves a 1.14× (1,800 vs. 2,049 ms at GA $= 4$) to 1.42× (469.9 vs. 668.4 ms single-microbatch; 1.18×–1.74× with 1-kernel $\mu = 0$) end-to-end 12.5B speedup over Tokamax Mosaic KDA on 64-chip TPU v7x (alongside 1.24×–1.63× over MaxText KDA).

3. **Full 64-Chip `EP=128` MoE Cube Serving and `FP8Fused` Scan (Table 21b):** At $B_{\text{chip}} = 64$ (4,096 sequences/cube), our fused Pallas E4M3 kernel (`FP8Fused`) keeps row-scaled E4M3 DMA, dequantization, the 16-token delta scan, and requantization inside VMEM. On 16-step sequential decode (`seq_scan_t16`) at 8K, this lifts `15L SLA Iso-State` to 1,091.9K tok/s/cube (60.02 ms; +18.8% vs. `15L KDA FP32` 919.4K, +14.5% vs. JAX `15L KDA FP8` 953.6K, and +14.6% vs. pure-JAX SLA FP8 952.9K) and `15L SLA-R2` to 1,047.2K–1,052.3K tok/s/cube (+10.3% vs. JAX `KDA FP8`, +28.1% vs. pure-JAX SLA FP8 817.6K). Applying the same `FP8Fused` recurrent scan kernel to KDA reaches 1,083.0K (within 0.8% of `SLA Iso-State`, as step-by-step recurrent decode has no chunk WY solve). With parallel `Tree-Neumann-1` ($T_{\text{ver}} = 16$), raw verification reaches 1.121M–1.158M tok/s/cube at 8K and sustains 1.037M–1.042M tok/s/cube in 16.24–16.47 GiB/chip (−93.6% vs. `0L:1A` OOM) at 131K.

4. **Real-Checkpoint Lossless Speculative Acceptance incl. Doubly-Clamped KDA (Table 21c):** Under exact lossless rejection [39] (draft length $K = 16$, $N = 64$ prompts/suite), zero-clamp `0.35B SLA` → `12.5B 7L SLA-R4` (iso-param, 60.4B) lifts mean accepted tokens $\bar{\tau}$ under greedy decoding on 32K/64K `HotpotQA` from 2.63/2.14 (unclamped KDA) to 5.86/4.80 (2.23 × /2.24×; 5.78/4.66 in FP8 PTQ). Against *doubly-clamped* KDA (`0.35B KDA TrainMaxTau=8K` → `12.5B 7L KDA` with $\tau_{\max} = 8$K, Appendix F: 4.42/4.09 greedy, 1.72/1.65 stochastic), SLA remains 1.33 × /1.17× higher greedy and 1.06 × /1.10× higher stochastic on OOD 32K/64K `HotpotQA` (379.8K vs. 287.6K greedy and 118.0K vs. 111.9K stochastic accepted `tok/s/cube` at 32K), while KDA edges in-distribution 8K Dolma (4.84–5.30 vs. 4.56).

Table 21: Measured 64-chip TPU v7x (4x4x4, 128 TensorCores) end-to-end 16-layer 12.5B sparse MoE (0.79B active, 256 experts $k = 8$) (a) fixed-CP=EP=128 F+B latency and tok/s/chip across $T = 128$K–1.05M ($T_{\text{loc}} = T/128$; at $T = 512$K, $T_{\text{loc}} = 4{,}096$ selects a faster XLA EP=128 MoE ragged-all-to-all tile schedule across all linear hybrids relative to $T = 256$K), (b) EP=128 cube verification throughput ($B_{\text{chip}} = 64$, comparing Tree-N1 [= Tree-Neumann-1, $T_{\text{ver}} = 16$] and FP8Fused seq_scan_t16), and (c) real-checkpoint lossless speculative acceptance $\bar{\tau}$ ($K = 16$, $N = 64$, **greedy [stochastic]**).

(a) 12.5**B MoE** CP=EP=128 **F+B**	$T = 128$K (tok/s/ch)	$T = 256$K (tok/s/ch)	$T = 512$K (tok/s/ch)	$T = 1.05$M (tok/s/ch)	1M HBM / Mixer ICI
0L:1A Full-Attn Ring-CP	1,740.4 ms (1,177)	4,797.5 ms (854)	10,712.1 ms (765)	36,725.0 ms (446)	25.8 GiB / 32,512 MiB
7L:1A MaxText KDA Tree ($G = 128$)	1,232.2 ms (1,662)	2,573.3 ms (1,592)	2,704.5 ms (3,029)	7,377.0 ms (2,221)	25.73 GiB / 4,232 MiB
7L:1A **Iso-Param SLA Tree-CP R4**	**1,212.9** ms (**1,689**)	**2,500.8** ms (**1,638**)	**2,542.6** ms (**3,222**)	**7,023.9** ms (**2,333**)	25.92 GiB / 4,736 MiB
15L:1A MaxText KDA Tree ($G = 128$)	1,221.4 ms (1,677)	2,498.2 ms (1,639)	2,249.8 ms (3,641)	5,337.6 ms (3,070)	25.49 GiB / 2,760 MiB
15L:1A Ring-State LASP R2	1,575.5 ms (1,299)	2,744.8 ms (1,492)	2,346.3 ms (3,491)	5,204.6 ms (3,148)	25.46 GiB / 17,392 MiB
15L:1A **SLA Tree-CP R2 NoMom / +Mom**	**1,174**/**1,189** ms	**2,331**/**2,368** ms	**1,965**/**2,003** ms (**5.45**× / 5.35×)	**4,738**/**4,904** ms (**3,458**/**3,341**)	**25.45** GiB/**2,992** MiB (−**83**%)
(b) 12.5**B MoE** EP=128 **Cube Decode**	8K, $B_{\text{chip}} = 64$ (/cube)	32K, $B_{\text{chip}} = 64$ (/cube)	131K, $B_{\text{chip}} = 64$ (/cube)	8K seq_scan_t16 (/cube)	131K KV+State / Chip
0L:1A Full-Attn GQA (Tree-N1, T_{ver}=16)	15,381 /ch (984.4K)	10,677 /ch (683.4K)	OOM ($B_{\text{chip}} = 8$: 250.7K)	n/a (softmax)	256.0 GiB (OOM)
7L:1A KDA Base FP32 (Tree-N1, T_{ver}=16)	17,212 /ch (1.102M)	16,258 /ch (1.041M)	13,106 /ch (838.8K)	14,180 /ch (907.5K)	32.88 GiB
7L:1A **SLA Iso-State BF16/FP8Fused**	**17,513** /ch (**1.121M**)	**16,544** /ch (**1.059M**)	**13,255** /ch (**848.3K**)	**16,628** /ch (**1,064.2K**)	**32.44**/**32.22** GiB
15L:1A KDA Base FP32 (Tree-N1, T_{ver}=16)	17,513 /ch (1.121M)	17,140 /ch (1.097M)	15,912 /ch (1.018M)	14,366 /ch (919.4K)	16.94 GiB
15L:1A KDA Base FP8 (JAX / FP8Fused)	18,089 /ch (1.158M)	17,792 /ch (1.139M)	16,285 /ch (1.042M)	14,901 /ch (953.6K/1,083K)	16.24 GiB
15L:1A **SLA Iso-State FP8Fused** (T_{ver}=16)	**17,960** /ch (**1.149M**)	**17,510** /ch (**1.121M**)	**16,220** /ch (**1.038M**)	**17,061** /ch (**1,091.9K**, +**18.8**%)	**16.23** GiB (−**93.7**%)
15L:1A **SLA-R2 FP8Fused** (T_{ver}=16)	**17,948** /ch (**1.149M**)	**17,499** /ch (**1.120M**)	**16,200** /ch (**1.037M**)	**16,363** /ch (**1,047.2K**, +**13.9**%)	**16.47** GiB (−**93.6**%)
15L:1A **SPD-SLA-W Polar-2+ FP8** (T_{ver}=16)	**17,994** /ch (**1.152M**)	**17,009** /ch (**1.089M**)	**16,210** /ch (**1.037M**)	12,801 /ch (819.3K JAX)	**16.47** GiB (−**93.6**%)
(c) Lossless Speculative $\bar{\tau}$ ($K = 16$)	Dolma 8K $\bar{\tau}$	Dolma 32K $\bar{\tau}$	SQuAD 32K/64K $\bar{\tau}$	HotpotQA 32K/64K $\bar{\tau}$	Accepted Hotpot@32K / Cube
0.35B KDA → 12.5B 7L KDA (raw)	4.84 [2.53]	3.75 [1.94]	3.72/3.55 [1.49/1.19]	2.63/2.14 [1.36/1.22]	170.7K [88.4K] tok/s
0.35B KDA ($\tau_{\max}$=8K) → 12.5B KDA ($\tau_{\max}$=8K)	**5.30** [**2.49**]	5.14 [3.04]	**4.62**/4.11 [1.65/1.56]	4.42/4.09 [1.72/1.65]	287.6K [111.9K] tok/s
0.35**B SLA** → 12.5**B** 7L **Iso-SLA** (60.4**B**)	4.56 [2.43]	**5.27** [**3.18**]	4.44/**4.20** [**1.69**/**1.71**]	**5.86**/**4.80** [**1.82**/**1.81**] (**1.33**×)	**379.8K** [**118.0K**] (**1.32**×)
0.35**B SLA** → 12.5**B** 7L **Iso-SLA (FP8 PTQ)**	4.67 [2.44]	5.14 [**3.18**]	4.45/4.13 [**1.69**/**1.70**]	**5.78**/**4.66** [**1.82**/**1.82**]	**373.9K** [**117.7K**]
0.35**B SLA** → 63.1**B** 4L **SLA-R4** (120.8**B**)	4.20 [2.46]	2.75 [1.91]	3.66/3.47 [1.71/1.78]	3.17/2.36 [1.32/1.25]	n/a (512c; 10.7× draft span)

F Extended Pre-Training XProf, Causal τ_{eff} Clamp Audit, 4-Rung Ladder, and 9-Task `lm-eval`

This section collects the hardware `XProf` HLO breakdowns (Tables 22–23), the multi-scale per-channel retention-horizon (τ_{eff}) dispersion and clamping audit (Table 24), the complete 45-row 4-rung pre-training and 8K $\rightarrow$ 32K extension ladder (Table 25), and the 39-row per-subtask 9-task `lm-eval` and 8K $\rightarrow$ 32K $\rightarrow$ 64K `RULER` QA breakdown (Table 26) summarized in Tables 7–8.

Hardware XProf operator attribution on 64-chip TPU v7x. XProf attributes most of SLA's step-time saving to linear attention ops (Figure 7a). Table 22 provides the complete hardware `XProf` HLO category breakdown (`xplane` traces captured over steps 15–17 across all 8 TensorCores of host 0, normalized by 24 core-steps to ms/step) for 63.1B MoE (`FSDP=2, EP=64`) on 64 TPU v7x chips. Mapping directly to Figure 7a (where the top bar groups ICI `async-done` wait with MoE/dense matmuls: $366.6+245.8 = 612.4$ ms, $354.1+266.8 = 620.9$ ms, and $343.1+237.2 = 580.3$ ms), the three direct linear attention op categories save -218.7 ms (87.2%) and reduced mixer memory/sync traffic saves -23.5 ms (9.4%) of exposed ICI `async-done` wait, totaling -242.2 ms (96.6%) of the -250.8 ms step-time reduction between `SLA-R4` and KDA Base (with -8.6 ms [3.4%] in MoE/FFN matmul overlap).

Table 23 complements this with a symmetric 4-arm `GA=4` (8.39M tok/step) `XProf` audit on 12.5B MoE (`EP=128`): deferring causal-conv weight `AllReduce` outside the `GA=4` microbatch loop yields 0.0% speedup on MaxText KDA (2,118.9 $\rightarrow$ 2,119.7 ms) while cutting `SLA-R4` (iso-param) to 1,686.0 ms in 3-step `XProf` (1,661.9 ms unprofiled min; $1{,}812.0 \pm 14.1$ ms over full 7,200-step $n = 2$ training runs due to host checkpoint/validation wrapper hooks) via a $1.96\times$ faster Pallas linear attention kernel (528.4 $\rightarrow$ 269.1 ms).

Table 22: Hardware `XProf` per-step HLO category breakdown (ms/step) on 64-chip TPU v7x 63.1B MoE (`FSDP=2, EP=64`, 2.10M tok/step).

HLO Category (`xplane` Measured)	`4L:1A` KDA Base	`4L:1A` SPD-SLA-W (R2)	`4L:1A` SLA-R4	Δ (SLA-R4 vs. KDA)
Pallas Linear Attention `custom-call`	703.0 ms	597.2 ms	552.0 ms	-151.0 ms (-21.5%)
VPU Gate / Norm `loop fusion`	190.4 ms	156.8 ms	143.7 ms	-46.7 ms (-24.5%)
Causal Short `convolution fusion`	220.0 ms	207.4 ms	199.0 ms	-21.0 ms (-9.5%)
ICI Collective `async-done` Wait	366.6 ms	354.1 ms	343.1 ms	-23.5 ms (-6.4%)
MoE Expert `ragged-dot` + Dense Matmul	245.8 ms	266.8 ms	237.2 ms	-8.6 ms (-3.5%)
Total Steady-State Step Time	**1,725.8** ms	**1,582.3** ms	**1,475.0** ms	**-250.8** ms (**1.170**$\times$)
Steady-State Throughput (`tok/s/chip`)	18,988	20,710 (1.09$\times$)	**22,216** (**1.17**$\times$)	**+3,228** (**+17.0**%)

At 6.3B tokens (RoPE validation checkpoint on the same 64-chip slice), `4L:1A` KDA Base and `4L:1A SLA-R4` reach 9.90 vs. 10.15 8K Dolma PPL, 10.27 vs. **9.75** 32K Tail PPL, and 14.5% vs. **37.5**% 32K 8-Needle accuracy.

Table 23: Symmetric 4-arm `GA=4` (8.39M tok/step) hardware `XProf` leaf-op breakdown (ms/step) on 64-chip TPU v7x 12.5B MoE (`EP=128`).

12.5B `GA=4` Arm (3 profiled steps)	Step Mean (ms)	`linear_mixer _pallas`	`short_conv _fusion`	Mixer + Conv	`full_attn _splash`	`all_gather _rs`	`moe_ragged _a2a`
`7L:1A` MaxText KDA (default AR)	2,118.9	508.1	91.7	599.8	69.8	191.9	264.7
`7L:1A` MaxText KDA + `defer_ar`	2,119.7 (1.00$\times$)	528.4	91.6	620.0	70.0	191.6	234.3
`7L:1A` SLA (Fused) R2+Mom + `defer_ar`	**1,797.4** (1.18$\times$)	**340.9** (-187.5)	**65.4**	**406.3** (-213.7)	69.9	**152.5**	282.5
`7L:1A` Iso-Param SLA (Fused) R4 + `defer_ar`	**1,686.0** (1.26$\times$)	**269.1** (1.96$\times$)	**61.7**	**330.8** (-289.2)	69.7	**152.6**	265.3

Causal τ_{eff} dispersion and NoPE `SSMax`. Table 24 audits per-channel retention horizons

$$\tau_{\text{eff}}(i) = \mathbb{E}_t\left[-\frac{1}{\ln f_{t,i}}\right]$$

(in tokens, for channel i with learned decay $f_{t,i}$) and Within-Tile Std (the RMS standard deviation of token-mean log-decay $\mathbb{E}_t[\ln f_{t,i}]$ across channels i within each $d_k/4$ tile block, equal to 0 when $G \in \{1, 4\}$) alongside zero-retraining inference interventions across 6.5B, 12.5B, and 63.1B MoEs.

Across all three scales, $G = d_k$ KDA develops 0.8%–1.3% ultra-slow $\tau_{\text{eff}}(i) \approx 131$K channels (within-tile log-decay RMS std = 3.84–4.97) that inject OOD state spikes into hybrid layers. Clamping $\tau_{\text{eff}}(i) \leq \tau_{\max} = 8{,}192$ at inference with zero retraining cuts 63.1B KDA 32K Tail PPL by $45.9\times$

Table 24: Multi-scale causal τ_{eff} dispersion, zero-retraining $\tau_{\max} = 8{,}192$ inference clamping, live `TrainMaxTau=8K` saturation, post-hoc $G = 4$ projection, and NoPE `SSMax` ($\alpha = 0.5$ on SLA, $\alpha = 0.25$ on KDA) on 32K Tail PPL, 32K 8-N, and 64K `RULER` QA ($N = 128$ unless noted N=512). The clamp evaluator (`a3_causal_clamp_eval`) scores its own held-out Dolma slice and paired $N = 128$ QA subset, so absolute Tail PPL and QA values differ from the standalone scorers of Tables 8, 7, 25, and 26 (by 1–2 QA items, 0.8–1.6 pts) while within-row intervention deltas remain exact.

Scale, Horizon, & Trained Checkpoint	G	Within-Tile Std	Unclamped 32K Tail / 8-N	Intervened 32K Tail / 8-N (Proj. G=4)	Unclamped 64K `SQuAD` / `Hotpot`	Intervened 64K QA (Proj. G=4 / Note)
6.5B `4L:1A` MaxText KDA (15.1B)	128	4.41	14.64/14.1%	**11.77**/**69.9**% (70.47)	53.1%/42.2%	**75.0**%/**79.7**% (24.2/39.8%)
6.5B `4L:1A` KDA + TrainMaxTau (15.1B)	128	$\leq \tau_{\max}$	**12.36**/**48.4**%	**12.36**/**48.4**% (co-capped)	**69.5**%/**73.4**%	**69.5**%/**73.4**% (68.8/70.3% N=512)
6.5B `4L:1A` Iso-Param SLA-R4 (15.1B)	4	**0.00**	**11.94**/**48.8**%	**11.94**/**48.8**% (11.94)	**67.2**%/**76.6**%	**67.2**%/**76.6**% (invariant)
12.5B `7L:1A` MaxText KDA (60.4B, `s0`)	128	4.58	12.35/15.2%	**10.26**/**57.8**% (57.85)	47.7%/44.5%	**78.1**%/**75.0**% (28.1/45.3%)
12.5B `7L:1A` KDA + TrainMaxTau (60.4B, $n = 2$)	128	$\leq \tau_{\max}$	**10.75**/**49.6**%	**10.75**/**49.6**% (co-capped)	**74.5**%/**72.8**% (N=512)	**74.5**%/**72.8**% (N=512)
12.5B `15L:1A` MaxText KDA (60.4B)	128	4.97	10.48/15.2%	**9.66**/**42.2**% (63.30)	59.4%/61.7%	**73.4**%/**78.9**% (25.8/34.4%)
12.5B `7L:1A` **Iso-SLA (Fused) R4** (60.4B, `s0`)	**4**	**0.00**	**9.80**/**63.7**%	**9.80**/**63.7**% (**9.80**)	**73.4**%/**77.3**%	**75.0**%/**77.3**% (+1.6/0.0 pts)
63.1B `4L:1A` MaxText KDA (120.8B, NoPE)	256	3.84	353.42/15.6%	**7.70**/**48.4**% (78.06)	34.4%/25.8%	**76.6**%/**82.0**% (29.7/25.0%)
$+\tau_{\max} = 8\text{K}$ + `SSMax` ($\alpha = 0.25$, N=512)	256	$\leq \tau_{\max}$	380.66/18.8% (`SSMax` only)	**8.04**/**74.2**% ($\tau_{\max}$+`SSMax`)	27.3%/27.0% (`SSMax` only)	**90.6**%/**86.5**% (77.5/76.4% norm)
63.1**B** `4L:1A` **SLA-R2** (120.8**B, NoPE**)	**4**	**0.00**	**13.78**/18.6%	**12.07**/21.9% ($\tau_{\max}$ only)	**57.0**%/**53.1**%	**61.7**%/**60.9**% (55.7/56.8% N=512)
63.1**B** `4L:1A` **SLA-R2 +** `SSMax` (120.8**B**)	**4**	**0.00**	13.78/18.6%	**8.45**/**84.4**% (`SSMax`, 0 clamp)	57.0%/53.1%	– (zero linear clamp)
63.1**B** `4L:1A` **SLA-R4 +** `SSMax` (120.8**B**)	**4**	**0.00**	22.78/18.0%	**8.41**/**83.2**% (`SSMax`, 0 clamp)	51.6%/45.1% (N=512)	**75.6**%/**79.1**% (62.3/66.6% norm)

(353.42 → 7.70), lifts 32K 8-Needle across every scale (14.1% → 69.9% at 6.5B, 15.2% → 57.8% at 12.5B `7L`, 15.2% → 42.2% at 12.5B `15L`, 15.6% → 48.4% at 63.1B), and recovers +14.0 to +56.2 pts on 64K QA; by contrast, naive contiguous $G = 4$ tile-mean projection without retraining collapses (57.8–78.1 Tail PPL) because KDA channels are unordered across the head, whereas sorting channels by learned τ_{eff} (co-permuting W_q, W_k, and `Conv1D`) and partitioning into $G = 4$ contiguous tiles by dynamic programming cuts $\log \tau$ and 1K-lag retention RMSE by 3.7× (1.97 → 0.53 and 0.165 → 0.044 at 12.5B with DP widths $91, 23, 9, 5$; 1.99 → 0.56 and 0.153 → 0.053 at 63.1B with DP widths $178, 40, 28, 10$). Live training-time softplus saturation (`TrainMaxTau=8K`) likewise stabilizes OOD retrieval at both 6.5B (69.5/73.4% QA) and 12.5B (60.4B tok, $n = 2$: 10.75 Tail, 49.6% 32K 8-N, 74.5/72.8% $N = 512$ QA) while retaining $G = 128$'s VPU gate tax and 4× FP32 state.

Conversely, on 63.1B `SLA-R4` (120.8B tok, NoPE), 24-layer state-norm telemetry confirms linear states contract cleanly from 8K to 32K ($\|S_t\|_F = 2.04 \to 1.62$, max row norm 1.59 → 1.15, 0.00% rows near E4M3 448), and rescaling each global softmax query q on the 6 NoPE full-attention layers at evaluation sequence length T via zero-retraining `SSMax` [45],

$$q \leftarrow q\left(1 + \alpha \ln\left(\frac{T}{T_{\text{train}}}\right)\right)$$

(with $T_{\text{train}} = 8{,}192$ and SSMax scale $\alpha = 0.5$, unrelated to α_{safe}; zero linear clamp), cuts 32K Tail PPL to 8.45 (`SLA-R2`) and 8.41 (`SLA-R4`) and restores 32K 8-Needle to 84.4% (`SLA-R2`) and 83.2% (`SLA-R4`). By contrast, linear $\tau_{\max} = 8{,}192$ clamping without `SSMax` moves 63.1B `SLA-R2` 32K 8-Needle only from 18.6% to 21.9%, and `SSMax` alone on unclamped KDA still diverges (380.66 Tail PPL, 18.8% 8-Needle), cleanly separating global NoPE softmax dilution from KDA's linear gate drift.

G Sorted-DP KDA Conversion, Woodbury Tile-RLS Boundary, and Pretrained Checkpoint Co-Design

Table 27 summarizes our extended TPU v7x verification of: (1) zero-shot and fast-healed conversion of per-channel KDA into 4-tile SLA via permutation-equivariant dynamic programming (`Sorted_DP_4Tile_SLA`), (2) intra-chunk Woodbury Recursive Least Squares (`Woodbury_RLS`) vs. `SLA_4Tile` and damped Neumann updates, and (3) real-checkpoint co-design across pretrained GatedDeltaNet (`Qwen3-Next-80B`, `Qwen3.5-0.8B`), per-channel KDA (`Kimi-Linear-48B`, `KDA-Hybrid-1.3B`), and Softmax GQA (`Qwen3-0.6B`) checkpoints:

1. **Sorted-DP KDA-to-SLA Conversion on TPU v7x ($n = 3$ Paired Seeds):** Starting from 8L per-channel KDA teachers (154.71 FineWeb PPL / 186.26 OOD PPL at 98.86 ms/step), sorting key channels by mean log-decay (co-permuting W_q, W_k, and `Conv1D`, exact algebraic permutation max logit diff 0.0741) and partitioning into G=4 tiles cuts within-tile decay SSE by **4.22**× (48.75 → 11.55; 11.78× unconstrained DP to 4.14). On 0-shot FineWeb-Edu evaluation, `Sorted_DP_4Tile_SLA` achieves **173.16** PPL (230.45 OOD) vs. **224.02**

Table 25: Matched RoPE ($\theta_{\max} = 5 \times 10^5$) 4-rung scaling ladder (0.35B to 63.1B MoE; see Table 8 for 12.5B Rung 2), 0.35B 3-seed & 30.2B (4.4× Chinchilla), 31.5B at 36.2B tok (1.0× Chinchilla), 12.5B 8K → 32K continued pre-training, and 12.5B `15L:1A` dilution control on 64-chip TPU v7x ([§]under `remat_32k` with 4-step gradient accumulation [`GA=4`], KDA/`0L:1A` hit `NaN` at steps 1–43 and required reruns without rematerialization, whereas 2-seed `SLA-R4` (iso-param) completed all 4,300 steps at 3,040 ms with 12.54 → 12.08 Tail PPL, 59.2%/48.2% 32K 4/8-N, and 58.5/62.6% N=512 64K QA, though router-remat spikes raised 8K PPL to 17.99).

MoE Scale (Active / Total, Tokens)	Architecture	Step (ms)	11-Set PPL (or 32K-Ext CE)	8K Tail PPL	32K Tail PPL	32K 4-N / 8-N	8-N Gain vs. KDA
0.35B/6.5B (15.1B tok, 2.2× Chin., `4L:1A` interval 5, `s0`)	`0L:1A` GQA+RoPE (LR = 4.5×10^{-4})	361.6	19.582	12.83	13.18	18.4%/12.9%	0.91×
	`4L:1A` MaxText KDA Base (FP32)	351.8	**18.573**	**12.42**	15.72	32.0%/14.1%	1.00×
	`4L:1A` KDA + TrainMaxTau ($\tau_{\max}$ = 8K, FP32)	351.2	19.035 (`s1`; 2.54 CE)	12.43	**12.36**	**56.6%/48.4%**	**3.43×**
	`4L:1A` **SLA-R2** (BF16)	**320.0**	18.773	12.62	**12.30**	**63.3%/55.1%**	**3.91×**
	`4L:1A` **SLA (Fused) R2+Mom** (BF16)	**308.5**	18.890	12.60	**12.28**	**60.9%/51.6%**	**3.66×**
	`4L:1A` **Iso-Param SLA-R4** (BF16)	**302.4**	19.093	12.83	**12.51**	**57.8%/48.8%**	**3.46×**
	`4L:1A` **SPD-SLA-W** (BF16)	344.8	19.104	12.64	**12.32**	**66.4%/56.2%**	**3.99×**
0.35B/6.5B 3-Seed & 30.2B Tok (4.4× Chinchilla)	3-Seed `4L:1A` MaxText KDA (15.1B, $n = 3$)	352.0	18.880 ± 0.267	**12.59**	15.89	32.9%/19.9%	1.00×
	3-**Seed** `4L:1A` **SLA-R2** (15.1**B,** $n = 3$)	**314.8**	19.077 ± 0.265	12.72	**12.77**	**49.9%/37.6%**	**1.89×**
	`0L:1A` GQA+RoPE (30.2B, 4.4× Chin.)	361.5	17.104	11.43	12.02	25.0%/15.2%	0.90×
	`4L:1A` MaxText KDA (30.2B, 4.4× Chin.)	351.5	**16.161**	11.03	620.87	31.6%/16.8%	1.00×
	`4L:1A` **SLA-R2** (30.2**B,** 4.4× **Chin.)**	**315.0**	16.250	**11.02**	**13.53**	35.2%/21.1%	1.26×
	`4L:1A` **Iso-Param SLA-R4** (30.2**B,** 4.4× **Chin.)**	**297.9**	16.326	11.15	**11.09**	**59.4%/45.3%**	**2.70×**
1.80B/31.5B (5.03B & 36.2B tok, `4L:1A`)	`0L:1A` GQA+RoPE (5.03B, LR = 4.5×10^{-4})	867.5	**20.380**	**12.98**	13.73	24.6%/14.8%	1.15×
	`4L:1A` MaxText KDA Base (5.03B, FP32)	807.7	21.228	13.99	47.16	28.9%/12.9%	1.00×
	`4L:1A` **SLA (Fused) R2+Mom** (5.03B, BF16)	**747.2**	21.439	14.19	**13.88**	**48.8%/33.2%**	**2.57×**
	`4L:1A` **SLA (Fused) R4+Mom** (5.03B, BF16)	**707.4**	21.709	14.44	**14.20**	**59.8%/46.1%**	**3.57×**
	`4L:1A` MaxText KDA (36.2B, 4,300 st., no remat)[§]	3,526	15.264	10.57	51.68	24.6%/12.5%	1.00× (39.1/34.4% QA)
	`4L:1A` KDA + TrainMaxTau (36.2B, $\tau_{\max}$ = 8K, no remat)[§]	3,307	**15.211**	**10.36**	**16.77**	**27.3%/14.8%**	**1.18× (59.6/53.3%** N=512 QA)
3.74B/63.1B (5.03B tok, RoPE)	`0L:1A` GQA+RoPE (LR = 4.5×10^{-4})	1,749.6	**19.373**	**11.66**	11.94	25.0%/11.7%	0.81×
	`4L:1A` MaxText KDA Base (FP32)	1,746.0	20.076	12.25	18.24	23.1%/14.4%	1.00×
	`4L:1A` GDN-2 + RoPE (FP32, `EP=128`, 14,143 tok/s/ch)	2,316.9	19.731	12.27	14.22	25.8%/13.3%	0.92×
	`4L:1A` **SLA (Fused) R2+Mom** (BF16, 20,386 tok/s/ch)	**1,607.4**	20.201	12.37	**12.47**	**35.2%/22.3%**	**1.55×**
	`4L:1A` **SLA (Fused) R4+Mom** (BF16, 20,847 tok/s/ch)	**1,571.8**	20.423	12.50	**12.80**	**27.0%/13.7%**	0.95×
	`15L:1A` **SLA (Fused) R2+MS** (BF16)	**1,547.2**	20.890	13.17	**13.01**	**59.4%/48.4%**	**3.36×**
	`15L:1A` **SPD-SLA-W+MS** (BF16, 20,473 tok/s/ch)	**1,600.6**	20.506	12.49	**12.42**	**41.4%/27.0%**	**1.88×**
0.79B/12.5B 8K → 32K Ext. (5.0–10.1B at 32K)	`0L:1A` GQA+RoPE (32K-ext, 1.00× speed)	4,697.0	2.2089 CE	11.29	10.95	25.4%/11.7%	0.18×
	`7L:1A` MaxText KDA (32K-ext, 1.53× vs. `0L`)	3,061.4	2.2006 **CE**	11.20	**10.78**	71.5%/63.3%	1.00×
	`7L:1A` **SLA-R2** (32**K-ext,** 1.66× **vs.** `0L`)	**2,828.0**	2.2048 CE	11.26	10.82	**73.8%/64.1%**	**1.01×**
	`7L:1A` **Iso-Param SLA-R4** (32**K-ext,** 1.77×)	**2,661.0**	2.2255 CE	11.46	11.01	60.9%/49.6%	0.78×
	`7L:1A` MaxText KDA (60.4B → 32K, +5.03B)	3,103.6	2.0907 **CE**	**10.38**	**9.98**	55.5%/47.3%	1.00×
	`7L:1A` **Iso-SLA (Fused) R4** (60.4**B** → 32**K,** 1.15×)	**2,688.7**	2.1142 CE	10.60	10.15	**64.8%/57.0%**	**1.21×**
	`15L:1A` MaxText KDA (32K-ext, uniform)	2,827.0	2.2092 CE	**11.10**	10.79	56.2%/40.6%	1.00×
	`15L:1A` SLA-R2+MS / SPD-SLA-W+MS (32K-ext)	**2,608 / 2,715**	2.238/2.241	11.38/11.42	11.17/11.18	23.8%/16.0, 12.9%	0.39×
0.79B/12.5B (15.1B & 60.4B `15L:1A` dilution controls)	`15L:1A` MaxText KDA Base (15.1B, FP32, $n = 3$)	556.4	17.658 ± 0.297	**11.64**	11.57	34.0%/21.4%	1.00×
	`15L:1A` **SLA-R2** (15.1B, uniform, BF16, $n = 3$)	**486.7 ± 11.7**	**17.609 ± 0.246**	11.70	**11.42**	**45.1%/34.8%**	**1.63×**
	`15L:1A` **SPD-SLA-W Polar-2+** (15.1B, uniform, `s0`)	509.3	**17.575**	11.65	**11.38**	**50.0%/39.8%**	**1.86×** (1.29× `s0`)
	`15L:1A` SLA-R2+MS (15.1B, staggered, $n = 3$)	492.2	17.882 ± 0.294	11.93	11.72	23.6%/13.0%	0.61×
	`15L:1A` KDA + `TrainMaxTau=8K` (60.4B, FP32, `s0`)	2,199.0	15.579	10.47	**10.30**	37.5%/26.6%	1.75×
	`15L:1A` **SLA-R2+MS** (60.4B, staggered, `s0`)	**1,934.0**	**15.443**	10.61	**10.38**	27.3%/19.9%	1.31×
	`15L:1A` SPD-SLA-W+MS (60.4B, staggered, `s0`)	2,006.0	15.624	10.67	10.49	23.1%/12.1%	0.79×
0.79B/12.5B (15.1B, additional controls & 128c TPU v4)	`7L:1A` Scalar GDN ($G = 1$, FP32, $n = 2$)	555.9	18.142	11.85	**11.51**	**– / 60.8%**	**4.50×**
	`7L:1A` KDA + `TrainMaxTau=8K` (FP32, `s0`)	560.6	17.924	11.66	11.60	58.2%/48.1%	3.56×
	`Pure-Linear` (∞:1) KDA / **SLA-R2+MS** (`s0`)	545.8 / **475.6**	18.096 / **17.642**	12.02 / 11.95	12.05 / **11.80**	n.r. (55.5/59.4% `SQuAD`)	n.r. (56.2/53.9% `Hotpot`)
	TPU v4 (128**c**) `7L:1A` **SLA-R4** (NoPE, $n = 2$)	**992.0** (1.10×)	−0.017 nats vs. KDA	10.18	**10.05**	**– / 66.4%**	**2.42×**

Table 26: Zero-shot 9-task `lm-eval` accuracy (11,793 real validation items; `norm9` macro-average using byte-length normalized `acc_norm` on multiple-choice tasks and exact-word `acc` on `WinoGrande`/`LAMBADA`, with zero-shot row-scaled FP8 state `norm9` in parentheses for SLA rows; 1 SE ≈ 0.65 pts) and real-text 4-candidate `RULER` QA accuracy (`SQuAD` 1-hop & `HotpotQA` 2-hop, $N = 128$/length, 25% chance, 1 SE ≈ 4.0 pts) across 8K $\rightarrow$ 32K $\rightarrow$ 64K against `0L:1A` GQA+RoPE and MaxText KDA baselines.

Model, Horizon, & Mixer Architecture	HS	ARC-E/C	PIQA	WG	SciQ	MMLU	BoolQ	LAMB	`lm-eval` (FP8)	`SQuAD` / `HotpotQA` (8K $\rightarrow$ 32K $\rightarrow$ 64K)
6.5B `0L:1A` GQA+RoPE (15.1B, BF16)	34.8	45.8 / 25.8	63.8	50.7	71.7	28.0	62.2	33.4	46.24%	64.8 $\rightarrow$ 54.7 $\rightarrow$ 56.2% / 67.2 $\rightarrow$ 52.3 $\rightarrow$ 50.0%
6.5B `4L:1A` MaxText KDA (15.1B, FP32)	35.4	45.7 / 25.8	64.8	50.5	70.3	27.1	63.8	37.2	46.73%	77.3 $\rightarrow$ 64.1 $\rightarrow$ 52.3% / 78.9 $\rightarrow$ 60.2 $\rightarrow$ 42.2%
6.5B `4L:1A` KDA + TrainMaxTau ($\tau_{\max} = 8$K)	34.9	46.4 / 25.6	65.2	51.1	69.1	27.6	63.6	36.3	46.64%	**75.0 $\rightarrow$ 71.1 $\rightarrow$ 69.5% / 79.7 $\rightarrow$ 74.2 $\rightarrow$ 73.4%**
6.5**B** `4L:1A` **SLA-R2** (15.1**B, BF16/FP8**)	34.6	45.1 / 24.7	65.4	52.2	70.6	28.5	63.8	38.2	**47.01%** (46.92%)	**71.9 $\rightarrow$ 69.5 $\rightarrow$ 71.1% / 71.1 $\rightarrow$ 71.1 $\rightarrow$ 74.2%**
6.5**B** `4L:1A` **SLA (Fused) R2** (15.1**B, BF16/FP8**)	35.8	46.4 / 24.3	64.5	51.6	71.4	28.5	63.8	36.4	46.97% (**47.08%**)	**75.8 $\rightarrow$ 72.7 $\rightarrow$ 68.8% / 72.7 $\rightarrow$ 72.7 $\rightarrow$ 70.3%**
6.5B `0L:1A` GQA+RoPE (30.2B, 4.4$\times$ Chin.)	37.7	47.3 / 26.7	66.3	52.2	73.0	29.6	63.5	35.8	48.01%	72.7 $\rightarrow$ 53.1 $\rightarrow$ 54.7% / 75.8 $\rightarrow$ 55.5 $\rightarrow$ 50.0%
6.5B `4L:1A` MaxText KDA (30.2B, 4.4$\times$ Chin.)	38.8	50.4 / 27.1	66.4	52.2	78.3	30.6	63.5	39.6	**49.66%**	79.7 $\rightarrow$ 53.1 $\rightarrow$ 32.0% / 85.2 $\rightarrow$ 39.8 $\rightarrow$ 20.3%
6.5**B** `4L:1A` **SLA-R2** (30.2**B, BF16/FP8**)	38.4	50.7 / 26.5	67.1	51.7	75.1	28.3	63.9	43.6	49.47% (49.44%)	78.1 $\rightarrow$ 59.4 $\rightarrow$ 56.2% / 79.7 $\rightarrow$ 58.6 $\rightarrow$ 51.6%
6.5**B** `4L:1A` **Iso-SLA-R4** (30.2**B, BF16/FP8**)	37.6	49.4 / 26.1	67.4	49.0	74.8	29.7	63.8	44.4	49.13% (48.97%)	**72.7 $\rightarrow$ 71.9 $\rightarrow$ 70.3% / 82.8 $\rightarrow$ 74.2 $\rightarrow$ 74.2%**
12.5B `0L:1A` GQA+RoPE (15.1B, $n = 3$ mean)	36.1	47.4 / 25.3	65.6	51.5	73.1	29.2	63.3	39.6	47.90 $\pm$ 0.36%	72.9 $\rightarrow$ 54.9 $\rightarrow$ 57.3% / 70.8 $\rightarrow$ 57.0 $\rightarrow$ 51.3%
12.5B `7L:1A` MaxText KDA (15.1B, $n = 2$ mean)	37.2	47.2 / 26.7	66.4	50.4	72.5	29.9	62.6	37.8	47.83 $\pm$ 0.48%	71.9 $\rightarrow$ 62.5 $\rightarrow$ 56.6% / 76.2 $\rightarrow$ 62.9 $\rightarrow$ 57.8%
12.5B `7L:1A` Scalar GDN ($G = 1$, 15.1B, `s1`)	37.0	47.6 / 25.5	65.2	51.5	74.1	28.8	62.8	40.6	48.11%	**74.2 $\rightarrow$ 72.7 $\rightarrow$ 74.2% / 77.3 $\rightarrow$ 76.6 $\rightarrow$ 76.6%**
12.5B `7L:1A` Wide-Key KDA ($G = 256$, $n = 2$ mean)	37.2	47.3 / 26.4	65.8	52.2	73.3	27.8	63.0	41.8	48.30%	74.2 $\rightarrow$ 66.8 $\rightarrow$ 55.9% / 76.2 $\rightarrow$ 68.0 $\rightarrow$ 54.3%
12.5**B** `7L:1A` **A2** ($G = 4$ **Exact Inv., FP32,** `s0`)	38.1	47.9 / 25.7	65.0	51.1	73.8	28.7	59.8	42.3	48.05% (48.04%)	71.9 $\rightarrow$ 65.6 $\rightarrow$ 62.5% / 73.4 $\rightarrow$ 70.3 $\rightarrow$ 68.0%
12.5**B** `7L:1A` **A3** ($G = 4$ `1-Step Neumann`, **FP32,** $n = 2$)	37.0	46.8 / 25.8	65.9	50.0	72.6	28.6	63.6	38.9	47.66% (47.47%)	**76.6 $\rightarrow$ 71.9 $\rightarrow$ 71.9% / 74.6 $\rightarrow$ 72.3 $\rightarrow$ 71.9%**
12.5**B** `7L:1A` **SLA-R2** (15.1**B,** $n = 3$ **mean**)	37.6	47.9 / 26.1	66.2	52.3	72.2	29.3	63.1	42.1	**48.55 $\pm$ 0.48%** (**48.64%**)	**73.7 $\rightarrow$ 68.8 $\rightarrow$ 69.5% / 74.0 $\rightarrow$ 71.9 $\rightarrow$ 73.7%**
12.5**B** `7L:1A` **SLA (Fused) R2** (15.1**B, BF16/FP8,** `s0`)	37.6	47.7 / 26.3	66.3	50.8	71.6	28.0	60.0	45.3	48.18% (48.21%)	**73.4 $\rightarrow$ 67.2 $\rightarrow$ 64.1% / 76.6 $\rightarrow$ 75.8 $\rightarrow$ 75.0%**
12.5**B** `7L:1A` **Iso-SLA-R4** (15.1**B,** $n = 3$ **mean**)	36.3	48.8 / 25.9	66.3	50.6	72.3	28.7	63.0	43.0	48.33 $\pm$ 0.45% (48.39%)	**74.5 $\rightarrow$ 72.7 $\rightarrow$ 73.2% / 74.5 $\rightarrow$ 75.5 $\rightarrow$ 72.9%**
12.5B `15L:1A` MaxText KDA (15.1B, FP32)	36.0	48.6 / 26.6	66.0	50.2	71.9	29.4	64.2	40.7	48.18%	76.6 $\rightarrow$ 68.0 $\rightarrow$ 66.4% / 75.8 $\rightarrow$ 70.3 $\rightarrow$ 68.0%
12.5**B** `15L:1A` **SLA-R2 Uniform** ($n = 3$ **mean**)	37.0	50.1 / 27.0	65.6	50.5	74.5	29.2	63.5	41.6	**48.79 $\pm$ 0.22%** (**48.75%**)	**72.6 $\rightarrow$ 71.9 $\rightarrow$ 69.0% / 72.2 $\rightarrow$ 65.9 $\rightarrow$ 65.1%** (68.2% N=512)
12.5**B** `15L:1A` **SPD-SLA-W Polar-2+ Uniform** (`s0`)	35.9	49.9 / 26.2	65.7	50.2	74.1	29.8	63.8	41.5	48.57% (**48.79%**)	71.9 $\rightarrow$ 71.1 $\rightarrow$ 65.6% / 71.9 $\rightarrow$ 68.8 $\rightarrow$ 65.6% (69.5% N=512)
12.5B `15L:1A` SLA-R2+MS (15.1B, BF16/FP8)	36.6	50.8 / 27.3	65.8	51.1	74.1	27.5	63.5	37.2	48.20% (48.18%)	57.0 $\rightarrow$ 57.0 $\rightarrow$ 57.0% / 67.2 $\rightarrow$ 57.8 $\rightarrow$ 53.9%
12.5B `Pure-Linear` KDA (15.1B, FP32)	37.0	49.1 / 25.5	65.6	51.2	70.1	27.7	63.6	40.1	47.77%	57.8 $\rightarrow$ 57.8 $\rightarrow$ 55.5% / 62.5 $\rightarrow$ 56.2 $\rightarrow$ 56.2%
12.5B `Pure-Linear` SLA-R2+MS (15.1B, BF16/FP8)	36.4	50.3 / 26.7	66.6	49.9	72.6	27.9	63.3	40.2	48.21% (48.30%)	60.2 $\rightarrow$ 57.8 $\rightarrow$ 59.4% / 62.5 $\rightarrow$ 56.2 $\rightarrow$ 53.9%
12.5B `0L:1A` GQA+RoPE (32K-ext, 10.1B)	37.5	49.0 / 24.2	65.8	51.7	73.7	28.9	63.8	40.6	48.35%	75.0 $\rightarrow$ 64.8 $\rightarrow$ 56.2% / 73.4 $\rightarrow$ 65.6 $\rightarrow$ 57.8%
12.5B `7L:1A` MaxText KDA (32K-ext, 10.1B)	37.9	47.0 / 26.6	66.8	51.8	72.0	30.0	59.2	37.2	47.62%	73.4 $\rightarrow$ 73.4 $\rightarrow$ 71.9% / 78.1 $\rightarrow$ 78.1 $\rightarrow$ 72.7%
12.5**B** `7L:1A` **SLA-R2** (32**K-ext, BF16/FP8**)	38.9	48.4 / 26.5	67.1	51.5	71.9	29.4	62.6	42.7	48.78% (48.86%)	74.2 $\rightarrow$ 72.7 $\rightarrow$ 71.9% / 71.9 $\rightarrow$ 77.3 $\rightarrow$ **75.8%**
12.5**B** `7L:1A` **Iso-SLA-R4** (32**K-ext, BF16/FP8**)	37.1	49.8 / 26.0	65.7	50.7	75.5	30.0	61.9	44.7	**49.04%** (**49.18%**)	**78.1 $\rightarrow$ 75.0 $\rightarrow$ 75.8%** / 78.1 $\rightarrow$ 72.7 $\rightarrow$ 74.2%
12.5B `0L:1A` GQA+RoPE (60.4B, BF16)	41.4	49.2 / 27.7	68.3	51.2	76.9	29.6	62.7	43.5	50.06%	70.3 $\rightarrow$ 59.4 $\rightarrow$ 58.6% / 76.6 $\rightarrow$ 58.6 $\rightarrow$ 53.1%
12.5B `7L:1A` MaxText KDA (60.4B, $n = 3$ mean)	41.4	51.4 / 27.4	68.2	51.4	75.7	30.4	62.9	44.6	50.36 $\pm$ 0.53%	79.2 $\rightarrow$ 69.3 $\rightarrow$ 54.9% / 75.0 $\rightarrow$ 65.6 $\rightarrow$ 50.3%
12.5B `7L:1A` KDA + TrainMaxTau (60.4B, $n = 2$)	40.9	51.2 / 27.6	67.8	51.9	75.8	30.2	62.8	44.5	50.30%	**78.1 $\rightarrow$ 75.0 $\rightarrow$ 74.5% / 77.3 $\rightarrow$ 73.8 $\rightarrow$ 72.8%** (N=512)
12.5**B** `7L:1A` **SLA-R2** (60.4**B,** $n = 3$ **mean**)	41.4	52.5 / 27.2	68.6	52.0	76.7	29.9	63.2	46.9	**50.94 $\pm$ 0.47%** (**50.93%**)	**78.1 $\rightarrow$ 75.3 $\rightarrow$ 75.5% / 78.6 $\rightarrow$ 78.7 $\rightarrow$ 77.9%** (77.4% FP8)
12.5**B** `7L:1A` **SLA-R4** (60.4**B, BF16/FP8**)	41.3	49.9 / 25.7	67.6	51.1	76.1	30.3	63.2	47.3	50.28% (50.13%)	**72.7 $\rightarrow$ 75.8 $\rightarrow$ 72.7% / 72.7 $\rightarrow$ 78.9 $\rightarrow$ 75.0%** (76.6% FP8)
12.5**B** `7L:1A` **Iso-SLA-R4** (60.4**B, BF16/FP8**)	41.2	48.9 / 28.1	67.1	47.7	76.7	30.0	62.3	49.3	50.13% (50.31%)	**78.1 $\rightarrow$ 75.0 $\rightarrow$ 76.6% / 78.1 $\rightarrow$ 80.5 $\rightarrow$ 77.3%** (77.3% FP8)
12.5**B** `7L:1A` **Iso-SLA (Fused) R4** (60.4**B,** $n = 3$ **mean**)	41.5	52.6 / 27.8	67.9	51.7	77.5	29.5	63.1	45.8	**50.82 $\pm$ 0.22%** (**50.84%**)	**77.1 $\rightarrow$ 73.2 $\rightarrow$ 74.7% / 75.0 $\rightarrow$ 73.2 $\rightarrow$ 74.5%** (73.2% FP8)
12.5B `15L:1A` MaxText KDA (60.4B, FP32)	41.6	49.5 / 27.2	68.0	51.7	77.2	30.1	62.7	45.6	50.41%	77.3 $\rightarrow$ 69.5 $\rightarrow$ 62.5% / 76.6 $\rightarrow$ 71.1 $\rightarrow$ 60.2%
12.5**B** `15L:1A` **SLA-R2 Uniform** (60.4**B, BF16/FP8**)	41.1	53.3 / 28.4	68.0	52.6	77.6	30.3	63.6	46.3	**51.24%** (**51.08%**)	**78.9 $\rightarrow$ 75.8 $\rightarrow$ 72.7% / 76.6 $\rightarrow$ 72.7 $\rightarrow$ 70.3%** (70.3% FP8)
12.5B `15L:1A` SLA-R2+MS (60.4B, BF16/FP8)	42.8	56.4 / 28.4	68.2	49.8	78.5	28.9	63.8	40.6	50.82% (51.02%)	67.2 $\rightarrow$ 64.1 $\rightarrow$ 64.8% / 71.1 $\rightarrow$ 57.8 $\rightarrow$ 57.0% (57.0% FP8)
12.5B `15L:1A` SPD-SLA-W+MS (60.4B, BF16/FP8)	42.0	53.5 / 28.9	68.5	50.3	77.1	30.1	63.6	40.8	50.54% (50.59%)	60.9 $\rightarrow$ 59.4 $\rightarrow$ 57.8% / 66.4 $\rightarrow$ 58.6 $\rightarrow$ 56.2% (56.2% FP8)
63.1B `0L:1A` GQA+RoPE (30.2B, RoPE)	40.6	50.2 / 27.0	68.4	51.6	74.0	28.9	63.8	42.3	49.64%	71.9 $\rightarrow$ 53.9 $\rightarrow$ 45.3% / 71.9 $\rightarrow$ 56.2 $\rightarrow$ 39.8%
63.1B `4L:1A` GDN-2 (30.2B, RoPE, 2,728 ms)	37.9	50.8 / 26.9	67.7	51.3	72.6	29.7	63.1	40.7	48.96%	69.5 $\rightarrow$ 62.5 $\rightarrow$ 55.5% / 72.7 $\rightarrow$ 66.4 $\rightarrow$ 64.1%
63.1**B** `4L:1A` **SLA (Fused) R4** (30.2**B, RoPE**)	37.7	47.3 / 26.1	67.5	51.7	73.4	29.2	59.6	44.2	48.53% (48.51%)	68.0 $\rightarrow$ 59.4 $\rightarrow$ **59.4%** / 71.1 $\rightarrow$ 59.4 $\rightarrow$ 54.7%
63.1**B** `15L:1A` **SPD-SLA-W+MS** (30.2**B, RoPE**)	38.1	48.4 / 25.2	66.4	50.8	71.5	30.3	63.0	42.0	48.40% (48.65%)	**71.1 $\rightarrow$ 64.1 $\rightarrow$ 57.8% / 72.7 $\rightarrow$ 65.6 $\rightarrow$ 64.8%**
63.1B `4L:1A` MaxText KDA (120.8B, NoPE)	**54.4**	61.5 / **35.4**	**73.2**	**54.6**	86.9	34.1	**64.1**	56.2	**57.83%**	85.9 $\rightarrow$ 39.1 $\rightarrow$ 33.6% / **89.1** $\rightarrow$ 28.9 $\rightarrow$ 27.3%
63.1**B** `4L:1A` **SLA-R2** (120.8**B, NoPE**)	53.4	**62.4** / 34.6	72.2	53.7	**87.5**	**34.4**	**64.1**	55.4	57.52%	**88.3 $\rightarrow$ 60.2 $\rightarrow$ 57.0%** / 84.4 $\rightarrow$ **61.7 $\rightarrow$ 51.6%**
63.1**B** `4L:1A` **SLA-R4** (120.8**B, NoPE**)	53.8	57.1 / 34.6	73.1	51.9	85.5	34.3	63.0	**57.2**	56.72% (56.74%)	85.2 $\rightarrow$ 57.0 $\rightarrow$ 53.9% / 87.5 $\rightarrow$ 51.6 $\rightarrow$ 43.0% (43.8% FP8)

for `Unsorted_4Tile_SLA` (-50.86 PPL, $t = -34.94, p = 0.00082$) and $\mathbf{228.39}$ for scalar `GDN` (G=1, -55.24 PPL, $t = -35.18, p = 0.00081$). After 200 lightweight healing steps, `Sorted_DP_4Tile_SLA` reaches $\mathbf{171.79}$ PPL (218.45 OOD) while running $\mathbf{1.734}\times$ faster than KDA ($\mathbf{57.01}$ vs. 98.86 ms/step), outperforming healed `Unsorted_4Tile_SLA` (174.67, $p = 0.00737$) and `Scalar_GDN` (178.91).

2. **Woodbury Tile-RLS vs. 4-Tile SLA Boundary Audit:** On synthetic high-coherence bursts ($\rho = 0.90$ and 0.98), exact intra-chunk `Woodbury_RLS` maintains $\mathbf{100.0}\%$ recall (0.0550 and 0.1980 MSE) where `KDA_PerChannel`, `SLA_4Tile`, and `Neumann_Damped` drop to 50.0% (0.7604 and 0.9436 MSE). On natural 8L FineWeb-Edu pre-training ($n = 3$ seeds, 0.0 future Jacobian leak), however, `SLA_4Tile` **wins outright** ($\mathbf{165.14 \pm 1.45}$ FW PPL, $\mathbf{195.67}$ OOD, $\mathbf{57.18}$ ms/step) over `KDA_PerChannel` (166.62 ± 1.51, 206.64 OOD, 98.88 ms), `Damped_1Step_Neumann` (166.85 ± 0.94, 36.33 ms, beating undamped 170.41 at $p = 0.0245$), and `Woodbury_RLS_M128` (181.67 ± 0.40, 148.93 ms, $p = 0.00145$; sparse budgets `M64` 210.39, `M32` 227.53).

3. **Real-Checkpoint Co-Design Across Five Open Models:** Across five public HuggingFace checkpoints, conversion cost follows an exact algebraic hierarchy: (a) *Scalar GDN (G=1 $\rightarrow$ G=4):* `Qwen/Qwen3-Next-80B-A3B-Instruct` (`L0,5,13,21`) and full 24L `Qwen/Qwen3.5-0.8B-Base` lift to 4-tile `SLA` (C=128) as an *exact zero-shot algebraic identity* (NMSE $= 1.38$–2.46×10^{-13} and 3.78×10^{-14}; 0.0000 PPL delta at 18.7741 FW / 19.6594 PG-19; 19.0806 FW healed across n=3 seeds after 4-tile symmetry breaking, with $\mathrm{SU}(2)^8$ Tile-0 holonomy reaching $\mathbf{99.80}\%$ $4\times$-OOD A_5 vs. 2.02% diagonal and 4.02% Abelian $\mathrm{SO}(2)^8$), whereas $2\times$-head packing ($16\times128 \rightarrow 8\times256$) incurs 1.885 0-shot NMSE from cross-head interference before healing $12.0\times$ to 0.157; (b) *Per-Channel KDA (G=128 $\rightarrow$ G=4):* co-permuting $W_q, W_k,$ `Conv1D`, $f_1,$ `dt_bias` by mean channel log-decay is exact function-preserving (6.7×10^{-7} max error on 27L `moonshotai/Kimi-Linear-48B-A3B-Base`; $\mathbf{0.000000}$ full-24L PPL delta at 12.9224 on `openeurollm/kda-sigmoid-hybrid-1.3B-100B`), and Fisher 1D DP partitioning into G=4 MXU tiles cuts within-tile decay SSE $4.55 \times$ –$9.54\times$, cuts 0-shot layer NMSE -85.1% on `KDA-1.3B` (0.2841 $\rightarrow$ 0.0423) and -66.6% on `Kimi-Linear-48B` (0.1458 $\rightarrow$ 0.0487), and cuts full-24L `KDA-1.3B` 0-shot FineWeb/PG-19 PPL from $20.1180/37.3492$ (`Scalar_GDN`) and $19.4997/35.9863$ (`Unsorted_4Tile`) to $\mathbf{14.0369}/\mathbf{25.3191}$ ($\mathbf{-6.08}/\mathbf{-12.03}$ PPL), healing in 40 layer steps (n=3) to $\mathbf{13.5537 \pm 0.0016}$ FW (21.3809 PG-19) vs. 16.0095 `Unsorted` (p=2.5×10^{-6}) and 17.3990 `Scalar_GDN` (p=1.1×10^{-5}); (c) *Pure Softmax GQA (100% KV):* converting $21/28$ layers of `Qwen/Qwen3-0.6B-Base` (18.1712 FW / 19.5784 PG-19) to a 3:1 hybrid (-75% KV cache) collapses zero-shot without distillation (2.88×10^6 PPL; 0.9149 layer NMSE), but 60 layer distillation steps recover 0.0108–0.0229 layer NMSE (0.989–0.995 cosine) and bring full-28L perplexity to $\mathbf{126.60 \pm 3.11}$ FW / $\mathbf{139.22}$ PG-19 for `4-Tile SLA`, outperforming `Scalar_GDN` ($161.36/179.97$) by $\mathbf{-34.77}$ FW PPL (p=0.0165).

Table 27: Extended TPU v7x verification of `Sorted-DP` KDA-to-SLA conversion, `Woodbury_RLS` boundary, and 5-model pretrained checkpoint co-design (`Qwen3-Next-80B`, `Qwen3.5-0.8B`, `Kimi-Linear-48B`, `KDA-Hybrid-1.3B`, `Qwen3-0.6B`).

Experiment / Model	Variant / Conversion	0-Shot FW / SSE	Healed / Final FW PPL	OOD / PG-19 PPL	Step (ms) / Speedup	Key Diagnostic / p-Value
TPU v7x 8L KDA to 4-Tile SLA (n=3)	`KDA_PerChannel` Teacher (G=256)	154.71 (SSE 48.75)	154.71 (4,000 st.)	186.26	98.86 ms (1.000×)	teacher reference
	`Scalar_GDN_G1` (G=1)	228.39 (SSE 49.27)	178.91 (+200 st.)	251.78	54.05 ms (1.829×)	-55.24 0-shot vs. Sorted
	`Unsorted_4Tile_SLA` (G=4)	224.02 (SSE 48.75)	174.67 (+200 st.)	238.18	57.01 ms (1.734×)	-50.86 0-shot vs. Sorted
	`Sorted_DP_4Tile_SLA` (G=4)	**173.16** (**4.22**× ↓)	**171.79** (+200 st.)	**218.45**	**57.01** ms (**1.734**×)	$p = 0.00082$ (0-shot), $p = 0.0074$
TPU v7x 8L RLS vs. SLA (n=3)	`SLA_4Tile` (G=4, Exact Inv.)	ρ=0.98: 50% rec.	$\mathbf{165.14 \pm 1.45}$	**195.67**	**57.18** ms (**1.73**×)	best natural-text LM & OOD
	`Damped_1Step_Neumann`	ρ=0.98: 50% rec.	166.85 ± 0.94	224.20	**36.33** ms (**2.72**×)	beats undamped 170.41 (p = .0245)
	`KDA_PerChannel` (G=256)	ρ=0.98: 50% rec.	166.62 ± 1.51	206.64	98.88 ms (1.00×)	8K/32K ker 17.6/68.7 ms
	`Woodbury_RLS_M128` (exact RLS)	ρ=0.98: **100%** rec.	181.67 ± 0.40	219.95	148.93 ms (0.66×)	0.0 leak; over-whitens natural text
Pretrained GDN (G=1 → G=4)	`Qwen3-Next-80B` (`L0,5,13,21`)	NMSE 1.4–2.5e-13	Exact 0-shot parity	–	9.83 ms (32K `SLA`)	exact algebraic lift
	`Qwen3.5-0.8B` (24L, n=3)	**18.7741** (Δ=0.0000)	**19.0806** (+35 st.)	**20.5719** PG-19	A_5: **99.80**% (4×)	$\mathrm{SU}(2)^8$ Tile-0 vs. 2.02% diag.
Pretrained KDA (G=128 → G=4)	`Kimi-Linear-48B` (7 layers)	NMSE **0.0487** (6.8× ↓)	0.0383 (+60 st.)	cos **0.9754**	perm err 6.7e-7	-66.6% 0-shot NMSE vs. G=1
	`KDA-Hybrid-1.3B` Teacher (G=128)	12.9224 (Δ_{perm}=0.0)	12.9224 (24L base)	17.8733 PG-19	1.00× ref.	exact sorted perm Δ=0.0000
	`KDA-1.3B Scalar_GDN` (G=1)	20.1180 (0.98×)	17.3990 (+40 st.)	34.7446 PG-19	-3.85 behind DP	collapses multi-scale decay
	`KDA-1.3B Unsorted_4Tile` (G=4)	19.4997 (1.00×)	16.0095 (+40 st.)	29.1106 PG-19	-2.46 behind DP	mixes fast/slow channels
	`KDA-1.3B Sorted_DP_4Tile`	**14.0369** (**4.55**× ↓)	**13.5537** (+40 st.)	**21.3809** PG-19	**1.73**× kernel	$\mathbf{-6.08}$ 0-shot; p=2.5e-6
`Qwen3-0.6B` GQA → 3:1 (-75% KV)	`28L Full GQA` / 0-Shot Swap	$18.17 / 2.88 \times 10^6$	18.1712 (teacher)	19.5784 / 2.73e6	100% KV ref.	0-shot GQA swap requires distill
	`3:1 Hybrid Scalar_GDN` (G=1)	NMSE 0.914 → 0.0108	161.36 ± 5.81	179.97 PG-19	-75% KV cache	60 layer distill steps (n=3)
	`3:1 Hybrid 4-Tile SLA` (G=4)	NMSE 0.915 → 0.0229	$\mathbf{126.60 \pm 3.11}$	**139.22** PG-19	-75% KV cache	$\mathbf{-34.77}$ FW vs. G=1 (p=.0165)

H Shared-β_t Isolation, Open-Source Solver Audits, and Exact vs. Approximate Inversion

This appendix provides the TPU v7x control experiments that back our credit to open-source MXU chunk solvers, isolate SLA's shared-β_t value-packing speed and quality trade-off at fixed layer shape, test $C = 256$ Mosaic compilation, and document our exploratory study of exact versus approximate chunk inversion. All 3-seed error bars in this section report sample standard deviations (ddof $= 1$).

H.1 Tokamax BF16 MXU vs. FP32 Sequential Substitution Audit

In open-source Tokamax [23], the Mosaic TPU KDA kernel dispatches to an MXU solver (inverting 8-token diagonal blocks via a Neumann series and combining blocks by Horner doubling with FP32 accumulation) when QKV inputs are BF16, and falls back to sequential forward substitution when QKV inputs are FP32. Table 28 isolates both paths in the same single-chip TPU v7x harness ($B = 1$, L2-normalized QK, 20 timed repeats after 3 warmups, excluding causal `Conv4`). Switching QKV inputs from FP32 to BF16 speeds up Tokamax by $1.68\times$–$2.25\times$ in core forward, $2.41\times$–$2.58\times$ in core F+B, and $1.63\times$–$1.71\times$ in full-layer F+B while agreeing with the FP32 path to 4.14×10^{-3}–4.19×10^{-3} relative L_2 error ($\cos = 0.999992$, all outputs finite). Our Table 5 benchmark already cast QKV inputs to BF16 before calling Tokamax Mosaic KDA (with the table's FP32 label referring to its FP32 recurrent state), so all reported SLA-vs-Tokamax speedups compare against Tokamax's fast BF16 MXU path. Even in this isolated single-chip harness where Tokamax omits causal `Conv4` and `SLA-R4` includes it, fused 1-kernel `SLA-R4` (Table 30, $2.912/9.952$ ms at $D = 2048$ and $3.896/14.247$ ms at $D = 1792$) remains $1.23\times$–$1.72\times$ faster than BF16-input Tokamax ($3.580/12.962$ and $6.413/24.462$ ms).

Table 28: **Tokamax Mosaic KDA input-precision audit on single-chip TPU v7x** ($B = 1$, $C = 64$, 20 repeats, L2-normalized QK, all core and full outputs finite). BF16 QKV inputs activate Tokamax's MXU Neumann–Horner solver; FP32 QKV inputs take its sequential forward-substitution fallback.

Geometry ($H_k \times d_k \times d_v$)	Seq T	QKV Path	Core Fwd (ms)	Core F+B (ms)	Full Fwd (ms)	Full F+B (ms)	FP32 / BF16 (Core / Full F+B)	Peak HBM (MiB)	BF16 vs. FP32 Rel. L_2
$D = 2048$ ($8 \times 128 \times 256$)	8K	**BF16 (MXU Neumann–Horner)**	**0.797**	**1.573**	**1.457**	**3.580**	2.44× / 1.63×	960.1	0.004189 (cos 0.999992)
		FP32 (Seq. Substitution)	1.653	3.832	2.365	5.834		927.9	ref
$D = 2048$ ($8 \times 128 \times 256$)	32K	**BF16 (MXU Neumann–Horner)**	**2.651**	**5.640**	**4.995**	**12.962**	2.58× / 1.68×	3,122.3	0.004194 (cos 0.999992)
		FP32 (Seq. Substitution)	5.965	14.560	8.544	21.818		2,763.9	ref
$D = 1792$ ($8 \times 256 \times 512$)	8K	**BF16 (MXU Neumann–Horner)**	**1.551**	**2.995**	**2.678**	**6.413**	2.44× / 1.71×	1,582.1	0.004138 (cos 0.999992)
		FP32 (Seq. Substitution)	2.611	7.298	3.867	10.947		1,580.9	ref
$D = 1792$ ($8 \times 256 \times 512$)	32K	**BF16 (MXU Neumann–Horner)**	**5.652**	**12.095**	**9.835**	**24.462**	2.41× / 1.70×	5,571.6	0.004142 (cos 0.999992)
		FP32 (Seq. Substitution)	10.055	29.102	14.980	41.559		5,051.6	ref

H.2 TPU v7x MXU Solver Accuracy vs. FP64 Across Chunk Sizes

To test how existing MXU solver formulations behave under real chunk coupling norms, we execute six inversion algorithms directly on TPU v7x against an FP64 reference across three chunk banks: real FineWeb-Edu chunks extracted from Qwen3.5-0.8B and Qwen3.5-2B (288 chunks each at $C \in \{32, 64, 128\}$, where 77.4%–86.5% of chunks have $\|M_c\|_2 > 1$ and $\max \|M_c\|_2$ reaches 61.6–62.3), plus high-coherence $G = 4$ SLA chunks with Householder gates $\beta_t \in (0, 2)$ (192 chunks at $C \in \{32, 64, 128, 256\}$, 100% with $\|M_c\|_2 > 1$). Table 29 shows that FP32 recursive block-doubling stays at machine rounding (1.4×10^{-7}–3.0×10^{-7} max relative error at every chunk size through $C = 256$, matching FP32 forward substitution). Even with BF16 operands, recursive block-doubling remains bounded below 1.5×10^{-2} with 0% chunks exceeding 0.1. By contrast, Tokamax's 8-token blocked Neumann–Horner formulation is accurate at $C \leq 64$ in FP32 (1.2×10^{-6}–2.8×10^{-6}), but with BF16 MXU operands at $C = 128$ its worst-case error grows to 10.6–12.7 (16.7%–17.4% of Qwen3.5 chunks above 0.1), and unblocked Neumann products diverge on 21.5%–23.3% of Qwen3.5 chunks and 100% of high-coherence $C \geq 128$ chunks (consistent with the warning in MaxText's GDN solver docstring).

H.3 Same-Shape Isolation of Shared β_t vs. Per-Value-Head Gates

Changing the packing factor from $R = 1$ to $R = 4$ alters QK projections, convolutions, and Gram matrices together with the number of chunk solves (on the identical Exact FP32 $C = 128$ kernel at $D = 2048, T = 8$K, packing $R = 1 \to 2 \to 4$ cuts core forward from $2.927 \to 1.631 \to 1.002$ ms

Table 29: **TPU v7x MXU solver accuracy vs. FP64 on real Qwen3.5 chunks and high-coherence SLA chunks** ($\beta_t \in (0, 2)$). Each entry reports worst-case relative Frobenius error $\max_c \left\| \hat{A}_c - A_c^{\text{FP64}} \right\|_F / \left\| A_c^{\text{FP64}} \right\|_F$ and the share of chunks exceeding 0.1 error ($> 10^{-1}$).

Chunk Bank (n chunks)	C	$\|M_c\|_2 > 1$ (max)	**Block-Doub. FP32**	Fwd. Sub. FP32	Block-Doub. BF16	Blocked N+H FP32	Blocked N+H BF16	Unblocked Prod. FP32
Qwen3.5-0.8B (288)	32	78.5% (17.0)	$\mathbf{1.35 \times 10^{-7}}$ (0%)	2.60×10^{-7} (0%)	6.73×10^{-3} (0%)	1.41×10^{-6} (0%)	8.91×10^{-2} (0%)	4.41 (3.1%)
	64	81.3% (32.8)	$\mathbf{1.78 \times 10^{-7}}$ (0%)	3.83×10^{-7} (0%)	6.73×10^{-3} (0%)	2.29×10^{-6} (0%)	8.91×10^{-2} (0%)	8.05×10^{8} (10.1%)
	128	85.1% (62.3)	$\mathbf{2.53 \times 10^{-7}}$ (0%)	4.61×10^{-7} (0%)	7.19×10^{-3} (0%)	6.72×10^{-4} (0%)	10.64 (17.4%)	3.00×10^{27} (21.5%)
Qwen3.5-2B (288)	32	77.4% (16.7)	$\mathbf{1.44 \times 10^{-7}}$ (0%)	2.33×10^{-7} (0%)	6.07×10^{-3} (0%)	1.19×10^{-6} (0%)	9.00×10^{-2} (0%)	1.18 (1.4%)
	64	80.9% (32.0)	$\mathbf{1.74 \times 10^{-7}}$ (0%)	3.34×10^{-7} (0%)	6.41×10^{-3} (0%)	2.76×10^{-6} (0%)	9.00×10^{-2} (0%)	5.20×10^{8} (9.0%)
	128	86.5% (61.6)	$\mathbf{1.80 \times 10^{-7}}$ (0%)	4.22×10^{-7} (0%)	6.97×10^{-3} (0%)	4.60×10^{-4} (0%)	12.69 (16.7%)	1.39×10^{26} (23.3%)
SLA $G = 4$ High-Coh. ($192, \beta_t \in (0, 2)$)	32	100% (7.8)	$\mathbf{1.83 \times 10^{-7}}$ (0%)	2.85×10^{-7} (0%)	1.14×10^{-2} (0%)	2.02×10^{-6} (0%)	0.312 (1.6%)	0.224 (2.1%)
	64	100% (9.0)	$\mathbf{2.71 \times 10^{-7}}$ (0%)	3.44×10^{-7} (0%)	1.14×10^{-2} (0%)	2.02×10^{-6} (0%)	0.310 (3.6%)	5.60×10^{7} (81.8%)
	128	100% (9.0)	$\mathbf{3.02 \times 10^{-7}}$ (0%)	3.44×10^{-7} (0%)	1.27×10^{-2} (0%)	2.92×10^{-6} (0%)	0.310 (6.8%)	1.96×10^{22} (100%)
	256	100% (9.3)	$\mathbf{3.02 \times 10^{-7}}$ (0%)	3.44×10^{-7} (0%)	1.49×10^{-2} (0%)	2.90×10^{-6} (0%)	0.310 (9.9%)	∞ (100%)

[2.92×] and core F+B from 5.742 → 3.191 → 2.105 ms [2.73×]). To isolate the effect of sharing β_t alone without changing QK head count or tensor shapes, Table 30 fixes $H_k = 2, d_k = 256, H_v = 8, G = 4, C = 128$ (`Exact Inv.`) and compares three gating regimes inside the identical fused kernel:

1. `SLA-R4 shared-`β_t: shared g and shared β_t per key head (2 solves per chunk, packed $4d_v$ value slice).
2. `GVA shared-`g`, per-`H_v β_t: identical QK projections, convolutions, and shared g, but per-value-head β_t (8 solves per chunk, matching Qwen3-Next's gate factorization).
3. `GVA per-`H_v g `and` β_t: identical QK projections and convolutions, with independent (g, β_t) per value head (8 solves per chunk).

Under equal gates, all three paths match to 0.0 relative L_2 difference across all 12 rows. Simply sharing β_t at fixed shape accelerates core forward by 2.45×–3.28×, core F+B by 2.25×–3.00×, and full-layer F+B by 1.87×–2.48×. In a 3-seed paired from-scratch pre-training control (223.2M–223.5M parameters, 12-layer `3L:1A` hybrid, 800 steps = 52.4M FineWeb-Edu tokens, identical shape $H_k = 2, d_k = 256, H_v = 8, G = 4$), `SLA-R4 shared-`$\beta_t$ runs 1.68×–1.69× faster (72.92 ± 0.75 vs. 122.43 ± 1.04 and 123.38 ± 1.37 ms/step), with validation perplexity within seed noise (74.24 ± 3.56 vs. 72.66 ± 1.45 and 72.89 ± 1.89; `SLA-R4` beats `GVA per-`H_v g, β_t on seeds 1 and 2 by -1.14 and -1.53 PPL while trailing on seed 0 by $+6.71$). Exact $C = 256$ (`order=-3`) trains at 74.11 ± 0.17 ms/step (1.65×–1.66× faster) with 75.00 ± 3.11 Val PPL.

Table 30: **Unconfounded same-shape shared-β_t isolation on TPU v7x** (fixed $H_k = 2, d_k = 256, H_v = 8, G = 4, C = 128$, `Exact Inv.`). Top: single-core kernel and full-layer F+B timings (20 repeats; 0.0 rel. L_2 discrepancy when gates match). Bottom: 3-seed paired from-scratch pre-training (223M params, 52.4M FineWeb-Edu tokens, sample std ddof = 1).

(a) Same-shape kernel and full-layer speed ($B = 1$, **fixed** $H_k = 2, d_k = 256, H_v = 8, G = 4, C = 128$)

Geometry / Seq T	Gating Arm	Solves / Chunk	Core Fwd (ms)	Core F+B (ms)	Full Fwd (ms)	Full F+B (ms)	Speedup (Core Fwd / Core F+B / Full F+B)
$D = 2048$, 8K	`SLA-R4` **shared-**β_t	**2**	**1.034**	**2.159**	**1.230**	**2.912**	2.84× / 2.75× / 2.25×–2.30×
	`GVA` shared-g, per-H_v β_t	8	2.940	5.944	3.144	6.565	1.00× ref
	`GVA` per-H_v g and β_t	8	2.938	5.920	3.175	6.700	1.00× ref
$D = 2048$, 32K	`SLA-R4` **shared-**β_t	**2**	**3.337**	**7.199**	**4.008**	**9.952**	3.28× / 2.99×–3.00× / 2.47×–2.48×
	`GVA` shared-g, per-H_v β_t	8	10.953	21.622	11.755	24.542	1.00× ref
	`GVA` per-H_v g and β_t	8	10.930	21.547	11.947	24.713	1.00× ref
$D = 1792$, 8K	`SLA-R4` **shared-**β_t	**2**	**1.258**	**2.812**	**1.545**	**3.896**	2.45×–2.48× / 2.27× / 1.87×
	`GVA` shared-g, per-H_v β_t	8	3.080	6.377	3.387	7.286	1.00× ref
	`GVA` per-H_v g and β_t	8	3.120	6.377	3.414	7.297	1.00× ref
$D = 1792$, 32K	`SLA-R4` **shared-**β_t	**2**	**4.255**	**10.210**	**5.150**	**14.247**	2.70×–2.74× / 2.25×–2.28× / 1.91×–1.92×
	`GVA` shared-g, per-H_v β_t	8	11.674	23.295	12.862	27.328	1.00× ref
	`GVA` per-H_v g and β_t	8	11.483	23.007	12.486	27.247	1.00× ref

(b) Same-shape 3-seed from-scratch pre-training (223**M params,** 12**L** `3L:1A`, 52.4**M tokens, ddof** $= 1$)

Pre-training Arm	Params	Solves / Chunk	Chunk C	Val PPL ($n = 3$)	Per-Seed Val PPL (`s0, s1, s2`)	OOD PG-19 PPL	Step Time (ms) / Speedup
`SLA-R4` **shared-**β_t	223.18M	**2**	128	74.24 ± 3.56	[78.11, 71.12, 73.49]	85.15 ± 3.99	$\mathbf{72.92 \pm 0.75}$ (1.68×–1.69×)
`GVA` shared-g, per-H_v β_t	223.24M	8	128	72.66 ± 1.45	[72.33, 71.40, 74.24]	82.63 ± 1.42	123.38 ± 1.37 (1.00×)
`GVA` per-H_v g and β_t	223.46M	8	128	72.89 ± 1.89	[71.40, 72.26, 75.02]	82.30 ± 2.40	122.43 ± 1.04 (1.00×)
`SLA-R4` exact $C = 256$ (`order=-3`)	223.18M	**2**	256	75.00 ± 3.11	[78.57, 73.55, 72.87]	84.99 ± 1.98	74.11 ± 0.17 (1.65×–1.66×)

H.4 Exact vs. Matmul-Free Solvers at $C = 256$ (`kVectorMove`-Safe Lowering)

In our original kernel lowering (`order=0`, `order=-2`, and banded block solvers), compiling $C = 256$ on TPU v7x aborts inside Mosaic's graph optimizer (`llo_xlu_graph_optimizer.cc:2483`:

kVectorPack vs. kVectorMove) because sub-block packing emits short transposes over trailing zeros. Replacing those reshapes with a kVectorMove-safe lowering (order=-3 for Exact FP32 and order=-4 for Exact BF16) resolves the compiler crash cleanly across all three SLA shapes (Table 31), showing that the earlier $C = 256$ failure was a lowering issue rather than an inherent barrier to exact inversion. At $C = 256$, kVectorMove-safe exact FP32 matches the XLA FP32 reference to 4.1×10^{-3} relative L_2 error (BF16 I/O precision floor), whereas matmul-free n1row and eaadj are $1.41\times$–$2.03\times$ faster in isolated core forward ($1.20\times$–$1.45\times$ core F+B) at 8.2%–10.4% relative L_2 error.

Table 31: **Single-core TPU v7x** $C = 256$ **kernel check** ($B = 1, T = 8$K, 20 repeats, all compiled rows have finite outputs and gradients). Unmodified exact block-doubling (order=0,-2) hits a Mosaic kVectorPack assertion; kVectorMove-safe exact block-doubling (order=-3,-4) compiles cleanly across all shapes.

Solver at $C = 256$	order	SLA-R4 ($D = 2048$)		SLA-R2 ($D = 2048$)		SLA-R4 ($D = 1792$)	
		Fwd / F+B (ms)	Rel. L_2	Fwd / F+B (ms)	Rel. L_2	Fwd / F+B (ms)	Rel. L_2
Exact FP32 / BF16 (unmodified)	$0/-2$	Mosaic SIGABRT	–	Mosaic SIGABRT	–	Mosaic SIGABRT	–
Exact FP32 (kVectorMove-safe)	-3	1.070/2.054	**0.00411**	1.721/3.108	**0.00412**	1.314/2.708	**0.00410**
Exact BF16 (kVectorMove-safe)	-4	0.785/1.677	**0.00414**	1.091/2.401	**0.00415**	0.988/2.298	**0.00412**
Row-Scaled 1-Step (n1row)	1	**0.649/1.550**	0.1025	**0.848/2.151**	0.1045	**0.917/2.206**	0.1029
Signed Erase-Aware (eaadj)	201	0.688/1.606	0.0816	0.887/2.308	0.0955	0.938/2.269	0.0817
No-Solve Bound ($A_c = I_C$, fwd only)	-1	0.654 / –	0.2482	0.873 / –	0.2553	0.890 / –	0.2495

H.5 Exploratory Approximate Solvers: Collinear Over-Subtraction, eaadj, Banded Exact, and Pre-Training

We investigated whether matmul-free or banded approximations of $(I_C + M_c)^{-1}$ can replace exact block-doubling in inference or pre-training.

Root cause of zero-shot failure: collinear-key over-subtraction. Replacing $(I_C + M_c)^{-1}$ zero-shot on exact-trained Qwen3.5-0.8B/2B/9B checkpoints with raw Neumann-1 ($I_C - M_c$), raw Neumann-2 ($I_C - M_c + M_c^2$), or unbounded block-Neumann (bn1b16) collapses FineWeb-Edu perplexity to 866–8.1×10^5 (Table 32a). Diagnostics on FineWeb-Edu ($T = 512$) isolate the cause: deep linear layers learn strongly correlated consecutive keys (max lag-1 cosine 0.802–0.862, top-1 singular energy share 0.751–0.811), pushing $\|M_c\|_2 > 1$ in 74.4%–83.4% of all chunks (up to 100% in deep layers, with $\max \|M_c\|_2 = 29.6$–76.2). On a collinear sequence, intervening writes $i \in (j, t)$ already partially erase token j before token t reads the state, so an undamped $-M_{t,j}$ term over-subtracts exponentially.

Signed erase-aware 1-step operator (eaadj) and prefix-sum VJP. To damp collinear echoes without matrix multiplications while supporting signed Householder step sizes $\beta_t \in (0, 2)$, define adjacent-key alignment $a_1 = 1$, $a_i = \mathrm{clip}(\langle k_i, k_{i-1}\rangle, 0, 1)$ for $i \geq 2$, intervening retention factors $r_i = 1 - \beta_i a_i$, and

$$\Phi_{t,j} = \mathbf{1}[t > j] \prod_{j<i<t} r_i = \mathbf{1}[t > j]\,(-1)^{\sum_{j<i<t} \mathbf{1}[r_i<0]} \exp\Big(\sum_{j<i<t} \log\max(|r_i|, \varepsilon)\Big),$$
$$A_{\text{EA}} = I_C - M_c \odot \Phi, \tag{18}$$

with $\varepsilon = 10^{-4}$ for real-valued TPU vector prefix sums.

Theorem H.1 (Scalar-Decay Collinear Exactness of Signed eaadj)**.** *Suppose a chunk has collinear unit keys $k_1 = \cdots = k_C = u$ ($\|u\|_2 = 1$), scalar decay per head $F_t = \gamma_t I_{d_k}$ ($\gamma_t \in (0, 1]$, as in GDN / Qwen3.5 or SLA with $G = 1$), and arbitrary step sizes $\beta_t \in (0, 2)$. Then $a_i = 1$ for all i and $A_{EA} = I_C - M_c \odot \Phi$ equals $(I_C + M_c)^{-1}$ identically (under both column-scaled and row-scaled conventions). The exactness condition does* not *hold when SLA uses $G = 4$ tiles with distinct decay rates, where the scalar Gram sum no longer factorizes multiplicatively.*

Proof. Under collinear unit keys and scalar decay ($G = 1$), $(M_c)_{t,j} = \mathbf{1}[t > j]\,\beta_j e^{\Lambda_t - \Lambda_j}$ satisfies $(M_c)_{t,i}(M_c)_{i,j} = \beta_i (M_c)_{t,j}$ for every $j < i < t$. Thus for $t > j$,

$$\big((I_C + M_c)A_{\mathrm{EA}}\big)_{t,j} = (M_c)_{t,j}\left[1 - \Phi_{t,j} - \sum_{i=j+1}^{t-1} \beta_i \prod_{j<s<i} (1 - \beta_s)\right] = 0,$$

because the sum telescopes to $1 - \Phi_{t,j}$. Diagonal conjugation preserves Hadamard multiplication by Φ, so the row-scaled convention is likewise exact. On TPU v7x across $C \in \{32, 64, 128, 256\}$ with random $\beta_t \in (0.05, 1.95)$, measured max absolute error against exact FP32 is 5.09×10^{-6}, 7.07×10^{-6}, 1.63×10^{-5}, and 4.18×10^{-5} for $G = 1$ scalar decay, versus $0.0108, 0.00823, 0.0116$, and 0.0154 when SLA's $G = 4$ tiles carry distinct decay rates. □

Differentiating A_{EA} also requires zero matrix multiplications. Given upstream gradient $dA_c \in \mathbb{R}^{C\times C}$, let $W_{t,j} = -\mathbf{1}[t > j]\,(dA_c)_{t,j}\,(M_c)_{t,j}\,\Phi_{t,j}$ with column sums $c_j = \sum_{t>j} W_{t,j}$ and row sums $\rho_t = \sum_{j<t} W_{t,j}$. Summing $\partial\Phi_{t,j}/\partial r_i = \Phi_{t,j}/r_i$ over all pairs $j < i < t$ collapses from $\mathcal{O}(C^3)$ to two $\mathcal{O}(C^2)$ prefix sums:

$$\begin{aligned}(dM_c)_{t,j} &= -\mathbf{1}[t > j]\,(dA_c)_{t,j}\,\Phi_{t,j}, \qquad f_i \equiv \sum_{t>i}\sum_{j<i} W_{t,j} = \sum_{j<i} c_j - \sum_{t\le i} \rho_t, \\ \frac{\partial \mathcal{L}}{\partial r_i} &= \mathbf{1}[|r_i| > \varepsilon]\,\frac{f_i}{\operatorname{sign}(r_i)\max(|r_i|, \varepsilon)},\end{aligned} \tag{19}$$

giving $d\beta_i^{\mathrm{EA}} = -a_i(\partial\mathcal{L}/\partial r_i)$ and $dk_i^{\mathrm{EA}} = g_i k_{i-1} + g_{i+1}k_{i+1}$ where $g_i = -\mathbf{1}[i \ge 2 \wedge 0 < \langle k_i, k_{i-1}\rangle < 1]\,\beta_i(\partial\mathcal{L}/\partial r_i)$.

Zero-shot swap, 300-step KL healing, and from-scratch pre-training. Table 32 summarizes zero-shot and 300-step FineWeb-Edu KL-healed inference (1.23M tokens, linear-attention parameters only) on Qwen3.5, including stratified $N = 1{,}425$ MMLU (SE ≈ 1.3 pp) for 0.8B/2B zero-shot and 0.8B healed runs alongside $N = 285$ (SE ≈ 3.0 pp) for 2B/9B and 3-seed 0.8B healing:

- *Zero-shot deployment:* Restricting to exact 16×16 diagonal blocks plus either one adjacent off-diagonal block (`bt1b16`) or erase-damped off-diagonal blocks (`bea1b16`) cuts KL to 0.0024–0.0034 at $C = 64$ and matches exact $N = 1{,}425$ MMLU within 0.00–0.42 pp (48.77%/48.63% vs. 49.05% on 0.8B; 55.44%/55.37% vs. 55.44% on 2B). By contrast, matmul-free `eaadj` and `n1row` still lose 10.0–24.0 pp on zero-shot MMLU.
- 300-*step KL healing:* Healing lowers `eaadj` KL by 8×–102× and beats `n1row` by 1.7×–4.5× lower KL across all five 0.8B and 2B regimes (with 3-seed 0.8B $C = 64$ KL 0.01538 ± 0.00017 vs. 0.06973 ± 0.00111). However, on $N = 1{,}425$ MMLU, healed 0.8B still loses 4.6 pp at $C = 64$ (44.35% vs. 48.98% exact; `n1row` 32.56%) and 16.8 pp at $C = 256$ (32.00% vs. 48.77% exact; `n1row` 35.65%), while healed weights evaluated under exact recurrent decoding (`EXDC`) recover full accuracy (48.98% and 52.42%). At 9B, multi-cluster non-adjacent key echoes allow `n1row` to reach lower healed KL than `eaadj` (0.1496 vs. 0.2426).
- *From-scratch pre-training (Table 33):* On 3-seed paired 204.3M `SLA-R4` runs (52.4M FineWeb-Edu tokens, zero NaNs across all 18 runs), `exact_c128` achieves the lowest mean validation perplexity (69.94 ± 1.15) and tightest seed spread, while saving only 3.8%–7.1% step time under matmul-free solvers (62.54–64.78 vs. 67.35 ms/step). Against `n1row_c256` (71.85 ± 5.96), exact wins on 2/3 seeds (seed 0: 67.44 vs. 69.39), and `eaadj` trails `n1row` at both chunk sizes (79.61 ± 3.27 at $C = 128$, 75.95 ± 3.16 at $C = 256$) while suppressing $\beta_t > 1$ reflections (25.7% vs. 32.9%). Together with the $+0.22$ PPL cost of 1-step Neumann in our 12.5B Dolma ablation (Table 8b), these results confirm exact FP32 block-doubling as SLA's default.

I Extended TPU v7x Control Studies

To probe six methodological questions (out-of-sample roofline predictiveness, exact shape-vs.-fusion-vs.-grouping attribution against open-source Tokamax, intrinsic slow-channel behavior on public `Kimi-Linear-48B-A3B` checkpoints, statistically separated $G = 4$ vs. $G = 1$ evaluation, isolated multi-seed chunk-momentum validation, and equal-wall-clock training against open-source Tokamax

Table 32: **Exploratory approximate solvers on exact-trained Qwen3.5 checkpoints:** (a) zero-shot swap ($N = 1{,}425$ stratified MMLU on 0.8B/2B at $C = 64$; $N = 285$ on 9B and $C = 128$); (b) 300-step FineWeb-Edu KL healing (1.23M tokens, linear layers only), showing $N = 1{,}425$ MMLU where run (*) and 3-seed sample std (ddof = 1) for 0.8B $C = 64$.

(a) Zero-shot solver swap on exact-trained Qwen3.5 weights

Model / Chunk	Exact Base (FW / MMLU)	Raw `n1` (FW / KL)	`bn1b16` (FW / KL)	`n1row` (FW / KL / MMLU)	`eaadj` (FW / KL / MMLU)	`bt1b16` (FW / KL / MMLU)	`bea1b16` (FW / KL / MMLU)
0.8B, $C = 64^*$	16.08 / 49.05%	8.1×10^5 / 10.8	1020 / 4.11	48.84 / 1.10 / 25.05%	19.76 / 0.204 / 33.19%	**16.15 / 0.0034 / 48.63%**	**16.15 / 0.0031 / 48.77%**
2B, $C = 64^*$	12.83 / 55.44%	1.4×10^5 / 9.27	–	26.40 / 0.715 / 31.23%†	14.28 / 0.104 / 45.40%	**12.88 / 0.0032 / 55.37%**	**12.88 / 0.0029 / 55.44%**
9B, $C = 64^\dagger$	9.51 / 68.77%	2.1×10^5 / 10.0	865.8 / 4.51	24.94 / 0.972 / 44.91%	110.0 / 2.451 / 31.58%	**9.54 / 0.0025 / 68.42%**	**9.54 / 0.0024 / 68.07%**
0.8B, $C = 128^\dagger$	16.08 / 47.02%	diverges	1.9×10^5 / 9.34	740.7 / 3.819 / 31.58%	48.87 / 1.112 / 23.16%	16.67 / 0.0365 / 48.77%	16.55 / 0.0294 / 47.72%
9B, $C = 128^\dagger$	9.51 / 68.77%	diverges	–	129.1 / 2.618 / 33.68%	431.6 / 3.802 / 25.61%	9.69 / 0.0186 / 69.12%	9.67 / 0.0168 / 68.77%

(b) 300-step FineWeb-Edu KL healing (1.23M tokens, linear-attention parameters only)

Model, Chunk, LR	Solver	Zero-Shot FW / MMLU	Healed FW PPL	Healed KL (nats)	Top-1 Agree	OOD PG-19 PPL	Healed MMLU (`EXDC` MMLU)
0.8B, $C = 64$, `1e-5` (s_0: $N = 1{,}425^*$; $n = 3^\dagger$)	`exact` (s_0^*)	16.09 / 48.98%	16.079	0.00182	97.95%	17.834	48.63%* (–)
	`eaadj` (s_0^* / n=3†)	19.76 / 33.12%	16.271 ± 0.010	$\mathbf{0.01538 \pm 0.00017}$	93.46 ± 0.18%	18.213 ± 0.011	44.35%* [$43.16 \pm 0.70^\dagger$] (48.98%*)
	`n1row` (s_0^* / n=3†)	48.78 / 25.05%	17.110 ± 0.029	0.06973 ± 0.00111	86.04 ± 0.24%	19.789 ± 0.026	32.56%* [$33.45 \pm 0.89^\dagger$] (47.65%*)
0.8B, $C = 256$, `3e-5` (s_0: $N = 1{,}425^*$)	`exact`	16.09 / 48.77%	16.191	0.00890	95.13%	17.906	50.11%* (–)
	`eaadj`	270.0 / 23.93%	16.470	**0.03599**	90.07%	18.705	32.00%* (52.42%*)
	`n1row`	2921 / 25.12%	17.107	0.07581	85.49%	19.866	35.65%* (51.23%*)
2B, $C = 64$, `1e-5`†	`eaadj`	14.28 / 44.91%	12.972	**0.01304**	94.57%	13.786	51.58%† (52.98%†)
	`n1row`	26.36 / 30.18%	13.328	0.04179	89.50%	14.612	52.63%† (55.44%†)
2B, $C = 256$, `3e-5`†	`eaadj`	262.5 / 24.56%	13.087	**0.02937**	91.35%	14.086	40.70%† (60.00%†)
	`n1row`	1198 / 28.42%	13.372	0.05088	88.59%	14.735	47.02%† (61.40%†)
9B, $C = 64$, `1e-5`†	`eaadj`	110.1 / 30.88%	11.791	0.24264	82.04%	13.680	68.77%† (72.63%†)
	`n1row`	24.99 / 44.56%	10.907	**0.14956**	86.05%	12.279	69.82%† (72.28%†)

*Stratified $N = 1{,}425$ MMLU (SE $\approx$ 1.3 pp); †stratified $N = 285$ MMLU (SE $\approx$ 3.0 pp). At $C = 64$ ($N = 1{,}425$), smaller 8×8 blocks degrade zero-shot MMLU (`bea1b8`: 41.40% / 50.32%; `bt1b8`: 36.98% / 45.26% on 0.8B/2B), while at $C = 128$ widening blocks to 32×32 (`bt1b32`/`bea1b32`) restores KL to 0.0025–0.0026 (0.8B) and 0.0016–0.0017 (9B).

Table 33: **Paired from-scratch `SLA-R4` pre-training across exact and matmul-free solvers** (204.3M params, 12L `3L:1A`, $H_k = 2, H_v = 8, d_k = 256, d_v = 128, G = 4$, 52.4M FineWeb-Edu tokens, zero NaNs across all 18 runs; 3-seed sample std ddof = 1).

Regime	Solver & Chunk	Val PPL (Train Solver)	Per-Seed Val PPL (`s0, s1, s2`)	Val PPL (Exact Dec.)	Exact vs. Train KL	OOD PG-19 PPL	Step (ms) / $\beta_t > 1$
$\mu = 0.0$ ($n = 3$ seeds)	`exact_c128`	$\mathbf{69.94 \pm 1.15}$	[69.39, 71.26, 69.17]	$\mathbf{69.94 \pm 1.15}$	**0.00000**	$\mathbf{81.72 \pm 2.67}$	67.35 / 32.9%
	`n1row_c128`	77.58 ± 4.76	[75.94, 82.94, 73.86]	80.21 ± 6.82	0.02843	89.63 ± 4.21	63.10 / 32.1%
	`eaadj_c128`	79.61 ± 3.27	[77.28, 83.34, 78.22]	82.33 ± 2.60	0.03081	98.53 ± 7.83	64.78 / 25.7%
	`n1row_c256`	71.85 ± 5.96	[67.44, 78.63, 69.48]	72.80 ± 6.53	0.01342	83.13 ± 8.07	63.04 / 29.1%
	`eaadj_c256`	75.95 ± 3.16	[72.42, 76.92, 78.53]	78.26 ± 5.06	0.02627	89.84 ± 4.73	**62.54** / 26.5%
$\mu = 0.25$ ctrl ($n = 1$, `s0`)	`exact_c128`	**66.85**	[66.85]	**66.85**	**0.00000**	**79.28**	69.35 / 36.9%
	`n1row_c128`	77.79	[77.79]	80.74	0.03224	86.87	**64.55** / 33.2%
	`eaadj_c128`	76.50	[76.50]	78.00	0.01936	89.00	67.23 / 26.0%

Mosaic KDA), we ran six controlled single-chip studies plus a 24-run multi-slice 64-chip MoE campaign on TPU v7x (JAX `0.12`, Tokamax `0.0.14`). Tables 34–39 report all verified results alongside explicit confounders and negative diagnostics.

Table 34: **Out-of-sample predictive roofline validation on single-chip TPU v7x** (full-layer F+B p_{50} ms; 0/22 failures). Peak rooflines were locked *a priori* before benchmarking 4 anchor shapes and 18 strictly held-out designs; calibrated predictions multiply peak time by a single anchor-fitted factor ($1/\eta_{\text{cal}} = 2.950$, $\eta_{\text{cal}} = 0.339$, anchor $R^2 = 0.987$, anchor MAPE 10.5%). Representative held-out rows are shown; summary lines aggregate all 18 held-out shapes and the 12 held-out shapes in the target packed systolic regime ($R \geq 4, d_k \geq 256$).

Split	Configuration ($D = 2048$ unless noted; `Exact Inv.`)	(R, d_k, d_v, C, T)	Peak Roofline (ms)	Calibrated Pred. (ms)	Measured p_{50} (ms)	Rel. Error
Anchors ($n = 4$)	Standard `SLA-R4`, 8K ($\mu = 0.25$)	(4, 256, 256, 128, 8K)	0.867	2.559	2.894	−11.6%
	Standard `SLA-R4`, 32K ($\mu = 0.25$)	(4, 256, 256, 128, 32K)	3.472	10.241	9.971	+2.7%
	Narrow width $D = 1792$, 8K ($\mu = 0.25$)	(4, 256, 256, 128, 8K)	0.789	2.328	2.787	−16.5%
	Standard `SLA-R4`, 8K ($\mu = 0$)	(4, 256, 256, 128, 8K)	0.867	2.559	2.883	−11.3%
Held-Out (selected of $n = 18$)	**Held-out wide/long-chunk ($d_v = 256$, `ex3`)**	**(8, 512, 256, 256, 8K)**	0.867	**2.558**	**2.472**	**+3.5%**
	Held-out wide/long-chunk ($d_v = 128$, `ex3` / `eaadj`)	(8, 512, 128, 256, 8K)	0.499	1.472	1.923 / 1.726	−23.5% / −14.7%
	Held-out $R = 8, d_k = 512, C = 128$	(8, 512, 256, 128, 8K)	0.933	2.752	2.593	+6.1%
	Held-out $R = 8, d_k = 256, C = 128$	(8, 256, 256, 128, 8K)	0.802	2.367	2.277	+3.9%
	Held-out $R = 4, d_k = 512, C = 128$	(4, 512, 256, 128, 8K)	1.058	3.120	3.266	−4.5%
	Held-out $R = 4, d_k = 256, C = 256$ (`ex3`)	(4, 256, 256, 256, 8K)	0.845	2.494	2.778	−10.2%
	Held-out 16K sequence length	(4, 256, 256, 128, 16K)	1.736	5.120	5.143	−0.4%
	Held-out wide model $D = 4096$	(4, 256, 256, 128, 8K)	2.989	8.816	6.694	+31.7%
	Unpacked control $R = 1, d_k = 256$	(1, 256, 256, 128, 8K)	1.262	3.723	6.370	−41.6%
	Unpacked control $R = 1, d_k = 128$	(1, 128, 256, 128, 8K)	0.992	2.927	6.261	−53.2%

Held-out summary ($n = 18$)**:** Zero-param Spearman $\rho = \mathbf{0.893}$ (`full_fb`), 0.874 (`full_fwd`), 0.785 (`core_fb`); 1-param MAPE 19.8% / MedAPE 18.6%.
Target packed regime ($R \geq 4, d_k \geq 256$**,** $n = 12$)**:** 1-param MAPE **13.9**% / MedAPE **12.5**% (unpacked $R = 1$ shapes incur unmodeled VPU/scheduling overhead).

Table 35: **Exact 3-way multiplicative attribution against open-source Tokamax Mosaic KDA** on single-chip TPU v7x (full-layer F+B p_{50} ms; $28/28$ rows finite). We build and verify a single-kernel VMEM-fused per-channel ($G = d_k$) KDA kernel at both $R = 1$ and $R = 4$ (vs. FP32 reference: forward cosine 0.999995, rel. L_2 0.0033, all 5 VJP gradients $\cos \geq 0.999984$). Fusing $R = 1$ per-channel KDA without packing regresses to $0.44\times$–$0.69\times$ of Tokamax; once $R = 4$ head packing is applied, VMEM fusion and $G = d_k \to G = 4$ grouping each add clean multiplicative wins.

Width, Length	Tokamax KDA $R = 1$ (ms)	Tokamax KDA $R = 4$ (Sharing)	Fused $G{=}d_k$ KDA $R = 1$ (vs. Tok $R1$)	Fused $G{=}d_k$ KDA $R = 4$ (Fusion)	**Fused** `SLA-R4` $G = 4, \mu = 0$ (Grouping)	Fused GDN $G = 1, R = 4$	**Total** `SLA` / Tok $R1$
$D = 2048$, 8K	4.807	3.409 (1.410×)	10.302 (0.467×)	3.068 (1.111×)	**2.857 (1.074×)**	2.747 (1.040×)	**1.683×**
$D = 2048$, 32K	17.741	12.625 (1.405×)	40.267 (0.441×)	11.522 (1.096×)	**9.626 (1.197×)**	9.295 (1.036×)	**1.843×**
$D = 1792$, 8K	4.141	3.249 (1.275×)	6.050 (0.685×)	2.956 (1.099×)	**2.728 (1.084×)**	2.657 (1.027×)	**1.518×**
$D = 1792$, 32K	15.155	11.844 (1.280×)	24.279 (0.624×)	10.622 (1.115×)	**9.221 (1.152×)**	8.905 (1.035×)	**1.643×**

Note: At $D = 2048$, enabling $\mu = 0.25$ adds 2.5%–4.7% (2.929 ms at 8K and 10.075 ms at 32K). Parenthesized factors give the exact multiplicative chain: Sharing × Fusion × Grouping = Total (e.g. $1.405 \times 1.096 \times 1.197 = 1.843\times$ at $D = 2048$, 32K); the $G = 1$ column reports $T_{\text{SLA}(G=4)}/T_{\text{GDN}(G=1)}$.

Table 36: **Public `Kimi-Linear-48B-A3B` slow-channel decay audit on TPU v7x.** (a) Static weight-projection Monte Carlo audit across all 81,920 KDA channels (20 KDA layers) in both public checkpoints. (b) Live 27-layer forward passes of `Kimi-Linear-48B-A3B-Instruct` on FineWeb-Edu ($20/20$ KDA layers finite with Tokamax floor $g \geq -5.0$). Even after long-context pre-training, the public checkpoint develops heavy > 1M-token decay channels and improves under zero-retraining decay clamping.

(a) Static weight-projection audit (81,920 KDA channels across 20 layers)

Public Checkpoint	τ_{MC} p_{50} / p_{95}	τ_{MC} p_{99} / $p_{99.9}$	Frac. $\tau > 8$K	Frac. $\tau > 64$K	Frac. $\tau > 1$M	Contig. $G{=}4$ Var. Share
`Kimi-Linear-48B-A3B-Base`	1.07 / 180.3	3,312 / 241,116	0.606%	0.198%	0.048%	97.59%
`Kimi-Linear-48B-A3B-Instruct`	1.05 / 151.8	2,506 / 155,226	0.508%	0.160%	0.037%	97.59%

(b) Live 27-layer FineWeb-Edu evaluation of `Kimi-Linear-48B-A3B-Instruct` (zero retraining)

Sequence Length	Evaluation Mode	Full PPL	Tail-2K PPL (Δ)	Mean $\tau > 8$K (Max Lyr)	Mean $\tau > 64$K / > 1M	Max Layer $p_{99}(\tau)$
$L = 2$K	Unclamped baseline	12.614	12.614	6.70% (15.12%)	1.14% / 0.034%	170,932
$L = 8$K	Unclamped baseline	9.510	9.971 (ref)	11.03% (25.30%)	3.52% / 0.265%	1,176,794
	Zero-retrain clamp $\tau_{\max} = 4{,}096$	**9.399**	**9.701 (−0.270)**	–	–	–
	Zero-retrain clamp $\tau_{\max} = 16{,}384$	9.484	9.914 (−0.057)	–	–	–
	Naive contiguous $G = 4$ avg. (no retrain)	25.019	30.337 (+20.37)	–	–	–
$L = 16$K	Unclamped baseline	9.504	11.340 (ref)	12.12% (27.79%)	4.19% / 0.338%	1,426,272
	Zero-retrain clamp $\tau_{\max} = 4{,}096$	**9.243**	**10.905 (−0.436)**	–	–	–
	Zero-retrain clamp $\tau_{\max} = 16{,}384$	9.461	11.297 (−0.043)	–	–	–
	Naive contiguous $G = 4$ avg. (no retrain)	26.706	36.160 (+24.82)	–	–	–

Table 37: **Statistically separated $G = 4$ vs. $G = 1$ dense-LM control ($n = 6$), isolated chunk-momentum ablation ($n = 3$) with 49-point Schur stability sweep, and synthetic MT-MQAR diagnostic ($n = 5$).** Parts (a)–(b) train a 274.7M-param 12L dense LM ($D = 1024$, $H_k = 2, d_k = 256, H_v = 8$) on 104.9M FineWeb-Edu tokens (0 non-finite steps across all 27 runs; sample std ddof $= 1$). Part (c) reports a synthetic 2-layer multi-timescale associative recall diagnostic ($n = 5$ seeds, calibrated lr $= 10^{-3}$, 2,000 steps).

(a) Paired $n = 6$ dense-LM decay-granularity control (`SLA-R4`, $\mu = 0$, 2K train $\rightarrow$ 8K/16K zero-shot)

Arm	Params	Step	Val 2K PPL	Paired Δ (t)	OOD Dolma	8K Full	16K Full / Tail (Clamp)
$G = 1$ (Scalar)	274.6M	72.6 ms	47.62 ± 0.34	ref	60.89 ± 2.87	46.85	47.63 / 48.07 (47.63)
$G = 4$ **(Default)**	274.7M	73.8 ms	$\mathbf{46.47 \pm 0.32}$	$\mathbf{-1.15}$ ($\mathbf{-5.78}$)	$\mathbf{55.61 \pm 1.05}$	**45.52**	**46.08 / 46.32 (46.10)**
$G = 256$ (Per-Ch.)	279.3M	75.9 ms	46.16 ± 0.44	$-0.31^{\dagger}$ (-2.19)	53.97 ± 0.39	45.35	46.47 / 47.24 (46.20)

(b) Isolated $n = 3$ dense-LM chunk-momentum ablation ($G = 4$) and 49-point Schur stability sweep ($\varphi \in [0.4, 0.99], \beta \in [0.3, 1.99]$)

Momentum Arm	Unstable Grid	Step	Val 2K PPL	Paired Δ (t)	OOD Dolma	16K Full / Tail	Learned $\bar{\beta}$ / $\bar{\varphi}_c$ (Guard)
$\mu = 0$ **(Default)**	**0/49**	**73.9** ms	$\mathbf{46.49 \pm 0.47}$	ref	55.86 ± 1.11	**45.90 / 46.12**	0.76–0.78 / 0.05–0.06 (–)
$\mu = 0.25$ Fixed	5/49 ($\rho \leq 1.09$)	74.1 ms	46.57 ± 0.13	$+0.09$ ($+0.41$)	54.95 ± 0.30	46.27 / 46.57	0.77–0.79 / 0.05 (–)
$\mu = 0.25$ Guarded	**0/49**	74.1 ms	46.53 ± 0.36	$+0.04$ ($+0.30$)	56.85 ± 1.19	46.19 / 46.59	0.77–0.79 / 0.05 (0.000)

(c) Diagnostic synthetic 2-layer Multi-Timescale MQAR ($n = 5$ seeds, lr $= 10^{-3}$; honest negative diagnostic for coarse G)

Head Regime	$G = 1$ Joint (Short / Long)	$G = 2$ Joint (Short / Long)	$G = 4$ Joint (Short / Long)	$G = 256$ Joint (Short / Long)
$H_k = 1$ (within-head)	$26.3 \pm 28.4\%$ (33.9% / 0.8%)	$13.4 \pm 0.3\%$ (17.1% / 0.7%)	$13.4 \pm 0.3\%$ (17.1% / 0.8%)	$\mathbf{84.1 \pm 25.5}\%$ (88.1% / 70.8%)
$H_k = 2$ (`SLA-R4`)	$51.8 \pm 34.7\%$ (67.1% / 0.8%)	$26.8 \pm 28.1\%$ (34.6% / 0.7%)	$13.8 \pm 0.1\%$ (17.7% / 0.8%)	$\mathbf{62.8 \pm 39.8}\%$ (66.9% / 49.4%)

Note: In (a), † compares $G = 256$ vs. $G = 4$; per-seed ΔPPL(G=4 $-$ G=1) is $[-0.53, -1.57, -1.80, -0.94, -1.29, -0.77]$ on Val 2K (-1.149 ± 0.487, 6/6) and $[-1.17, -2.23, -2.44, -0.93, -1.61, -2.10]$ on 16K tail (-1.746 ± 0.610, $t = -7.01$, 6/6; OOD $t = -4.96$, 8K $t = -5.32$, 16K full $t = -5.48$); at 16K unclamped tail, $G = 256$ trails $G = 4$ by $+0.913 \pm 1.566$. In (b), the 5 unstable fixed-μ grid points occur at $(\varphi, \beta) \in \{(0.90, 1.99), (0.95, 1.95), (0.95, 1.99), (0.99, 1.95), (0.99, 1.99)\}$ (64-chunk norm up to 224.7×); Schur guard $\mu_c = \min(\mu, (1 - \varphi_c)/(1 + \varphi_c))$ restores stability on all 49/49 points; $p_{95}(\beta) \leq 1.20$ and $p_{95}(\varphi_c) \leq 0.26$ across all runs. In (c), at uncalibrated lr $= 2 \times 10^{-3}$ ($H_k = 2$), $G = 4$ reached $25.1 \pm 28.1\%$ (one seed at 74.4%); per-channel $G = 256$ remains substantially stronger on this synthetic two-layer task.

Table 38: **Equal-chip-hour pre-training vs. open-source Tokamax Mosaic KDA on single-chip TPU v7x** ($D = 2048$, 12L `9L:3A`, $T = 4$K, $B = 16$, FineWeb-Edu + OOD Dolma). We compare open-source Tokamax KDA (unclipped and safe-clipped $g \geq -3.5$), single-kernel fused `SLA-R4` ($\mu = 0$), and a single-kernel fused `SLA-R4 + GDN-2` prototype (decoupled erase/write gates, $\mu = 0$) under equal post-compile wall-clock budgets (180 s: $n = 3$ paired seeds; 360 s and iso-token: $n = 2$ completed pairs `s0,s1`).

Architecture (Single-Chip v7x)	Params	Step ms (Speedup)	Non-Fin. Steps	$t = 180$ s ($n = 3$) Steps / Val / OOD	$t = 360$ s ($n = 2$) Steps / Val / OOD	Paired ΔVal PPL 180 s / 360 s	Iso-Token ($n = 2$) Wall / Val / OOD
Tokamax KDA (unclipped)	821.6M	271.6	1,093–1,149	662 / 7,111–20,896 / –	1,324 / 927–6,883 / –	diverges	diverges
Tokamax KDA (clip $g \geq -3.5$)*	821.6M	271.0 (1.00×)	0	664 / 90.45 / 126.81	1,327 / 58.44 / 76.33	ref / ref	360.1 s / 58.44 / 76.33
Fused `SLA-R4` ($\mu = 0$)	760.0M	**212.3** (**1.276**×)	**0**	849 / **76.00** / 100.55	1,694 / 48.76 / 60.40	$\mathbf{-14.45}$ ($t = -43.7$) / $\mathbf{-9.68}$	**281.0** s / 54.04 / 68.37
Fused `SLA-R4 + GDN-2` ($\mu = 0$)	763.9M	221.0 (**1.226**×)	**0**	815 / 76.76 / **97.17**	1,631 / **48.51** / **59.77**	$\mathbf{-13.70}$ ($t = -19.5$) / $\mathbf{-9.93}$	292.3 s / **52.87** / **63.84**

*__Confounder disclosures:__ (1) `SLA-R4` (760.0M) and `SLA-R4 + GDN-2` (763.9M) carry 61.6M and 57.7M *fewer* parameters than Tokamax KDA (821.6M) due to 4:1 shared key projections and $G = 4$ gates. (2) Tokamax KDA follows Kimi-Linear's official $\beta_t \in (0, 1)$ write gate, whereas `SLA-R4` uses $\beta_t \in (0, 2)$ and `SLA-R4 + GDN-2` uses decoupled erase/write gates. (3) Open-source Tokamax KDA (`use_gate_in_kernel=False`) required clipping $g \geq -3.5$ ($f \geq 0.030$) to avoid float32 `exp2` backward overflow. (4) Seed `s2` safe-clipped Tokamax KDA completed $t = 180$ s (666 steps, Val 90.67) before the reservation window closed, so $t = 180$ s is complete $n = 3$ paired while $t = 360$ s and iso-token (87.0M tokens) are $n = 2$ paired (`s0,s1`); on `s2` at 360 s, fused `SLA-R4` and `SLA-R4 + GDN-2` reached 50.48 and 51.33 Val PPL.

Table 39: **Per-seed breakdown of the** 24**-run multi-slice** 64**-chip TPU v7x MoE campaign** (128 TensorCores per run, EP $= 128$, $T_{\text{train}} = 8{,}192$, 2.10M tok/step, packed Dolma 3.5, 0 non-finite steps across all 24 runs; sample std ddof $= 1$). Parts (a) and (b) detail the $n = 3$ paired equal-chip-hour and iso-token comparisons against safe-clipped open-source Tokamax Mosaic KDA ($C = 64, g \geq -3.5$) summarized in Table 9. Part (c) details all $n = 6$ paired seeds (`s0..s5`) for the $[4, 256, 512]$ `Exact Inv.` FP32 ($\mu = 0$) $G = 1$ vs. $G = 4$ sweep in Table 10.

Arm (64-Chip v7x)	Seed	Steps (Tokens)	Measured p_{50} ms (TF/s/dev)	Iso-Token PPL Early / Late	Eq.-Chip-h 11-Set PPL	8K → 32K Tail PPL	32K 4-N / 8-N	`norm9` (N=512 64K SQuAD/Hotpot)
(a) 12.5**B MoE** (`7L:1A`, $D = 1024$, 40.0 **min** = 42.67 **chip-h; iso-token checkpoints at steps** 2,000 / 4,000)								
Tokamax KDA	s0	4,730 (9.92B)	517.5 (135.8)	25.110 / 20.437	20.004	13.02 → 14.06	38.7%/31.6%	46.31% (55.7/58.0%)
(C=64, clip $g \geq -3.5$,	s1	4,880 (10.23B)	514.0 (136.7)	24.700 / 19.980	19.480	12.62 → 14.58	61.3%/54.7%	46.85% (46.5/45.1%)
sh=1024, 12.53B)	s2	4,960 (10.40B)	504.7 (139.3)	24.529 / 19.735	19.066	12.62 → 13.28	29.7%/18.0%	46.01% (57.6/59.6%)
Fused Iso-Param `SLA-R4`	s0	5,520 (11.58B)	459.0 (165.1)	25.557 / 20.654	**19.641**	**12.74 → 12.30**	**71.9%/58.6%**	**46.88%** (**65.8**/**68.2**%)
(G=4, μ=0, sh=1280,	s1	5,470 (11.47B)	458.2 (165.4)	24.710 / 19.936	**18.999**	**12.35 → 12.06**	61.3%/54.3%	**48.00%** (**66.0**/**68.8**%)
−3.0M params)	s2	5,520 (11.58B)	**448.3** (**169.0**)	24.865 / 19.899	**18.798**	**12.40 → 12.13**	**64.1%/50.0%**	**47.06%** (**66.0**/**67.0**%)
Fused `SLA-R4 + GDN-2`	s0	5,150 (10.80B)	483.1 (154.4)	**24.572 / 19.912**	**19.211**	–*	–*	–*
(G=4, μ=0, sh=1024,	s1	5,220 (10.95B)	475.7 (156.8)	24.918 / 20.152	**19.375**	12.36 → 12.15	**67.2%**/**61.7%**	47.10% (65.0/**69.5**%)
+1.7M params)	s2	5,320 (11.16B)	478.9 (155.8)	**24.290 / 19.647**	**18.741**	–*	–*	–*
(b) 63.1**B MoE** (`4L:1A`, $D = 1792$, 35.0 **min** = 37.33 **chip-h; iso-token checkpoints at steps** 600 / 1,200)								
Tokamax KDA	s0	1,250 (2.62B)	1,694.0 (216.9)	37.075 / 26.963	26.766	16.87 → 18.59	24.6%/14.8%	40.98% (44.5/42.6%)
(C=64, clip $g \geq -3.5$,	s1	1,240 (2.60B)	1,693.4 (216.9)	36.027 / 26.349	26.186	16.74 → 20.11	25.0%/16.8%	41.10% (41.6/41.2%)
sh=3584, 63.13B)	s2	1,240 (2.60B)	1,705.0 (215.5)	36.956 / 26.867	26.716	17.14 → 19.66	39.8%/28.5%	41.59% (–*)
Fused Iso-Param `SLA-R4`	s0	1,410 (2.96B)	1,506.0 (261.6)	**36.579** / 25.988	**25.141**	**15.89 → 16.19**	**45.7%**/**31.6%**	**42.02%** (–*)
(G=4, μ=0, sh=4352,	s1	1,410 (2.96B)	**1,503.0** (**262.1**)	36.078 / **25.478**	**24.613**	**15.78 → 15.98**	**39.5%**/26.2%	–*
−14.4M params)	s2	1,390 (2.91B)	1,516.0 (259.9)	37.529 / 26.251	**25.482**	16.34 → **16.14**	**62.5%**/**50.8%**	–*
Fused `SLA-R4 + GDN-2`	s0	1,350 (2.83B)	1,546.0 (247.2)	36.614 / **25.796**	25.177	–*	–*	–*
(G=4, μ=0, sh=3584,	s1	1,370 (2.87B)	1,551.0 (246.5)	**35.812** / 25.621	24.919	16.06 → 16.18	37.1%/**30.9%**	–*
−33.5M params)	s2	1,350 (2.83B)	1,549.0 (246.8)	**35.956 / 25.932**	**25.329**	**16.22** → 16.57	42.2%/27.0%	–*
(c) 12.5**B MoE** (`7L:1A`, $[4, 256, 512]$ `Exact Inv.` **FP32,** $\mu = 0$**,** 7,200 **steps** = 15.10**B tok): paired** $G = 1$ **vs.** $G = 4$ **across** $n = 6$ **seeds**								
Paired Seed (Offset)	ΔPPL (G_4−G_1)	G=1 / G=4 11-Set	G=1 8K→32K Tail	G=4 8K→32K Tail	G=1 32K 4/8-N	G=4 32K 4/8-N	G=1 / G=4 `norm9`	G=1 vs. G=4 64K QA
s0 (0K)	**−0.336**	18.210 / **17.873**	11.78 → **11.37**	**11.77** → 11.39	**73.4%**/**66.4%**	55.9%/45.7%	48.22% / **48.77%**	66.4/69.5 vs. 66.2/**71.1**%
s1 (100K)	**−0.056**	17.681 / **17.625**	**11.50** → **11.24**	11.54 → 11.28	**65.2%**/52.0%	61.3%/**52.7%**	**48.06%** / 47.42%	67.6/**72.5** vs. **68.8**/71.3%
s2 (200K)	**−0.061**	17.178 / **17.117**	11.45 → **11.17**	**11.31** → 11.22	**72.7%**/**58.6%**	33.2%/17.2%	**48.63%** / 48.42%	**69.5**/**71.7** vs. 58.6/64.5%
s3 (300K)	**−0.656**	18.845 / **18.189**	11.71 → 11.41	**11.59** → **11.28**	**73.1%**/57.0%	66.0%/**61.7%**	47.24% / **49.34%**	64.5/69.3 vs. **69.0**/**72.9**%
s4 (400K)	**−0.074**	17.445 / **17.372**	11.68 → **11.45**	**11.57** → 11.75	**66.0%**/**62.1%**	25.0%/15.2%	48.16% / **48.74%**	**71.1**/**71.9** vs. 59.4/62.3%
s5 (500K)	**−0.210**	17.572 / **17.363**	11.73 → **11.39**	**11.61** → 11.63	**71.5%**/**61.7%**	36.3%/19.5%	48.15% / **48.69%**	**69.9**/**71.3** vs. 62.1/66.6%
Mean ± **SD** (n=6)	**−0.232** ± **0.235**	17.822 / **17.590**	11.64 → **11.34**	**11.57** → 11.43	**70.3%**/**59.6%**	46.3%/35.4%	48.08% / **48.56%**	**68.2**/**71.0** vs. 64.0/68.1%

*All 24 runs completed 100% of training steps and all scheduled in-loop 11-set OLMo v3 validation evaluations with 0 non-finite losses or gradients. Dashes mark post-training long-context/downstream stages interrupted after final checkpoint/validation logging (`12.5B SLA-R4 + GDN-2 s0,s2`: PJRT teardown abort `rc = -6` during post-train eval launch; `63.1B`: reservation window cutoff during post-train `lm-eval` / 64K RULER). Paired $n = 3$ statistics vs. Tokamax KDA: at 12.5B, `SLA-R4` step time -56.91 ± 1.41 ms ($t = -69.79$, $p = 2.1 \times 10^{-4}$), final PPL -0.370 ± 0.107 ($t = -6.02$, $p = 0.0265$), 32K Tail -1.81 ± 0.68 ($p = 0.044$), `norm9` $+0.92 \pm 0.31$ ($p = 0.036$), 64K SQuAD/Hotpot $+12.7 \pm 6.0$ / $+13.7 \pm 8.7$ pts; `SLA-R4 + GDN-2` step time -32.84 ± 6.33 ms ($p = 0.012$), final PPL -0.408 ± 0.352 (3/3 wins). At 63.1B, `SLA-R4` step time -189.1 ± 1.2 ms ($p = 10^{-5}$), step-1,200 PPL -0.820 ± 0.185 ($p = 0.0165$), final PPL -1.477 ± 0.212 ($t = -12.08$, $p = 0.0068$), 32K Tail -3.35 ± 0.88 ($p = 0.022$), 32K 4/8-N $+19.4$/ $+16.1$ pts ($p = 0.016$/0.050); `SLA-R4 + GDN-2` step-1,200 PPL -0.943 ± 0.219 ($p = 0.0175$), final PPL -1.414 ± 0.163 ($t = -15.01$, $p = 0.0044$). In (c), $G = 4$ step p_{50} is 499.8 ± 4.9 ms vs. 503.7 ± 4.6 ms for $G = 1$; two-sided paired t-test on 11-set PPL gives $t = -2.42$, $p = 0.060$ (exact two-sided sign test $p = 0.031$, 6/6 seed wins).