

000 DASH: DETERMINISTIC ATTENTION SCHEDUL- 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 DASH: DETERMINISTIC ATTENTION SCHEDUL- 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 TRAINING

006 **Anonymous authors**

007 Paper under double-blind review

ABSTRACT

013 Determinism is indispensable for reproducibility in large language model (LLM)
 014 training, yet it often exacts a steep performance cost. In widely used attention
 015 implementations such as FlashAttention-3, the deterministic backward pass can
 016 incur up to a 37.9% throughput reduction relative to its non-deterministic coun-
 017 terpart, primarily because gradient accumulation operations must be serialized to
 018 guarantee numerical consistency. This performance loss stems from suboptimal
 019 scheduling of compute and gradient-reduction phases, leading to significant hard-
 020 ware underutilization.

021 To address this challenge, we formulate the backward pass of deterministic at-
 022 tention as a scheduling problem on a Directed Acyclic Graph (DAG) and derive
 023 schedules that minimize the critical path length. Building on this formulation, we
 024 present DASH(Deterministic Attention Scheduling for High-Throughput), which
 025 encapsulates two complementary scheduling strategies: (i) Descending Q-Tile It-
 026 eration, a reversed query-block traversal that shrinks pipeline stalls in causal atten-
 027 tion, and (ii) Shift Scheduling, a theoretically optimal schedule within our DAG
 028 model that reduces pipeline stalls for both full and causal masks.

029 Our empirical evaluations on NVIDIA H800 GPUs demonstrate that DASH nar-
 030 rows the performance gap of deterministic attention. The proposed strategies im-
 031 prove the throughput of the attention backward pass by up to $1.28\times$ compared to
 032 the baseline, significantly advancing the efficiency of reproducible LLM training.

1 INTRODUCTION

033 The pursuit of consistent and verifiable outcomes is a cornerstone of rigorous scientific research and
 034 large-scale engineering. In the domain of large language model (LLM) training (Wu et al., 2024),
 035 where experiments span thousands of GPUs (Grattafiori et al., 2024; DeepSeek-AI et al., 2025)
 036 and incur enormous costs, this principle of reproducibility becomes indispensable. Reproducibility
 037 empowers practitioners to diagnose training instabilities, such as loss divergence, and to evaluate
 038 the impact of architectural modifications. Consequently, deterministic training, which guarantees
 039 bitwise identical results across runs, is increasingly adopted as a standard practice for industry.

040 The origin of the non-determinism in attention of LLM training can be traced back to a fundamental
 041 yet often overlooked characteristic of computer arithmetic: the non-associativity of floating-point
 042 (FP) operations (Villa et al., 2009). For instance, $(10^8 + 10^{-6}) - 10^8$ evaluates to 0.0 in single-
 043 precision, whereas $10^8 - 10^8 + 10^{-6}$ yields the correct 10^{-6} . This sensitivity is magnified in the
 044 massively parallel environment of GPUs (Shanmugavelu et al., 2024).

045 In high-performance attention (Vaswani et al., 2023) mechanisms like FlashAttention (Dao et al.,
 046 2022), the backward pass computation is parallelized across hundreds of GPU Streaming Multipro-
 047 cessors (SMs) (NVIDIA, 2022) to maximize throughput. Each SM, running a Cooperative Thread
 048 Array (CTA), accumulates a partial contribution to gradient tensors (e.g., the gradient for the query
 049 matrix, dQ). The default high-speed approach allows these CTAs to concurrently update the final
 050 gradient in global memory via non-deterministic atomicAdd operations, as shown in Figure 1 mid-
 051 dle. This creates a non-deterministic accumulation order: the final accumulated value depends on
 052 the uncontrolled completion order of the CTAs, leading to bit-wise variations between runs.

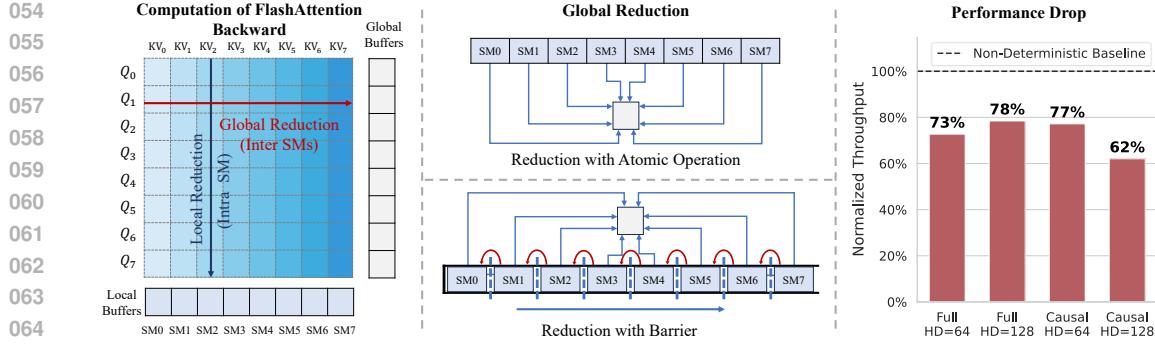


Figure 1: Overview of the deterministic FlashAttention. Left: Tiled computation structure of the backward pass, highlighting the local and global reductions. Middle: Comparison between the non-deterministic (atomic-based) and deterministic (ordered) global reduction. Right: Performance degradation under causal and full attention masks, HD stands for head dimension.

To enforce reproducibility, FlashAttention-3 (Shah et al., 2024) provides a deterministic mode. It enforces a fixed accumulation order by using synchronization barriers to force CTAs to perform their additions in a serialized order (e.g., ordered by CTA index). However, this guarantee of consistency imposes a significant performance penalty. As illustrated in Figure 1 right, enabling deterministic mode may lower throughput by up to 37.9%, leading to severe training costs when scaling LLMs across hundreds of thousands of GPUs.

This performance gap is not an inherent consequence of serialization itself. Instead, it stems from a direct conflict between the tile scheduling and a rigid, pre-determined accumulation order. As illustrated in the middle of Figure 1, the full mask scenario, commonly employed in multi-modal tasks, highlights a key inefficiency: the naive schedule creates a bottleneck by forcing reductions to start sequentially. An ideal schedule, however, would parallelize this process, allowing CTAs to begin reduction on different tiles concurrently. Crucially, this reveals that the computation schedule and the accumulation order are tightly coupled and cannot be optimized in isolation.

To address this, we introduce **Deterministic Attention Scheduling for High-throughput (DASH)**, a framework that formulates deterministic attention backward execution as an explicit scheduling optimization problem. We model the deterministic backward pass as a Directed Acyclic Graph (DAG), and formalize the objective as minimizing the DAG’s critical path length. Based on this model, we design two complementary scheduling strategies. The first, *Descending Q-Tile Iteration*, is a heuristic that processes query tiles in reverse order to advance dependency resolution and shrink pipeline bubbles in causal attention. The second strategy, a theoretically optimal algorithm we term *Shift Scheduling* is provably optimal under our DAG model. It employs a phase-shifted assignment of computational tasks to GPU multiprocessors, creating a perfectly staggered execution pattern. This ensures that the workload is perfectly balanced and that the serialized reduction operations proceed without contention while approaching the model’s theoretical utilization bound.

Our empirical evaluations on NVIDIA H800 GPUs show that DASH significantly narrows the performance gap relative to the FlashAttention-3 deterministic baseline. The two strategies deliver up to a $1.28\times$ speedup for the deterministic attention backward pass, significantly improving the efficiency of reproducible LLM training.

In summary, we made the following contributions in this paper:

- We identify the misalignment between tile execution and accumulation ordering as the principal source of performance degradation in deterministic attention.
- We provide the first DAG-based formalization of deterministic attention backward scheduling, enabling principled optimization of critical path length.
- We introduce two complementary scheduling strategies, Descending Q-Tile Iteration and Shift Scheduling, that achieve up to a $1.28\times$ speedup over the FlashAttention-3 deterministic baseline on H800 GPUs.

108 **2 BACKGROUND**110 **2.1 DETERMINISTIC FLASHATTENTION BACKWARD PASS**

111
 112 We first outline the core gradient computations in the FlashAttention backward pass: dQ , dK , and
 113 dV (Figure 1, left). During backpropagation, the gradients dK and dV are accumulated across all
 114 queries for each key (or value) position, i.e., they are reduced along the Q axis. In contrast, dQ
 115 requires a reduction across all key–value (KV) positions for each query, i.e., along the KV axis. To
 116 expose parallelism, the implementation partitions the KV dimension across SMs, allowing dK and
 117 dV to be computed within each SM via a local reduction. However, this strategy distributes partial
 118 contributions to dQ over multiple SMs, necessitating a global reduction to produce the final gradient.
 119 A conventional implementation performs this reduction using atomic additions (Figure 1, middle),
 120 which induces run-to-run variation because floating-point addition is non-associative. The resulting
 121 numerical nondeterminism undermines strict reproducibility in large-scale training. To guarantee
 122 determinism, one must enforce a prescribed accumulation order. FlashAttention-3 achieves this by
 123 performing a tile-wise sequential accumulation of dQ along the KV dimension.
 124

125 **2.2 GPU ARCHITECTURE**

126 On modern GPUs, the memory hierarchy comprises registers, shared memory, L2 cache, and global
 127 memory (NVIDIA, 2022), reflecting a fundamental capacity–latency trade-off: smaller and faster
 128 storage resides closer to the compute units. Shared memory is private to each SM, enabling low-
 129 latency intra-SM data reuse, whereas the L2 cache is globally shared, mediating inter-SM data ex-
 130 change and coherence. In datacenter-class GPUs, the L2 cache may be physically segmented, with
 131 each segment preferentially serving a subset of SMs; remote-segment accesses typically incur higher
 132 latency than local ones. This hierarchical organization materially shapes the attainable performance
 133 and the efficiency of memory-bound GPU kernels.
 134

135 **2.3 DETERMINISM IN OTHER OPERATIONS OF THE TRANSFORMER**

136 Other components, such as GEMMs, attention forward and normalization, also involve reduc-
 137 tion operations; however, the computational cost of enforcing determinism in these cases is gen-
 138 erally minimal during typical LLM training. GEMMs may exhibit nondeterministic behavior only
 139 when the reduction axis (i.e., the K-dimension) is partitioned across multiple blocks, as in split-
 140 K (NVIDIA Corporation, 2025) or stream-K (Osama et al., 2023) parallelization modes. In large-
 141 batch LLM training, parallelism along the M and N dimensions is typically sufficient to fully utilize
 142 the GPU, rendering split-K or stream-K modes unnecessary; therefore, disabling these modes gener-
 143 ally results in only a minor reduction in throughput. Similarly, other operations involving reduction,
 144 such as attention forward passes and normalizations, typically perform reductions within a single
 145 block, thereby ensuring a deterministic reduction order. Purely elementwise operations, including
 146 activation functions and bias additions, are inherently deterministic.
 147

148 **3 DASH: SCHEDULING STRATEGIES FOR DETERMINISTIC ATTENTION**

149 In this section, we introduce optimized scheduling strategies for deterministic attention. Without
 150 loss of generality, we assume that the number of KV tiles equals the number of SMs, denoted by
 151 n . When the actual number of KV tiles differs from the number of SMs, we conceptually refine or
 152 aggregate attention heads so that all SMs remain fully utilized under the same analytical framework.
 153

154 **3.1 PROBLEM FORMULATION**

155 We formalize the deterministic attention backward scheduling problem as an optimization over a
 156 directed acyclic graph (DAG), as shown in Figure 2. The DAG’s structure is constrained jointly by
 157 the dataflow of FlashAttention and the architectural characteristics of the target GPU. Our model
 158 represents a simplified abstraction of actual GPU execution; its primary purpose is to offer insights
 159 into more effective scheduling decisions, rather than to accurately predict real execution times. As
 160 such, there remain significant differences between our theoretical model and the complexities of
 161 real-world GPU behavior.
 162

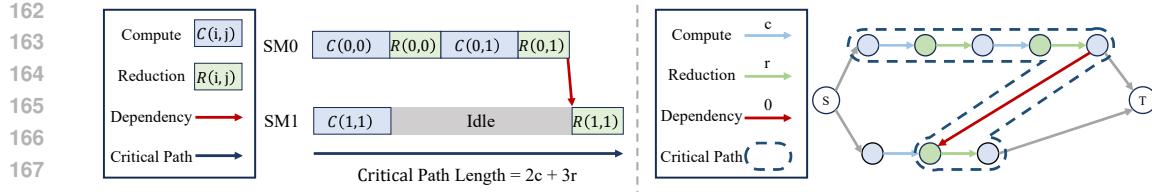


Figure 2: Visualization of the Deterministic Scheduling Problem. The Gantt chart (left) shows a naive execution schedule for a problem with two KV-tiles (i -index) and two Q-tiles (j -index). Each task consists of a compute phase $C(i, j)$ and a reduction phase $R(i, j)$. Local reductions enforce contiguous execution on a single SM (e.g., all tasks for $i = 0$ on SM0). A deterministic global reduction order introduces a cross-SM dependency (red arrow), forcing SM1 to idle and creating a pipeline bubble. The corresponding DAG (right) abstracts this schedule, where the critical path determines the end-to-end latency.

Graph Construction. Each tile-processing task is modeled as a linear path of nodes connected by edges that encode two successive phases: (i) the tile’s computation and (ii) the subsequent global reduction. These phase edges are weighted by their respective execution times, which are assumed to be constants. To encode legal accumulations and data dependencies across tiles, we insert zero-weight dependency edges between nodes of different task paths. In this way, edge weights capture quantitative duration, while the topology captures qualitative ordering constraints. The scheduling objective is to minimize the critical-path length of the resulting DAG, thereby reducing end-to-end latency and improving overall execution efficiency.

Optimization Constraint. Data movement across different memory levels incurs substantial overhead, while registers provide the fastest storage in GPUs. To leverage fast register-resident accumulation of dK and dV , all operations for a given KV tile must run contiguously on a single SM. Consequently, the edges associated with this tile form an unbroken chain, which imposes a key constraint on our optimization.

3.2 ANALYSIS OF FLASHATTENTION-3 DETERMINISTIC BACKWARD SCHEDULE

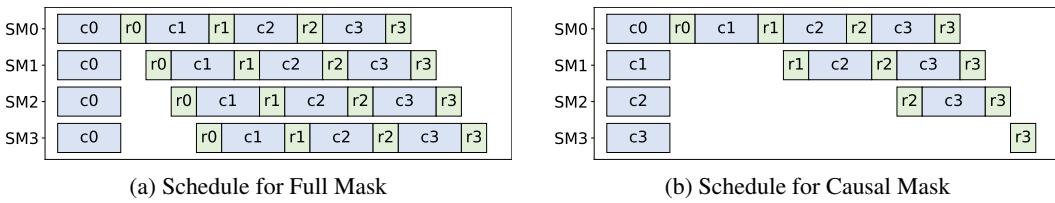


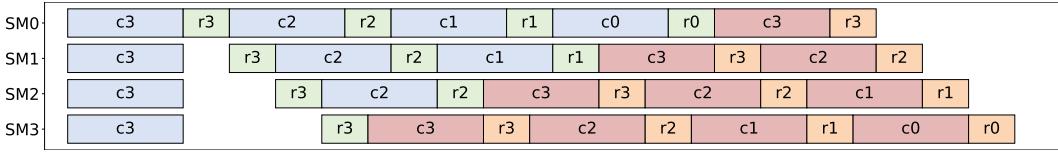
Figure 3: Backward scheduling of FlashAttention-3 for both mask shapes. Each colored segment denotes one block’s computation (cost c) followed by a reduction (cost r). Idle gaps correspond to pipeline bubbles. For clarity, since we assume the number of KV tiles equals the number of SMs, each SM processes exactly one KV tile; thus we omit the KV index in the visualization and show only the query index for each block.

Under a full attention mask, the FlashAttention-3 backward schedule achieves reasonable pipeline utilization (Figure 3a). Observable bubbles (SM idle periods) arise only during the startup phase of the first computation stage, before steady-state overlap is established. Let each stage incur a computation cost c followed by a reduction cost r . After the initial fill, each attention head sustains n sequential (computation + reduction) pairs, giving $T_{\text{steady}} = n \cdot (c + r)$ where n is the number of SMs. The startup overhead contributes an additional $(n-1) \cdot r$ due to staggered completion of the first sequence of reductions. Hence, for m heads, $T_{\text{full}} = mT_{\text{steady}} + T_{\text{startup}} = m \cdot n \cdot (c + r) + (n-1) \cdot r$, up to negligible control and synchronization overhead.

In contrast, when a causal mask is applied, the data dependencies inherent in the schedule lead to significant inefficiencies. As shown in Figure 3b, this schedule introduces a substantial bubble within the execution of **each** attention head, preventing effective pipelining. The critical path for a single

216 head becomes $T_{head_causal} = n \cdot (c + r) + (n - 1) \cdot r$. Since this inefficient pattern repeats for every
 217 head, the total execution time for m heads is approximately $T_{causal} = m \cdot T_{head_causal} + T_{startup} \approx$
 218 $m \cdot n \cdot (c + r) + (n - 1) \cdot r$.
 219

220 3.3 DESCENDING Q-TILE ITERATION: A ROBUST HEURISTIC FOR CAUSAL MASKS
 221

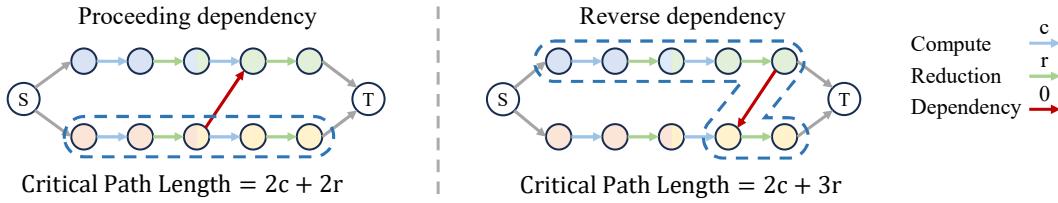


223 Figure 4: Descending (reverse-order) query tile schedule for the causal mask. Reversing the Q -block
 224 traversal accelerates dependency resolution. Colors distinguish attention heads in the pipeline.
 225

226 To mitigate the pipeline bubbles caused by causal masking, we propose a simple yet effective mod-
 227 ification: reversing the processing order of the query (Q) blocks. As illustrated in Figure 4, this
 228 reversed schedule allows most SMs to begin their computation earlier by resolving dependencies
 229 more quickly.

230 The crucial advantage of this approach is its impact on pipeline efficiency for subsequent atten-
 231 tion heads. By reversing the order, the short tasks are completed first, freeing up their SMs
 232 much earlier. Consequently, the second head can immediately begin to utilize these available re-
 233 sources, creating a tightly coupled pipeline that almost eliminates the idle gaps between heads.
 234 This sustained high utilization across an even number of m heads yields a total execution time of:
 235 $T_{reversed} \approx \frac{m \cdot (n+1)(c+r)}{2} + (n - 1) \cdot r$.
 236

237 3.4 SHIFT SCHEDULING
 238



240 Figure 5: Illustrative example for Lemma 1. Left: Added dependency (zero-weight) edges pre-
 241 serve non-decreasing depth order and do not lengthen the critical path. Right: A backward (depth-
 242 decreasing) dependency edge violates the lemma's condition and increases the critical path.
 243

244 Although the Descending Q-Tile Iteration significantly improves performance, it is natural to ask
 245 whether a theoretically optimal schedule exists. To address this, we examine the impact of introduc-
 246 ing reduction-induced inter-SM dependencies on the computation DAG's critical path.
 247

248 Disregarding (for the moment) the accumulation edges required for dQ updates, the graph decom-
 249 poses into n independent chains whose total time is minimized when their cumulative workloads
 250 are perfectly balanced. In this idealized scenario, all chains are also isomorphic, as they share an
 251 identical task structure and number of tasks. The core challenge is thus to insert the necessary
 252 zero-weight dependency edges without lengthening the original critical path. The lemma below
 253 characterizes precisely when this is possible; its proof is deferred to Appendix B for brevity.
 254

255 **lemma 1.** *Let $G_0 = (V, E_0)$ be a DAG consisting of a single source node s , a single sink node t , and
 256 $n \geq 1$ parallel, isomorphic chains connecting s to t . All edge weights in E_0 are strictly positive. Let
 257 the depth of a node v , denoted $depth(v)$, be the number of edges on the unique path from s to v within
 258 its chain in G_0 . Let a sequence of graphs G_1, \dots, G_k be generated such that $G_i = (V, E_{i-1} \cup \{e_i\})$,
 259 where each $e_i = (u_i, v_i)$ is a zero-weight edge. We add the explicit condition that every new graph
 260 G_i in the sequence must remain a DAG.*

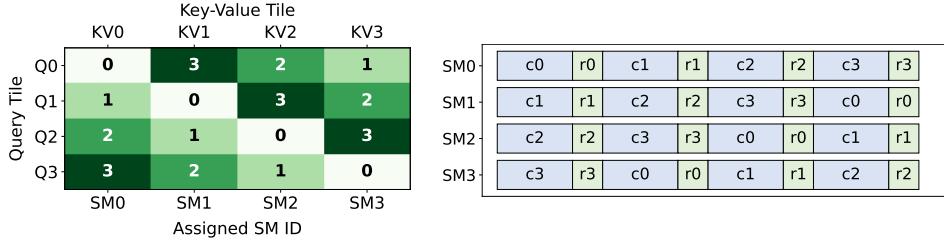


Figure 6: Optimal full-mask schedule via cyclic shifting. Left: Cyclic visiting order of Q tiles per SM; distinct timestamps (i.e., the value in each box) on each row induce a natural, conflict-free reduction sequence for every dQ block. Right: Simulated timeline showing fully balanced utilization without additional bubbles.

Under this condition, the critical path length of G_k is equal to that of G_0 if and only if for every added edge $e_i = (u_i, v_i)$ for $i \in \{1, \dots, k\}$, the condition $\text{depth}(u_i) \leq \text{depth}(v_i)$ holds.

As illustrated in Figure 5, Lemma 1 dictates that to preserve the original critical path length, any added dependency edge (u, v) must satisfy the condition $\text{depth}(u) \leq \text{depth}(v)$. This formal constraint translates to a critical physical limitation: for any given query tile Q_j , the tasks involving it cannot be executed in parallel.

A schedule that assigns two tiles contributing to the same dQ_j —say (KV_i, Q_j) and (KV_k, Q_j) —to execute concurrently on different SMs would create a resource conflict during their reduction phases. Resolving this conflict requires serializing the reductions, for instance, forcing the reduction for (KV_k, Q_j) to wait for the one from (KV_i, Q_j) to complete, or vice versa. Because the conflicting reduction tasks would otherwise start at the same depth in the DAG, this forced serialization introduces a dependency edge (u, v) where $\text{depth}(u) > \text{depth}(v)$ —from the completion of the first reduction to the start of the second. This directly violates the lemma’s condition and sub-optimally extends the critical path.

Our objective is thus twofold: first, to balance the workload across SMs, and second, to devise a conflict-free reduction order that adheres to the lemma’s constraint.

Optimal Schedule for Full Masks Under a full mask, per-KV-tile workloads are uniform, allowing for immediate balancing. To satisfy the second objective, we employ a *Shift Scheduling*, as illustrated in Figure 6. In this schedule, SM_i processes KV blocks in the order $(i, i+1, \dots, n-1, 0, \dots, i-1)$. This cyclical assignment inherently creates a conflict-free, sequential ordering for the reductions on any given dQ block, directly satisfying the lemma’s condition. As both workload balancing and conflict-free reduction are achieved, this schedule is theoretically optimal.

Symmetric Shift Scheduling for Causal Masks Causal masking induces a strongly imbalanced workload: early KV blocks participate in the full set of query interactions, whereas later blocks contribute progressively fewer operations, yielding workloads that decrease linearly across the sequence.

We address this by *Symmetric Shift Scheduling*. Its core is a symmetric pairing principle: SMs jointly handle KV blocks i and $n-1-i$, pairing the longest with the shortest, the second-longest with the second-shortest, and so forth. This pairing equalizes task chain lengths per SM, restoring near-perfect balance.

We operationalize symmetric pairing via a two-phase schedule. In Phase 1, a cyclic shift is applied to the dense lower-left rectangle, efficiently filling the pipeline. Phase 2 addresses the residual triangles using a purely analytical model of workload folding, where tasks from the lower-right are logically mapped to the upper-left’s masked slots to form a conceptual square without any data movement. The operational sequence—a top-down traversal of the left triangle and a bottom-up traversal of the right—is algebraically equivalent to a diagonal-initialized shift schedule on this conceptual square. This equivalence is key: it preserves workload balance, ensures contiguous computation for each KV block, enforces depth-monotone accumulation to satisfy Lemma 1, and ultimately eliminates all pipeline bubbles.

324
 325 **Summary of Optimal Performance** In summary, the
 326 proposed scheduling strategies achieve theoretical opti-
 327 mality for both scenarios. By perfectly balancing work-
 328 loads and eliminating pipeline bubbles, the total execu-
 329 tion time for m heads is: Full Mask: $T_{full_opt} = m \cdot n \cdot$
 $(c + r)$; Causal Mask: $T_{causal_opt} = \frac{m \cdot (n+1) \cdot (c+r)}{2}$

331 4 EXPERIMENTS

333 In this section, we empirically evaluate the performance
 334 of our proposed scheduling strategies under full and
 335 causal masks. We measure throughput under various se-
 336 quence lengths and analyze how architectural factors in-
 337 teract with different scheduling choices.

339 4.1 EXPERIMENTAL SETUP

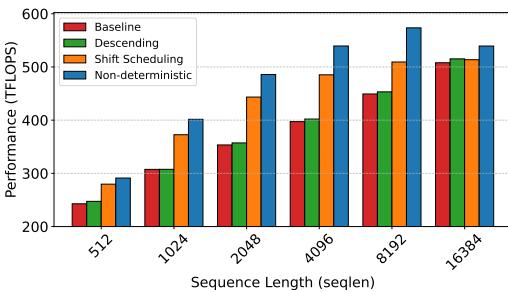
341 **Hardware and Software.** All experiments are con-
 342 ducted on a server equipped with NVIDIA H800 GPUs,
 343 CUDA version 12.6 and Triton (Tillet et al., 2019) ver-
 344 sion 3.4. All kernels are implemented by extending the
 345 FlashAttention-3 implementation.

347 **Baseline and Proposed Methods.** We compare our
 348 methods against the deterministic backward pass of
 349 FlashAttention-3, which serves as our primary baseline.
 350 We also benchmark against the Triton tutorial’s imple-
 351 mentation for causal attention, as its public version lacks a full-mask counterpart. We omit
 352 FlashAttention-2 because prior published benchmarks (Shah et al., 2024) on Hopper-class GPUs
 353 show it is consistently outperformed by FlashAttention-3, and thus it no longer constitutes a com-
 354 petitive baseline. The methods under evaluation are:

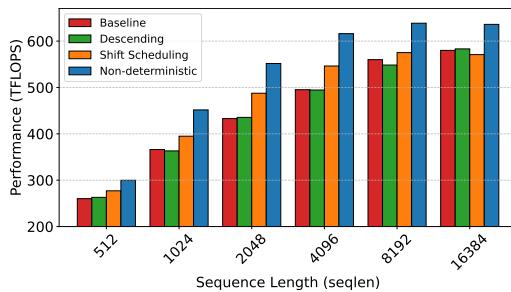
- 355 • **Descending Q-Tile Iteration** (for both masks)
- 356 • **Shift Scheduling** (for full masks)
- 357 • **Symmetric Shift Scheduling** (for causal masks)

359 **Benchmark Settings** Following the methodology of the FlashAttention-3 study, we evaluate per-
 360 formance by fixing the total number of tokens at 16,384 while varying the sequence length from 512
 361 to 16,384. Similarly, we fix the hidden dimension to be 2,048, and test different head dimensions in
 362 64 and 128. All the results are tested using BF16 precision random inputs.

364 4.2 PERFORMANCE ON FULL ATTENTION MASKS

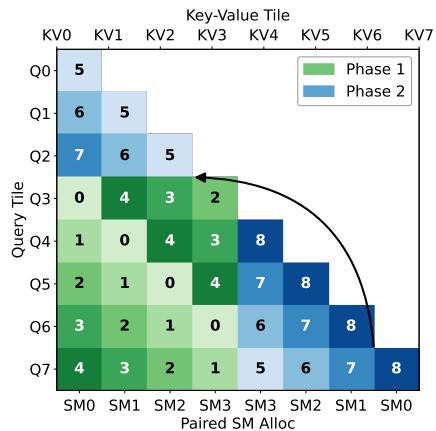


376 (a) Full mask, $headdim = 64$.



376 (b) Full mask, $headdim = 128$.

377 Figure 8: Backward-pass throughput under full attention masks.



378 Figure 7: Optimal causal-mask sched-
 379 ule using symmetric shift and two-phase work-
 380 load folding. Phase 1 processes the dense lower-left rect-
 381 angle; Phase 2 folds the remaining triangles into a
 382 logical square and traverses it starting from the main dia-
 383 gonal, first covering the upper-left portion before the lower-
 384 right, ensuring each KV block is exe-
 385 cuted contiguously.

378 Figure 8 presents the throughput comparison for the full attention mask scenario. Our Shift Schedul-
 379 ing consistently outperforms the FlashAttention-3 baseline across most sequence lengths, demon-
 380 strating the effectiveness of our theoretically optimal approach. However, a notable exception occurs
 381 at the maximum sequence length of 16,384, where its performance slightly degrades relative to the
 382 baseline.

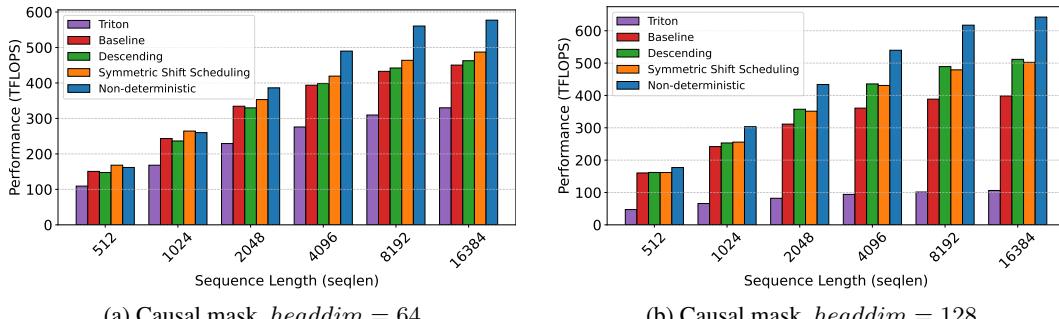
383 This phenomenon highlights a divergence between our theoretical model and practical hardware
 384 execution. Our model assumes zero-cost dependency edges, but in reality, inter-SM communication
 385 for synchronizing reduction operations is mediated by the L2 cache. This incurs significant latency,
 386 ranging from approximately 200 cycles for accesses to the local L2 cache segment to over 500 cycles
 387 for remote segment accesses on H800-class GPUs (Luo et al., 2025). This latency differential is a
 388 direct consequence of the distributed L2 cache architecture described in Section 2.

389 At a sequence length of 16,384 and a KV block size of 128, the computation for a single head is dis-
 390 tributed across 128 blocks, often mapped to 128 SMs. This high degree of parallelism necessitates
 391 frequent cross-SM communication to signal task completion. Given the large number of partici-
 392 pating SMs, a substantial portion of these synchronization signals must traverse the higher-latency
 393 links to a remote L2 cache segment. The Shift Scheduling, with its more intricate dependency graph
 394 compared to the simpler, linear dependency of the baseline, becomes more sensitive to this com-
 395 munication overhead at extreme parallelism. This increased synchronization cost, dominated by
 396 remote L2 accesses, ultimately outweighs the computational benefits of the schedule in this specific
 397 high-parallelism, long-sequence scenario, leading to the observed performance degradation.

398

399 4.3 PERFORMANCE ON CAUSAL ATTENTION MASKS

400

(a) Causal mask, $headdim = 64$ (b) Causal mask, $headdim = 128$

410

411 Figure 9: Backward-pass throughput under causal attention masks.

412

413 The performance evaluation for causal attention masks, presented in Figure 9, confirms the efficacy
 414 of our proposed methods. Both the Descending Q-Tile Iteration and our theoretically optimal Sym-
 415 metric Shift Scheduling demonstrate a throughput improvement over the FlashAttention-3 baseline
 416 across all tested configurations.

417

An interesting trade-off emerges when comparing our two proposed methods at different head di-
 418 mensions (headdim). At $headdim = 64$, the Symmetric Shift Scheduling achieves the highest
 419 performance, validating the benefits of its superior workload balancing. *However, the descending*
 420 *schedule does not perform very well in this case. This is because in the FlashAttention-3 causal*
 421 *backward kernel, the L2-aware LPT scheduler interleaves multiple heads across SMs. When head-*
 422 *dim = 64 and the sequence length is short, each head's L2 footprint remains small, allowing many*
 423 *heads to reside in cache with only 1-2 tiles in flight per head. Consequently, the causal stalls tar-*
 424 *getted by Descending Q-Tile Iteration are largely masked by cross-head interleaving, resulting in*
 425 *only marginal net performance gains.*

426

427 However, at $headdim = 128$, Symmetric Shift Scheduling's performance is surpassed by the sim-
 428 pler Descending Q-Tile Iteration. This performance inversion is attributable to a critical interaction
 429 between algorithmic complexity and GPU resource limitations, specifically register pressure. The
 430 Symmetric Shift Scheduling, while algorithmically optimal, requires a more complex implemen-
 431 tation to manage the state of the folded task space. This complexity translates to higher register usage
 432 per thread to maintain additional loop counters and intermediate states.

When $headdim = 128$, the base register requirement for storing accumulators and other intermediate values is already substantial. The additional overhead (around 10 registers) from our optimal schedule can push the total register count per thread beyond the hardware’s physical limit, as shown by Nsight Compute (NVIDIA Corporation, 2024). This forces the compiler to generate code that spills registers, offloading their contents to the much slower local memory. The high latency incurred by these spill-induced memory operations introduces significant execution stalls, which negate the algorithmic benefits of the more balanced workload and lead to degraded performance. In contrast, the simpler Descending Q-Tile Iteration operates below this critical register pressure threshold, thereby avoiding spilling and achieving better effective performance in this high-resource-demand scenario. Therefore, the two schedules for causal masks are complementary: Symmetric Shift is theoretically optimal under our DAG model, while Descending is the practically preferred choice for large head dimensions on current GPUs.

In the future, Symmetric Shift’s theoretical advantages are expected to be fully realized on newer architectures with greater on-chip resources (such as Blackwell GPUs with TMEM, or devices equipped with larger register files), or under kernel designs that are less constrained by register allocation than the present FlashAttention-3 implementation.

4.4 END-TO-END PERFORMANCE

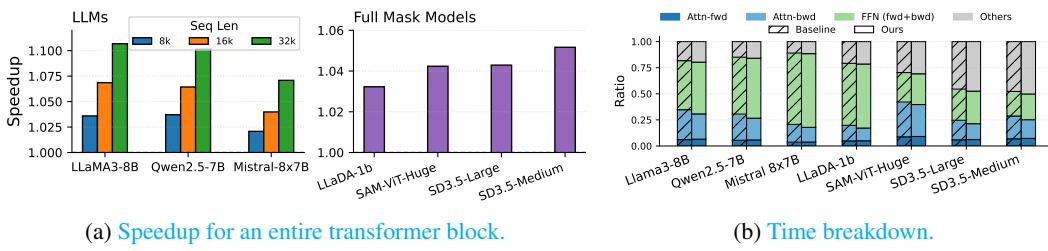


Figure 10: End-to-end performance of a transformer block.

To assess the performance gains delivered by DASH during training, we measured the runtime required to process an entire transformer block, accounting for both forward and backward passes.

We evaluated DASH across a range of widely adopted models. For causal mask scenarios, we selected famous LLMs: LLaMA3-8b (Grattafiori et al., 2024), Qwen2.5-7b (Qwen et al., 2025), and Mistral-8×7b (Jiang et al., 2024). For full mask scenarios, we included the vision model SAM-huge (Kirillov et al., 2023), the diffusion models StableDiffusion3.5 (medium and large) (AI, 2024), and the diffusion-based language model LLaDA-1b (Nie et al., 2025).

For LLMs, we employ a batch size of 1 with sequence lengths of 8k, 16k, and 32k. In the case of full mask models, a batch size of 16 is used, with the training sequence length fixed at 4k in accordance with standard architectural configurations. The relative speedup achieved by our approach compared to the baseline is illustrated in Figure 10a. For causal models, we observe end-to-end performance improvements ranging from 2% to 10%. Full mask models also exhibit a speedup of approximately 4%. In summary we achieved an average speedup of around 5%, which aligns with our internal training experience on thousands of GPUs. Additionally, Figure 10b provides a detailed breakdown of computation time across different kernel operations, with causal models evaluated at a sequence length of 16k.

4.5 IMPACT OF DETERMINISM ON NUMERICAL STABILITY

Table 1: Max gradient deviation averaged over 10 identical backward passes; $M_r = \max |g_r - g_{\text{ref}}|$.

Masking Scheme	Non-deterministic	Deterministic
Full	2.4×10^{-4}	0
Causal	4.9×10^{-4}	0

486 Our analysis of backward passes indicates that non-deterministic kernels cause run-to-run gradient
 487 deviations of $O(10^{-4})$, while deterministic ones guarantee bitwise identical outcomes (Table 1).
 488 Although small, this variability can accumulate, so determinism is key to achieving reproducibility.
 489

490 5 RELATED WORKS

493 **FlashAttention and Kernel-Level I/O Optimization** Early optimization of attention focused on
 494 mitigating the I/O bottleneck imposed by the quadratic attention matrix. FlashAttention (Dao et al.,
 495 2022) introduced an I/O-aware tiled and fused kernel that avoids materializing the full attention
 496 matrix in HBM. FlashAttention-2 and 3 (Dao, 2023; Shah et al., 2024) further improved utilization
 497 via refined work partitioning and leveraged specialized hardware for asynchronous data movement.
 498

499 **Low-Precision Attention** Low-precision methods further reduce bandwidth and memory cost.
 500 The SageAttention series (Zhang et al., 2025b;a;c) systematically explores progressively lower for-
 501 mats while maintaining accuracy.

502 **Inference-Oriented Attention Kernels** Inference-specialized kernels include FlashDecoding and
 503 FlashDecoding++(Dao et al., 2023; Hong et al., 2024) for autoregressive decoding, PodAtten-
 504 tion(Kamath et al., 2025) for mixed prefilling/decoding, and DeFT (Yao et al., 2025) and Fast-
 505 Tree (Pan et al., 2025) for tree-structured generation.
 506

507 **Distributed Cyclic Scheduling** Our shift-based scheduling is inspired by cyclic (ring-
 508 style) phase-shift patterns long used in distributed systems. Distributed attention algo-
 509 rithms—RingAttention (Liu et al., 2023), StripedAttention (Brandon et al., 2023), and Loong-
 510 Train (Gu et al., 2024)—adopt related cyclic schemes to overlap communication and computation
 511 across devices, whereas we apply a shift strategy intra-GPU to co-optimize deterministic accumula-
 512 tion and work balance.
 513

514 **Deterministic Implementations** Existing deterministic implementations either split dK , dV and
 515 dQ calculation into different passes (e.g., Triton tutorials (Tillet et al., 2019))—forcing a second K/V
 516 read—or materialize per-tile dQ partials for later consolidation (FlashAttention-2), adding memory
 517 footprint and an extra reduction kernel. These designs trade bandwidth or memory rather than co-
 518 optimizing execution and accumulation order, which is the focus of our approach.
 519

520 **Determinism in Inference** Determinism for inference has also been examined: He & Lab (2025)
 521 attribute non-reproducibility to lack of “batch invariance,” where outputs depend on batch size, and
 522 design batch-invariant kernels. Their goal differs from ours: we target training time run-to-run
 523 determinism, where batch configurations are fixed to ensure reproducibility.
 524

525 6 CONCLUSION

528 In this work, we addressed the significant performance penalty associated with the deterministic
 529 backward pass in modern attention mechanisms. By formulating the computation as a scheduling
 530 problem on a DAG, we introduced DASH, a framework featuring two distinct and complementary
 531 strategies. The first, Descending Q-Tile Iteration, provides a simple yet remarkably effective heuris-
 532 tic that accelerates causal attention. The second, derived from our conflict-free scheduling lemma,
 533 represents a theoretically optimal solution.

534 Our empirical evaluation not only demonstrates that DASH significantly narrows the performance
 535 gap, improving throughput by up to $1.28\times$ over the baseline, but more importantly, it reveals a
 536 crucial insight: theoretical optimality does not always translate to practical superiority. We iden-
 537 tified hardware realities, such as register pressure and inter-SM communication latency, as critical
 538 factors that can override the benefits of a more complex, algorithmically perfect schedule. By pro-
 539 viding a suite of solutions catering to different scenarios, DASH enables practitioners to achieve
 high throughput attention in reproducible LLM training.

540 REFERENCES

541

542 Stability AI. Stable diffusion 3.5, 2024. URL <https://github.com/Stability-AI/sd3>.
543 5. Accessed: Nov. 2025.

544

545 William Brandon, Aniruddha Nrusimha, Kevin Qian, Zachary Ankner, Tian Jin, Zhiye Song, and
546 Jonathan Ragan-Kelley. Striped attention: Faster ring attention for causal transformers, 2023.
547 URL <https://arxiv.org/abs/2311.09431>.

548

549 Tri Dao. Flashattention-2: Faster attention with better parallelism and work partitioning, 2023. URL
550 <https://arxiv.org/abs/2307.08691>.

551

552 Tri Dao, Daniel Y. Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. Flashattention: Fast and
553 memory-efficient exact attention with io-awareness, 2022. URL <https://arxiv.org/abs/2205.14135>.

554

555 Tri Dao, Daniel Haziza, Francisco Massa, and Grigory Sizov. Flash-decoding for long-
556 context inference. PyTorch Blog, October 2023. URL <https://pytorch.org/blog/flash-decoding/>.

557

558 DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Cheng-
559 gang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang,
560 Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting
561 Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Haowei Zhang, Honghui
562 Ding, Huajian Xin, Huazuo Gao, Hui Li, Hui Qu, J. L. Cai, Jian Liang, Jianzhong Guo, Jiaqi
563 Ni, Jiashi Li, Jiawei Wang, Jin Chen, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li,
564 Junxiao Song, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang,
565 Lecong Zhang, Lei Xu, Leyi Xia, Liang Zhao, Litong Wang, Liyue Zhang, Meng Li, Miaojun
566 Wang, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Mingming Li, Ning Tian, Panpan
567 Huang, Peiyi Wang, Peng Zhang, Qiancheng Wang, Qihao Zhu, Qinyu Chen, Qiushi Du, R. J.
568 Chen, R. L. Jin, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, Runxin Xu, Ruoyu Zhang,
569 Ruyi Chen, S. S. Li, Shanghao Lu, Shangyan Zhou, Shanhua Chen, Shaoqing Wu, Shengfeng
570 Ye, Shengfeng Ye, Shirong Ma, Shiyu Wang, Shuang Zhou, Shuiping Yu, Shunfeng Zhou, Shut-
571 ing Pan, T. Wang, Tao Yun, Tian Pei, Tianyu Sun, W. L. Xiao, Wangding Zeng, Wanjia Zhao,
572 Wei An, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, X. Q. Li, Xiangyue
573 Jin, Xianzu Wang, Xiao Bi, Xiaodong Liu, Xiaohan Wang, Xiaojin Shen, Xiaokang Chen, Xi-
574 aokang Zhang, Xiaosha Chen, Xiaotao Nie, Xiaowen Sun, Xiaoxiang Wang, Xin Cheng, Xin
575 Liu, Xin Xie, Xingchao Liu, Xingkai Yu, Xinnan Song, Xinxia Shan, Xinyi Zhou, Xinyu Yang,
576 Xinyuan Li, Xuecheng Su, Xuheng Lin, Y. K. Li, Y. Q. Wang, Y. X. Wei, Y. X. Zhu, Yang
577 Zhang, Yanhong Xu, Yanhong Xu, Yanping Huang, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui
578 Li, Yaohui Wang, Yi Yu, Yi Zheng, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Ying
579 Tang, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yu Wu,
580 Yuan Ou, Yuchen Zhu, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yukun Zha, Yunfan
581 Xiong, Yunxian Ma, Yuting Yan, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Z. F.
582 Wu, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhen Huang, Zhen Zhang, Zhenda
583 Xie, Zhengyan Zhang, Zhewen Hao, Zhibin Gou, Zhicheng Ma, Zhigang Yan, Zhihong Shao,
584 Zhipeng Xu, Zhiyu Wu, Zhongyu Zhang, Zhuoshu Li, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li,
585 Ziwei Xie, Ziyang Song, Ziyi Gao, and Zizheng Pan. Deepseek-v3 technical report, 2025. URL
586 <https://arxiv.org/abs/2412.19437>.

587

588 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad
589 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan,
590 Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Ko-
591 rennev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava
592 Spataru, Baptiste Roziere, Bethany Biron, Bin Tang, Bobbie Chern, Charlotte Caucheteux,
593 Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret,
594 Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius,
595 Daniel Song, Danielle Pintz, Danny Livshits, Danny Wyatt, David Esiobu, Dhruv Choudhary,
596 Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab
597 AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco

594 Guzmán, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind That-
 595 tai, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Kore-
 596 vaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra,
 597 Ivan Evtimov, Jack Zhang, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Ma-
 598 hadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu,
 599 Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jong-
 600 so Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasudevan Alwala,
 601 Karthik Prasad, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid
 602 El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Kushal Lakhotia, Lauren
 603 Rantala-Yearly, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin,
 604 Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi,
 605 Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Maria Tsimpoukelli, Mathew
 606 Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Ku-
 607 mar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoy-
 608 chev, Niladri Chatterji, Ning Zhang, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan
 609 Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan,
 610 Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ra-
 611 mon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohan Maheswari, Ro-
 612 hit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan
 613 Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell,
 614 Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Raparthy, Sheng
 615 Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer
 616 Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman,
 617 Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mi-
 618 haylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor
 619 Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vitor Albiero, Vladan Petrovic, Weiwei
 620 Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaofang
 621 Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao Jia, Xuewei Wang, Yaelle Gold-
 622 schlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning
 623 Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh,
 624 Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria,
 625 Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alexei Baevski, Allie Feinstein,
 626 Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu, Andres Alvarado, An-
 627 drew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, An-
 628 nie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel,
 629 Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leon-
 630 hardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu
 631 Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Mon-
 632 talvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changhan Wang, Changkyu Kim, Chao
 633 Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Cynthia
 634 Gao, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, David Adkins, David Xu, Davide
 635 Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkang Wang, Duc Le,
 636 Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily
 637 Hahn, Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smo-
 638 thers, Fei Sun, Felix Kreuk, Feng Tian, Filippos Kokkinos, Firat Ozgenel, Francesco Caggioni,
 639 Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia
 640 Swee, Gil Halpern, Grant Herman, Grigory Sizov, Guangyi Zhang, Guna Lakshminarayanan,
 641 Hakan Inan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harri-
 642 son Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Hongyuan Zhan, Ibrahim Damlaj,
 643 Igor Molybog, Igor Tufanov, Ilias Leontiadis, Irina-Elena Veliche, Itai Gat, Jake Weissman, James
 644 Geboski, James Kohli, Janice Lam, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jen-
 645 nifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang,
 646 Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Jun-
 647 jie Wang, Kai Wu, Kam Hou U, Karan Saxena, Kartikay Khandelwal, Katayoun Zand, Kathy
 648 Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang,
 649 Kunal Chawla, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell,
 650 Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madiyan Khabsa,
 651 Manav Avalani, Manish Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias
 652 Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Miao Liu, Michael L.

648 Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike
 649 Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari,
 650 Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan
 651 Singhal, Nick Egebo, Nicolas Usunier, Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong,
 652 Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent,
 653 Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar,
 654 Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Ro-
 655 driguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Rangaprabhu Parthasarathy,
 656 Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin
 657 Mehta, Sachin Siby, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon,
 658 Sasha Sidorov, Satadru Pan, Saurabh Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ra-
 659 maswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha,
 660 Shishir Patil, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal,
 661 Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satter-
 662 field, Sudarshan Govindaprasad, Sumit Gupta, Summer Deng, Sungmin Cho, Sunny Virk, Suraj
 663 Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo
 664 Koehler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook
 665 Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Ku-
 666 mar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov,
 667 Wei Li, Wencheng Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiao-
 668 jian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv Kleinman, Yanjun Chen, Ye Hu, Ye Jia,
 669 Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yu Zhao,
 670 Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhao-
 671 duo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. The llama 3 herd of models, 2024. URL
<https://arxiv.org/abs/2407.21783>.

672 Diandian Gu, Peng Sun, Qinghao Hu, Ting Huang, Xun Chen, Yingtong Xiong, Guoteng Wang,
 673 Qiaoling Chen, Shangchun Zhao, Jiarui Fang, Yonggang Wen, Tianwei Zhang, Xin Jin, and Xu-
 674 anzhe Liu. Loongtrain: Efficient training of long-sequence llms with head-context parallelism,
 675 2024. URL <https://arxiv.org/abs/2406.18485>.

676 Horace He and Thinking Machines Lab. Defeating nondeterminism in llm inference. *Thinking Machines Lab: Connectionism*, 2025. doi: 10.64434/tml.20250910.
 677 <https://thinkingmachines.ai/blog/defeating-nondeterminism-in-llm-inference/>.

678 Ke Hong, Guohao Dai, Jiaming Xu, Qiuli Mao, Xiuhong Li, Jun Liu, Kangdi Chen, Yuhan Dong,
 679 and Yu Wang. Flashdecoding++: Faster large language model inference on gpus, 2024. URL
<https://arxiv.org/abs/2311.01282>.

680 Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris
 681 Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gi-
 682 anna Lengyel, Guillaume Bour, Guillaume Lample, Lélio Renard Lavaud, Lucile Saulnier, Marie-
 683 Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le
 684 Scao, Théophile Gervet, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed.
 685 Mixtral of experts, 2024. URL <https://arxiv.org/abs/2401.04088>.

686 Aditya K. Kamath, Ramya Prabhu, Jayashree Mohan, Simon Peter, Ramachandran Ramjee, and
 687 Ashish Panwar. Pod-attention: Unlocking full prefill-decode overlap for faster llm inference. In
 688 *Proceedings of the 30th ACM International Conference on Architectural Support for Program-
 689 ming Languages and Operating Systems, Volume 2*, ASPLOS '25, pp. 897–912. ACM, March
 690 2025. doi: 10.1145/3676641.3715996. URL <http://dx.doi.org/10.1145/3676641.3715996>.

691 Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete
 692 Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick.
 693 Segment anything, 2023. URL <https://arxiv.org/abs/2304.02643>.

694 Hao Liu, Matei Zaharia, and Pieter Abbeel. Ring attention with blockwise transformers for near-
 695 infinite context, 2023. URL <https://arxiv.org/abs/2310.01889>.

702 Weile Luo, Ruibo Fan, Zeyu Li, Dayou Du, Hongyuan Liu, Qiang Wang, and Xiaowen Chu. Dis-
 703 secting the nvidia hopper architecture through microbenchmarking and multiple level analysis,
 704 2025. URL <https://arxiv.org/abs/2501.12084>.

705 Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai Lin,
 706 Ji-Rong Wen, and Chongxuan Li. Large language diffusion models, 2025. URL <https://arxiv.org/abs/2502.09992>.

709 NVIDIA. NVIDIA H100 Tensor Core GPU Architecture. Technical report, NVIDIA,
 710 mar 2022. URL <https://www.nvidia.com/content/dam/en-zz/Solutions/gtc22/data-center/h100/gtc22-whitepaper-hopper.pdf>. White paper.

713 NVIDIA Corporation. *NVIDIA Nsight Compute*, 2024. URL <https://developer.nvidia.com/nsight-compute>. Version 2022.4.

716 NVIDIA Corporation. Split-k gemm. https://github.com/NVIDIA/cutlass/tree/main/examples/06_splitK_gemm, 2025.

718 Muhammad Osama, Duane Merrill, Cris Cecka, Michael Garland, and John D. Owens. Stream-
 719 k: Work-centric parallel decomposition for dense matrix-matrix multiplication on the gpu. In
 720 *Proceedings of the 28th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel
 721 Programming*, PPoPP '23, pp. 429–431, New York, NY, USA, 2023. Association for Computing
 722 Machinery. ISBN 9798400700156. doi: 10.1145/3572848.3577479. URL <https://doi.org/10.1145/3572848.3577479>.

724 Zaifeng Pan, Yitong Ding, Yue Guan, Zheng Wang, Zhongkai Yu, Xulong Tang, Yida Wang, and
 725 Yufei Ding. Fasttree: Optimizing attention kernel and runtime for tree-structured LLM inference.
 726 In *Eighth Conference on Machine Learning and Systems*, 2025. URL <https://openreview.net/forum?id=BwvHcHZ3kJ>.

729 Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan
 730 Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang,
 731 Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin
 732 Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li,
 733 Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang,
 734 Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report, 2025.
 735 URL <https://arxiv.org/abs/2412.15115>.

736 Jay Shah, Ganesh Bikshandi, Ying Zhang, Vijay Thakkar, Pradeep Ramani, and Tri Dao.
 737 Flashattention-3: Fast and accurate attention with asynchrony and low-precision, 2024. URL
 738 <https://arxiv.org/abs/2407.08608>.

739 Sanjiv Shanmugavelu, Mathieu Taillefumier, Christopher Culver, Oscar Hernandez, Mark Coletti,
 740 and Ada Sedova. Impacts of floating-point non-associativity on reproducibility for hpc and deep
 741 learning applications, 2024. URL <https://arxiv.org/abs/2408.05148>.

743 Philippe Tillet, H. T. Kung, and David Cox. Triton: an intermediate language and compiler for
 744 tiled neural network computations. In *Proceedings of the 3rd ACM SIGPLAN International
 745 Workshop on Machine Learning and Programming Languages*, MAPL 2019, pp. 10–19, New
 746 York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450367196. doi:
 747 10.1145/3315508.3329973. URL <https://doi.org/10.1145/3315508.3329973>.

748 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez,
 749 Lukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2023. URL <https://arxiv.org/abs/1706.03762>.

752 Oreste Villa, Daniel Chavarría-Miranda, Vidhya Gurumoorthi, Andres Marquez, and Sriram Krishnamoorthy. Effects of floating-point non-associativity on numerical computations on
 753 massively multithreaded systems. Pacific Northwest National Laboratory (PNNL), Richland, WA
 754 (US), Cray User Group, Inc., Corvallis, OR, United States(US), 05 2009. URL <https://www.osti.gov/biblio/976992>.

756 Tongtong Wu, Linhao Luo, Yuan-Fang Li, Shirui Pan, Thuy-Trang Vu, and Gholamreza Haffari.
757 Continual learning for large language models: A survey, 2024. URL <https://arxiv.org/abs/2402.01364>.
759

760 Jinwei Yao, Kaiqi Chen, Kexun Zhang, Jiaxuan You, Binhang Yuan, Zeke Wang, and Tao Lin.
761 Deft: Decoding with flash tree-attention for efficient tree-structured llm inference, 2025. URL
762 <https://arxiv.org/abs/2404.00242>.
763

764 Jintao Zhang, Haofeng Huang, Pengle Zhang, Jia Wei, Jun Zhu, and Jianfei Chen. Sageattention2:
765 Efficient attention with thorough outlier smoothing and per-thread int4 quantization, 2025a. URL
766 <https://arxiv.org/abs/2411.10958>.
767

768 Jintao Zhang, Jia Wei, Haofeng Huang, Pengle Zhang, Jun Zhu, and Jianfei Chen. Sageattention:
769 Accurate 8-bit attention for plug-and-play inference acceleration, 2025b. URL <https://arxiv.org/abs/2410.02367>.
770

771 Jintao Zhang, Jia Wei, Pengle Zhang, Xiaoming Xu, Haofeng Huang, Haoxu Wang, Kai Jiang,
772 Jun Zhu, and Jianfei Chen. Sageattention3: Microscaling fp4 attention for inference and an
773 exploration of 8-bit training, 2025c. URL <https://arxiv.org/abs/2505.11594>.
774

775

776

777

778

779

780

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

810 A THE USE OF LARGE LANGUAGE MODELS
811

812 During the preparation of this manuscript, the authors employed a large language model (LLM) for
813 two primary purposes. First, the LLM was used as a tool to improve the grammar, spelling, and
814 overall clarity of the text. Second, it was used to assist in the initial stages of the literature search.
815 The role of the LLM was strictly that of an assistant. All language suggestions were reviewed and
816 edited by the authors to ensure they accurately reflected the intended scientific meaning. Further-
817 more, any literature identified with the assistance of the LLM was independently retrieved, reviewed,
818 and vetted for relevance and accuracy by the authors. All intellectual contributions, including the
819 conception of the research, methodology, and final conclusions, are the exclusive work of the human
820 authors, who take full responsibility for the final content of this paper.

821 B PROOF OF LEMMA 1
822

823 *Proof.* Let $LP_i(x)$ denote the length of the longest path from the source node s to node x in graph
824 G_i . The critical path length of G_i is $CP(G_i) = LP_i(t)$.

825 Due to the isomorphic structure of the chains in G_0 , all nodes at the same depth j have the same
826 longest path length from s . Let's denote this common length as $L_j = LP_0(v)$ for any node v with
827 $depth(v) = j$. Since all original edge weights in E_0 are strictly positive, it follows that for any two
828 depths j_1 and j_2 , if $j_1 < j_2$, then $L_{j_1} < L_{j_2}$. This implies $j_1 \leq j_2 \iff L_{j_1} \leq L_{j_2}$.

829 The proof proceeds by induction on the number of added edges, k .
830

831 **Base Case (k=1):** We prove the statement for the addition of a single edge $e_1 = (u, v)$ to G_0 to
832 form G_1 .

833 **Sufficient Condition (\implies):** Assume $depth(u) \leq depth(v)$. By the lemma's premise, we are
834 given that adding e_1 results in G_1 being a DAG. We must show that $CP(G_1) = CP(G_0)$.

835 The longest path to any node x in G_1 is given by the recurrence $LP_1(x) = \max_{(w,x) \in E_1} \{LP_1(w) +$
836 $weight(w, x)\}$. For node v , this becomes:

$$837 \quad LP_1(v) = \max(LP_0(v), LP_0(u) + 0)$$

838 By definition, $LP_0(v) = L_{depth(v)}$ and $LP_0(u) = L_{depth(u)}$. The condition $depth(u) \leq depth(v)$
839 implies $L_{depth(u)} \leq L_{depth(v)}$. Thus, $LP_1(v) = \max(L_{depth(v)}, L_{depth(u)}) = L_{depth(v)} =$
840 $LP_0(v)$. Since the longest path to v is unchanged, and this is the only modification, the longest
841 paths to all successors of v also remain unchanged. Therefore, $LP_1(x) = LP_0(x)$ for all $x \in V$,
842 which implies $CP(G_1) = CP(G_0)$.

843 **Necessary Condition (\iff):** Assume $CP(G_1) = CP(G_0)$ and (as per the lemma's premise) G_1 is
844 a DAG. We prove the contrapositive: if $depth(u) > depth(v)$, then $CP(G_1) > CP(G_0)$.

845 Since G_1 is a DAG, adding the edge (u, v) did not create a cycle. The longest path to v becomes:

$$846 \quad LP_1(v) = \max(LP_0(v), LP_0(u) + 0) = \max(L_{depth(v)}, L_{depth(u)})$$

847 Since we assume $depth(u) > depth(v)$ and all original edge weights are strictly positive, we have
848 $L_{depth(u)} > L_{depth(v)}$. This leads to $LP_1(v) = L_{depth(u)} > L_{depth(v)} = LP_0(v)$. The longest
849 path to v has strictly increased. This increase propagates to all successors of v , including the sink t .
850 Therefore, $LP_1(t) > LP_0(t)$, which means $CP(G_1) > CP(G_0)$. This contradicts our assumption.
851 Thus, the condition $depth(u) \leq depth(v)$ is necessary.

852 **Inductive Hypothesis (IH):** Assume for some $k \geq 1$, the lemma holds. That is, given that G_k
853 is a DAG, $CP(G_k) = CP(G_0)$ if and only if the condition $depth(u_i) \leq depth(v_i)$ held for all
854 $i \in \{1, \dots, k\}$. We make the stronger hypothesis that if the condition held, then $LP_k(x) = LP_0(x)$
855 for all nodes $x \in V$.
856

857 **Inductive Step:** We prove the lemma for the addition of the $(k + 1)$ -th edge, $e_{k+1} = (u, v)$, to G_k
858 to form G_{k+1} .

864 **Sufficient Condition (\implies):** Assume $\text{depth}(u) \leq \text{depth}(v)$. By the lemma's premise, we are
 865 given that G_{k+1} is a DAG. We must show $CP(G_{k+1}) = CP(G_k)$.
 866

867 The longest path to v in G_{k+1} is $LP_{k+1}(v) = \max(LP_k(v), LP_k(u) + 0)$. By the IH, since the
 868 conditions held for the first k edges, we have $LP_k(v) = LP_0(v) = L_{\text{depth}(v)}$ and $LP_k(u) = LP_0(u) =$
 869 $L_{\text{depth}(u)}$. The calculation is identical to the base case: $LP_{k+1}(v) = \max(L_{\text{depth}(v)}, L_{\text{depth}(u)}) =$
 870 $L_{\text{depth}(v)} = LP_k(v)$. The longest path to v is unchanged, and by propagation, $LP_{k+1}(x) = LP_k(x)$
 871 for all $x \in V$. This maintains our strong hypothesis and proves $CP(G_{k+1}) = CP(G_k) = CP(G_0)$.
 872

872 **Necessary Condition (\iff):** Assume $CP(G_{k+1}) = CP(G_k)$ and (as per the lemma's premise)
 873 G_{k+1} is a DAG. We prove the contrapositive: if $\text{depth}(u) > \text{depth}(v)$, then $CP(G_{k+1}) >$
 874 $CP(G_k)$.
 875

875 Since G_{k+1} is a DAG, adding (u, v) did not create a cycle. We compute $LP_{k+1}(v)$:
 876

$$LP_{k+1}(v) = \max(LP_k(v), LP_k(u) + 0)$$

878 Using the IH ($CP(G_k) = CP(G_0)$) implies the conditions held for the first k edges, so $LP_k(x) =$
 879 $LP_0(x)$ for all x :
 880

$$LP_{k+1}(v) = \max(LP_0(v), LP_0(u)) = \max(L_{\text{depth}(v)}, L_{\text{depth}(u)})$$

882 Since we assume $\text{depth}(u) > \text{depth}(v)$, we have $L_{\text{depth}(u)} > L_{\text{depth}(v)}$. This leads to $LP_{k+1}(v) =$
 883 $L_{\text{depth}(u)} > L_{\text{depth}(v)} = LP_0(v) = LP_k(v)$. The longest path to v strictly increases. This in-
 884 crease propagates to the sink node t , so $CP(G_{k+1}) > CP(G_k)$. This contradicts our assumption.
 885 Therefore, the condition is necessary.
 886

887 By the principle of induction, the lemma holds for any $k \geq 1$. \square
 888

889 C EXACT ALGORITHM AND MODIFICATIONS

891 We present the exact algorithm in Algorithm 1 in this section. The following pseudocode is adapted
 892 from the original FlashAttention-3 paper (Shah et al., 2024), with our unique modifications clearly
 893 highlighted.
 894
 895
 896
 897
 898
 899
 900
 901
 902
 903
 904
 905
 906
 907
 908
 909
 910
 911
 912
 913
 914
 915
 916
 917

918
919
920
921

Algorithm 1 DASH algorithm

923 **Require:** Matrices $\mathbf{Q}, \mathbf{K}, \mathbf{V}, \mathbf{O}, \mathbf{dO} \in \mathbb{R}^{N \times d}$ in HBM, logsumexp vector $L \in \mathbb{R}^N$ in HBM, block
 924 sizes B_c, B_r .

925 1: In a preprocessing kernel, compute $D = \text{rowsum}(\mathbf{dO} \circ \mathbf{O}) \in \mathbb{R}^d$ (pointwise multiply), write
 926 D to HBM and divide it into T_r blocks D_1, \dots, D_{T_r} of size B_r each.

927 2: Divide \mathbf{Q} into $T_r = \left\lceil \frac{N}{B_r} \right\rceil$ blocks $\mathbf{Q}_1, \dots, \mathbf{Q}_{T_r}$ of size $B_r \times d$ each, and divide \mathbf{K}, \mathbf{V} in to
 928 $T_c = \left\lceil \frac{N}{B_c} \right\rceil$ blocks $\mathbf{K}_1, \dots, \mathbf{K}_{T_c}$ and $\mathbf{V}_1, \dots, \mathbf{V}_{T_c}$, of size $B_c \times d$ each.

929 3: Divide \mathbf{dO} into T_r blocks $\mathbf{dO}_i, \dots, \mathbf{dO}_{T_r}$ of size $B_r \times d$ each, and divide L into T_r blocks
 930 L_i, \dots, L_{T_r} of size B_r each.

931 4: Initialize pipeline object to manage barrier synchronization with s -stage circular SMEM buffer.

932

933 5: **if** in producer warpgroup **then**

934 6: Deallocate predetermined number of registers.

935 7: Issue load \mathbf{K}_j and \mathbf{V}_j from HBM to shared memory.

936 8: Upon completion, commit to notify consumer of the load of \mathbf{K}_j and \mathbf{V}_j .

937 9: **for** i in assigned Q-tile schedule **do**

938 10: Wait for the $(i \% s)$ th stage of the buffer to be consumed.

939 11: Issue loads of $\mathbf{Q}_i, \mathbf{dO}_i$ from HBM to shared memory at the $(i \% s)$ th stage of the buffer.

940 12: Upon completion, commit to notify consumers of the loads of $\mathbf{Q}_i, \mathbf{dO}_i$.

941 13: **end for**

942 14: **else if** in consumer warpgroups **then**

943 15: Reallocate predetermined number of registers as function of number of consumer warps.

944 16: On-chip, Initialize $\mathbf{dK}_j = (0)_{B_c \times d}, \mathbf{dV}_j = (0)_{B_c \times d}$.

945 17: Wait for \mathbf{K}_j and \mathbf{V}_j to be loaded in shared memory.

946 18: **for** i in assigned Q-tile schedule **do**

947 19: Wait for \mathbf{Q}_i to be loaded in shared memory.

948 20: Load L_i, D_i from HBM to on-chip SRAM.

949 21: On chip, compute $\mathbf{S}_i^{(j)} = \mathbf{Q}_i \mathbf{K}_j^T \in \mathbb{R}^{B_r \times B_c}$ (SS-GEMM). Commit.

950 22: Wait for \mathbf{dO}_i to be loaded in shared memory.

951 23: On chip, compute $\mathbf{dP}_i^{(j)} = \mathbf{dO}_i \mathbf{V}_j^\top \in \mathbb{R}^{B_r \times B_c}$ (SS-GEMM). Commit.

952 24: On chip, wait for $\mathbf{S}_i^{(j)}$, then compute $\mathbf{P}_i^{(j)} = \exp(\mathbf{S}_i^{(j)} - L_i) \in \mathbb{R}^{B_r \times B_c}$.

953 25: On chip, wait for $\mathbf{dP}_i^{(j)}$, then compute $\mathbf{dS}_i^{(j)} = \mathbf{P}_i^{(j)} \circ (\mathbf{dP}_i^{(j)} - D_i) \in \mathbb{R}^{B_r \times B_c}$.

954 26: On chip, compute $\mathbf{dV}_j \leftarrow \mathbf{dV}_j + (\mathbf{P}_i^{(j)})^\top \mathbf{dO}_i \in \mathbb{R}^{B_c \times d}$ (RS-GEMM). Commit.

955 27: On chip, compute $\mathbf{dK}_j \leftarrow \mathbf{dK}_j + \mathbf{dS}_i^{(j)^\top} \mathbf{Q}_i \in \mathbb{R}^{B_c \times d}$ (RS-GEMM). Commit and wait
 956 for both \mathbf{dV}_j and \mathbf{dK}_j .

957 28: On chip, compute $\mathbf{dQ}_i^{(\text{local})} = \mathbf{dS}_i^{(j)} \mathbf{K}_j \in \mathbb{R}^{B_r \times d}$ (SS-GEMM), and write $\mathbf{dQ}_i^{(\text{local})}$ to
 958 smem. Notify the \mathbf{dQ} -writer.

959 29: **end for**

960 30: **else if** in \mathbf{dQ} -writer warp **then**

961 31: **for** i in assigned Q-tile schedule **do**

962 32: Wait for $\mathbf{dQ}_i^{(\text{local})}$ to be ready in smem.

963 33: **Wait until the global order grants this block its turn to reduce.**

964 34: Using a semaphore, atomically add $\mathbf{dQ}_i^{(\text{local})}$ to \mathbf{dQ}_i in global memory.

965 35: **Advance the global order.**

966 36: **end for**

967 37: **end if**

968

969

970

971
