Lean Attention: Hardware-Aware Scalable Attention Mechanism for the Decode-Phase of Transformers

Rya Sanovar   Srikant Bharadwaj   Renee St. Amant   Victor Rühle   Saravan Rajmohan

Microsoft

Abstract—Transformer-based models have emerged as one of the most widely used architectures for natural language processing, natural language generation, and image generation. The size of the state-of-the-art models has increased steadily reaching billions of parameters. These huge models are memory hungry and incur significant inference latency even on cutting edge AI-accelerators, such as GPUs. Specifically, the time and memory complexity of the attention operation is quadratic in terms of the total context length, i.e., prompt and output tokens. Thus, several optimizations such as key-value tensor caching and FlashAttention computation have been proposed to deliver the low latency demands of applications relying on such large models. However, these techniques do not cater to the computationally distinct nature of different phases during inference.

To that end, we propose LeanAttention, a scalable technique of computing self-attention for the token-generation phase (decode-phase) of decoder-only transformer models. LeanAttention enables scaling the attention mechanism implementation for the challenging case of long context lengths by re-designing the execution flow for the decode-phase. We identify that the associative property of online softmax can be treated as a reduction operation thus allowing us to parallelize the attention computation over these large context lengths. We extend the "stream-K" style reduction of tiled calculation to self-attention to enable parallel computation resulting in an average of 2.6x attention execution speedup over FlashAttention-2 and up to 8.33x speedup for 512k context lengths.

I. INTRODUCTION

Transformer-based [31] language models [12], [23], [26], [30], [37] have revolutionized the field of natural language processing (NLP) and found applications across diverse domains [18], [29]. These powerful models, fueled by massive amounts of data and sophisticated architectures, have become indispensable tools for tasks such as machine translation [15], question answering [7], text generation [7], and sentiment analysis.

The core of the transformer architecture is the powerful component of self-attention; however, execution of the self-attention mechanism is slow and suffers from a large memory footprint, especially when dealing with long context length. A standard implementation of self-attention demonstrates quadratic time and memory complexity with respect to total sequence length, which leads to scalability challenges as model sizes [10] and supported context lengths increase [8], [24], [36]. Despite these scalability challenges, we see a trend toward state-of-the-art models supporting greater and greater context lengths, with some production models supporting context lengths of

hundreds of thousands of tokens. Support for long context lengths can improve a model's utility by allowing for an increasingly rich context, which is particularly beneficial in a range of applications (e.g. RAG involving numerous or long documents) allowing improved relevance, coherence, and user experience.

To mitigate LLM scalability challenges, mechanisms like FlashAttention [14] and FlashAttention-2 [13] have been developed. FlashAttention brings IO-awareness to optimize computation in the attention mechanism in a way that reduces slow reads and writes to and from GPU high bandwidth memory [19] (via incrementally computing the softmax computation in SRAM, also known as tiling). It allows for parallelization over batch size and number of heads. FlashAttention-2 builds on FlashAttention to further optimize the attention mechanism by additionally reducing the number of non-matrix multiply operations to maximize GPU throughput, and it additionally enables parallelization across input sequence length (query length) as well. While these optimizations provide significant improvements, e.g. FlashAttention-2 realized 2x speedup over FlashAttention, these mechanisms only provide performance benefits for a subset of problem sizes (i.e. sequence length, batch size, and number of heads) because they overlook the distinct behavior of the attention mechanism during the decode phase versus the prefill-phase in decoder-only transformer models.

In decoder-only transformer models, the inference process for a single request involves multiple forward passes of the model where output tokens are generated sequentially [21]. This inference procedure inherently comprises two distinct computational phases due to the practice of reusing (caching) the key-value tensors of the attention mechanism of the previously computed tokens [28]. The first phase is the *prompt computation phase* (sometimes known as *prefill phase*) where all tokens from the input prompt undergo parallel forward passes through the model to generate the first output token. This phase is computationally intensive and demands high FLOPS/s (floating point operations per second) [21]. Following the prompt computation, the *decode phase* (sometimes known as *token-generation phase*) begins in an auto-regressive manner [31]. Each subsequent token is produced based on the forward pass of the preceding token and the cached context (*kv-cache*) from previous tokens in the sequence. With the push towards longer context lengths, this cached context can be long, exceeding

FlashAttention-2
SM Occupancy: 2/5

FlashAttention-2 w/ FixedSplit
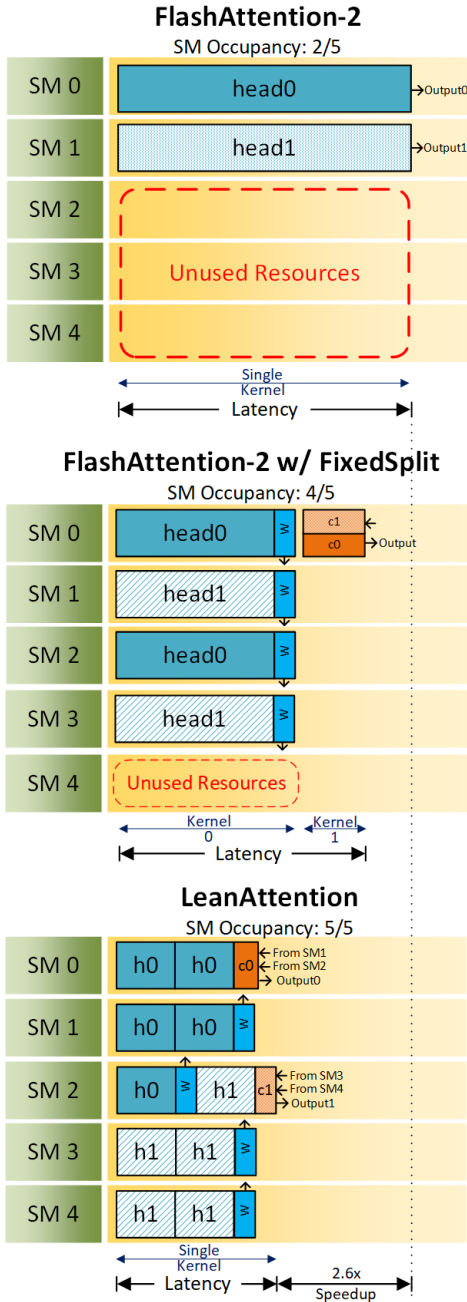SM Occupancy: 4/5

LeanAttention
SM Occupancy: 5/5

Fig. 1. Execution schedule of FlashAttention-2, FlashDecoding, and LeanAttention across a hypothetical five SM GPU executing attention of 2 heads. LeanAttention splits the context into optimal LeanTiles (shown here with 5 tiles per head).

work as it relates to the decode-phase of inference, which we find exhibits unique computation characteristics in comparison to the prefill phase. In summary, for the decode phase of LLM inference, FlashAttention-2 parallelizes computation along the number of heads in the mode and batch size. Adding fixed-split partitioning in FlashDecoding [4] additionally enables parallelization along context length, but this technique suffers from inefficient load balancing and unnecessary overheads.

To address these limitations, we introduce *LeanAttention*, a generalized *exact-attention* mechanism which enables parallelization across all modes of the attention matrix, ensures perfect quantization efficiency (i.e. GPU utilization), delivers a runtime speedup in attention computation for long context lengths, and is scalable to multi-GPU scenarios with tensor parallelism. As shown in Figure 1, we aim to execute attention in a single fused kernel with a greater degree of parallelization, i.e. across the context length as well, while ensuring peak processor occupancy and minimal reduction overheads that are independent of problem size (prefill and decode stage).

Overall, our contributions are as follows:

- Identify the limitations of state-of-the-art attention execution optimizations on GPUs during the decode-phase of transformer-based models. (subsection III-B)
- Approach the softmax operation in attention as a special reduction operation and extract it out of the inner loop of the attention algorithm, thus treating re-scaling of unscaled attention output tensors as an associative reduction operation. (subsection IV-A)
- Leverage a stream-K style [27] partitioning of attention that always provides equalized compute loads to every compute unit in the hardware system to maximize hardware occupancy and speedup across a wide range of problem sizes and hardware architectures. (subsection IV-C)
- Finally, define LeanAttention as a generalized attention partitioning mechanism that closely mirrors attention computation to the compute and memory resources of modern hardware systems and that is optimized for decode-phase problem sizes in addition to prefill-phase. (section IV)

As detailed in section VI, LeanAttention results in an average of 2.6x latency speedup over FlashAttention-2 for the decode phase of transformer-based models and up to 8.33x speedup for 512k context lengths, with near 100% GPU occupancy across the entire problem landscape.

II. BACKGROUND

In this section, we provide required background on Standard Attention [31] and FlashAttention-2 [13].

A. Standard Attention

We will first explain some elementary details of the standard attention implementation.

For a given input tensor with dimensions of batch size B, query sequence length N_q, key/value sequence length (also known as context length) N_k, and hidden dimension D, multi-head attention typically splits attention computation into h

more than hundreds of thousands of tokens in length [16], [22], [24], [36]. Despite state-of-the-art batching techniques [34] and attention partitioning mechanisms [13], [14], the sequential processing of this long context length makes the decode phase slow, bound by memory bandwidth [32] and capacity [21]. Importantly, as we discuss in section III, even when the prompt size is significantly larger than the number of output tokens, the majority of the overall processing time is consumed by the decode (token generation) phase.

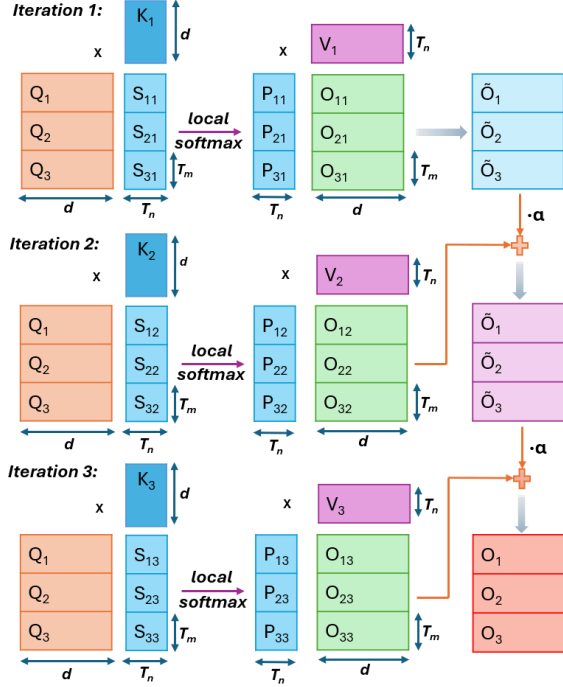In this work, we aim to address the limitations of previous

Fig. 2. Iterative update of output in FlashAttention-2 with $C_n = 3$ iterations.

number of heads along the hidden dimension, with each head responsible for computing attention independently for a head dimension of size $d = D/h$.

Unlike standard transformer execution, the query length and context lengths may not always be equal in size, where key-value tensors are cached [28]. For instance, the prefill-phase of generative decoder-only transformers such as GPT-4 [26] or Phi-2 [23] has sequence lengths $N_q = N_k = N$, but in their decode phase the context length increments by 1 after every autoregressive step of decode generation, while the query sequence length, for a given batch instance and head, is the singular token that was generated in the previous n-th time step, i.e., $N_q = 1$ and $N_k = N + n$.

The query matrix $Q \in R^{N_q \times d}$ and key K and value V matrices $\in R^{N_k \times d}$ are inputs to the following equation which is computed independently by the different batch instances and heads. The output matrix $O \in R^{N_q \times d}$ is obtained in essentially three steps as shown in Equation 1. Table I summarises the three operations involved in self-attention along with their corresponding dimensions involved in both decode and prefill-phase.

$$S = QK^T, \ P = softmax(\frac{S}{\sqrt{d}}), \ O = PV \quad (1)$$

Standard attention implementation involves computing the large intermediate matrices, namely the attention score matrix $S \in R^{N_q \times N_k}$ and the softmax matrix $P \in R^{N_q \times N_k}$ and storing them in global memory. It requires *a priori* knowledge of all tokens in a row of the attention matrix for computing

the softmax function. Specifically, the row-wise maximum and exponential sum must be computed beforehand, which necessitates examining all tokens in a given row of the attention matrix and thus necessitates storing intermediate results. The computational complexity of standard attention is on the order of $O(N_q N_k d)$, with the two matrix multiplications contributing to the majority of it. Due to slow memory access speeds, storing and retrieving these matrices [19] is costly in terms of latency and incurs a large memory footprint, both on the order of $O(N_q N_k)$.

Operation	Type	Operation Dimension	
		Prefill	Decode
$query \times key$	MatMul	$L \times d \times L$	$1 \times d \times L$
$softmax$	EleWise	$L \times L$	$1 \times L$
$attn\_score \times value$	MatMul	$L \times L \times d$	$1 \times L \times d$

TABLE I
OPERATIONS IN SELF-ATTENTION. MATRIX MULTIPLICATIONS ARE DESCRIBED IN THE $M \times N \times K$ FORMAT.

B. Flash Attention-2

To mitigate the memory footprint and access overhead [19] associated with storing the S and P matrices, FlashAttention introduced an adroit way of fusing all the three operations: $query \times key$ MatMul, softmax and $attn\_score \times value$ MatMul into a single kernel, requiring no intermediate global memory reads and writes. To this end, it employs two strategies: tiling and recomputation.

By utilizing the online softmax algorithm [25], FlashAttention only needs a single pass over an entire row of tokens to compute their softmax, bypassing the issue of *a priori* knowledge in standard attention computation. This helps leverage the tiling strategy which partitions the input matrices into smaller chunks that are easier to bring into shared memory (Algorithm1 Line 4-5). Line 6 in Algorithm 1 refers to a grid launch of C_m *cooperative thread arrays* (CTAs)[1], each computing a given block tile of rows of the output matrix O.

In summary, the three operations discussed in Equation 1 are fused together and computed locally for a given chunk. To ensure that the correct attention output is obtained in the end, the output block generated from each chunk is appropriately scaled by the right scaling factors on the fly, before moving on to the next chunk to be computed for updating a given output tile. This on-chip fused computation avoids the need of storing the intermediate attention matrix in global memory. This process is visualized in Figure 2 that depicts a case where $C_n = C_m = 3$, number of streaming multiprocessors (SMs) on the GPU is assumed to be 5, and α is a normalization parameter that is element-wise multiplied to the incoming partial output matrix. In addition to tiling, FlashAttention-2 utilizes a recomputation strategy to reduce memory reads and writes while calculating gradients in the backward pass. We omit that discussion here given that our focus is on inference optimization as opposed to training.

[1]Blocks of GPU threads are cocheduled in CTAs, which virtualize the hardware's streaming multiprocessor cores (SMs)

Algorithm 1 FlashAttention-2 [13]

1: **Require:** Load matrices $Q \in R^{N_q \times d}$ and $K, V \in R^{N_k \times d}$ into GMEM.
2: **Require:** Initialize Matrix O to $(0)_{N_q \times d} \in R^{N_q \times d}$ in GMEM.
3: Set block sizes T_m and T_n.
4: Partition Q, O as $Q_i, O_i \in R^{T_m \times d}$ where $i \in (1, C_m)$.
5: Partition K, V as $K_j, V_j \in R^{T_n \times d}$ where $j \in (1, C_n)$.
6: **for** $i = 1$ **to** C_m **do**
7:    Load Q_i from GMEM to SMEM.
8:    Initialize m_i to $(-\infty)_{T_m \times 1}$ and l_i to $(0)_{T_m \times 1} \in R^{T_m \times 1}$ in SMEM.
9:    **for** $j = 1$ **to** C_n **do**
10:      Load K_j, V_j from GMEM to SMEM.
11:      Compute on-chip:
12:      $S_i = Q_i K_j^T$ where $S_i \in R^{T_m \times T_n}$
13:      $m_i^{new} = max(m_i, rowmax(S_i))$
14:      $P_i = exp(S_i - m_i^{new})$ where $P_i \in R^{T_m \times T_n}$
15:      $l_i^{new} = e^{m_i - m_i^{new}} l_i + rowsum(P_i)$
16:      $O_i^{new} = P_i V_j + diag(e^{m_i - m_i^{new}}) O_i$
17:      $l_i = l_i^{new}, m_i = m_i^{new}, O_i = O_i^{new}$
18:    **end for**
19:    Compute $O_i = diag(l_i)^{-1} O_i$ and write to GMEM.
20:    Compute logexpsum $L_i = m_i + log(l_i)$ and write to GMEM.
21: **end for**

Fig. 3. Timeshare of prompt (prefill) stage compared to decode stage for different output token ratios with 1024 prompt tokens.

Along with parallelizing computation over batches and heads, FlashAttention-2 further parallelizes over the independent query blocks and this gives a speedup of 2x times that of vanilla FlashAttention. To save more global memory space, FlashAttention-2 also provides the minor optimization of storing a logarithmic exponential sum $L \in R^{N_q \times 1}$ rather than storing both l and m (Algorithm 1 Line 20). It additionally uses the algorithmic hack of not scaling a given output block O_i by the final correct exponential sum l_i until the very end to reduce non-matmul FLOPs, which takes more cycles to compute on the GPU compared to matmul FLOPs (Algorithm 1 Line 19).

Thus, the above two strategies and optimized work partitioning ensures that the extra global memory space required by FlashAttention-2 is $O(N_q)$ (needed to store the logexpsum L), an impressive improvement in memory footprint over the $O(N_q \times N_k)$ from traditional attention and the addtional partitioning helps it reach 50-70% of peak theoretical FLOPS/s.

While FlashAttention-2, and other related techniques, such as Ring Attention [24] and Striped Attention [9] are optimized for prefill-phase problem sizes, they suffers from longer latencies during the decode phase.

III. Challenges In The Decode Phase

Prior to outlining our methodology for LeanAttention, to set the stage for our approach, we delve into some of the challenges encountered in the decode phase of LLM inference, as well as the limitations of FlashAttention-2 optimizations in the decode phase.
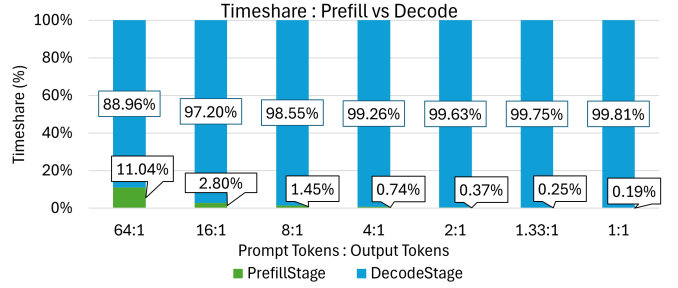
A. Time Spent in Decode Phase

Generative LLM inference comprises two distinct computational phases: the prompt-phase (sometimes called the prefill phase) and the decode phase (sometimes called the token-generation phase). In the prefill-phase, all tokens in the input prompt (aka query) undergo parallel forward passes through the model to generate the first output token. The query length, N_q, in this phase is the same as the context length, N_k, resulting in an NxN attention matrix ($N_q = N_k = N$). This phase is computationally intensive and demands high FLOPS/s.

Following the prefill phase, the decode phase begins generating each subsequent output token. The nature of this phase is auto-regressive, where the next output token is produced based on the forward pass of the preceding token and the cached context (KV cache) from previous tokens in the sequence. For each iteration of the decode phase, its query length is a single token, $N_q = 1$, and its context length, N_k, could be on the order of more than thousands of tokens depending on the auto-regressive step and input prompt length. This makes leveraging parallelism along context length (N_k) a crucial aspect in reducing decode phase processing time.

As the number of output tokens generated rises, the context length becomes longer and thus the proportion of time spent in the decode phase relative to the prefill-phase becomes larger. Figure 3 depicts this imbalance in processing time spent between the prefill-phase and the decode phase. Even with a prompt-input to output-token ratio of 64:1, more than 80% of the processing time is consumed by the decode phase, taking up to nearly 100% of the timeshare for longer output lengths.

B. Limitations of FlashAttention-2 for Decode

In both the prompt and decode phase, FlashAttention-2 computes *sequentially* along the context length (N_k) dimension, following the softmax introduced dependency among the tokens. While FlashAttention-2 does parallelize over query lengths (N_q) to increase SM occupancy, this additional mode of parallelism has limited parallelization capacity in the decode phase because query length is equal to a single token ($N_q = 1$). Adhering to this sequential nature and not partitioning attention along context length makes vanilla FlashAttention-2 [13] suffer from extremely low SM occupancy during decode as depicted in Figure 4. FlashAttention-2 is constrained by the sequential
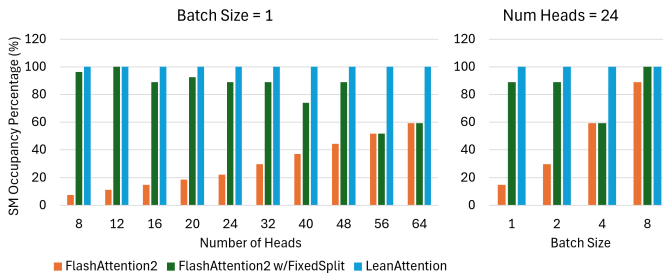
Fig. 4. Occupancy of SMs in a single Nvidia-A100-80GB GPU with different attention modes. FlashAttention2 has a low occupancy whereas FlashDecoding (FlashAttention2 w/ FixedSplit) has quantization efficiency issue with the 108 SMs on the GPU. Lean Attention occupies all the SMs available in the system

nature of rescaling and updating softmax as it traverses each key/value tile for a given query tile. This means that at any given point in time, the number of CTAs in flight on the GPU is directly proportional to the number of query tiles, and, therefore, to the query sequence length - regardless of the context length.

More explicitly, for a single batch instance, the maximum number of heads for state of the art LLMs barely occupy the compute resources of modern hardware architecture systems during the decode-phase where query length $N_q = 1$. For example, for a model with 128 heads, its decode phase would suffer from severe under-utilization of an 8 GPU A100 system that has 864 compute cores at its disposal. Unlike, the prefill-phase, decode-phase can offer parallelization only across batch size and number of heads for FlashAttention-2.

Processor occupancy in vanilla FlashAttention-2 could be improved by increasing the batch size or number of heads, the other two modes of parallelization addressed by FlashAttention-2. Intuitively, having larger batch sizes in the decode phase could provide enough work to every compute resource to fully occupy the GPU, but this brings other challenges and limitations. Due to increasingly large model sizes, the need to independently caching KV context for every batch instance would exceed the memory capacity of the hardware system. Moreover, scheduling overheads [35] for efficiently batching queries along with the challenges for batching low SLA queries would increase inference latency and challenge utilization.

Without having to resort to larger batch sizes as the sole solution to resolve the GPU occupancy issue (which is limited by available memory capacity), the large context length in the decode phase would benefit from partitioning its workload across different SMs efficiently. This motivates the need for smarter attention decomposition technique which can efficiently distribute the workload across the cores without the need to resort to larger batch sizes.

C. FlashDecoding

FlashDecoding, which is FlashAttention-2 with *fixed-split* partitioning, has recently been proposed [4], [5], [17], where attention computation is also partitioned along context length (N_k). *Fixed-split* is a general matrix multiplication decomposition scheme that we describe briefly here. Given a MatMul

computation problem with matrices A ($M \times K$) and B ($N \times K$) to obtain a matrix C ($M \times N$) where $C = AB^T$, to optimize concurrent computation, the fixed-split mechanism [6] partitions the K-mode of the A and B matrices into s batches based on a fixed splitting factor s provided dynamically at run time. This launches s times the CTAs (Cooperative Thread Arrays, equivalent to a threadblock) as launched without fixed-split, which are computing partial products of the C matrix concurrently. Fixed-split utilizes the associativity of addition in the inner product to later reduce or "fix-up" the partially computed C matrices to result in the final C matrix. The concurrency from fixed-split reduces latency and simultaneously increases hardware occupancy at the cost of an additional reduction at the end. FlashDecoding++ [17] builds on this technique and gets rid of synchronization costs by approximating a global max value in softmax to avoid final re-scaling. The intermediate local softmax calculation in the inner loop of FlashDecoding is avoided, and a final global softmax is computed once all the partial exponential sums have been determined. Additionally double buffering is used to hide memory access latencies.

Despite these improvements, fixed-split used in FlashDecoding and FlashDecoding++ is a non-optimal load balancing strategy. It requires launching an additional reduction kernel and thus suffers with kernel launch overheads as well as reduction or fix-up overheads which scale with problem size. Moreover, a fixed-split decomposition suffers from quantization inefficiencies, i.e. low GPU utilization, and loses out on performance gains it could get from the idle resources otherwise. These inefficiencies have been depicted in Figure 1.

The fixed-split style of partitioning of FlashDecoding along the context length does occupy a larger number of compute resources on the GPU compared to vanilla FlashAttention-2, but this occupancy varies greatly depending on problem size, i.e. number of heads, batch size, and context length. As shown in Figure 4, quantization efficiency, i.e GPU utilization, is heavily dependent on the problem size, splitting factor and number of compute units in the hardware system, making it unlikely for FlashDecoding to reach perfect quantization efficiency for all problem sizes and hardware systems.

In contrast, LeanAttention, with its stream-K-style decomposition discussed in section IV, will always provide well-balanced loads to each compute unit in the hardware system and reach near 100% GPU occupancy across all problem sizes and hardware architectures.

D. Multi-GPU Execution with Tensor Parallelism

Vanilla FlashAttention-2 not only severely underutilizes GPU cores in the decode phase, but is also not adaptable to multi-GPU scenarios due to its lack of support for tensor parallelism. This makes FlashAttention-2 less scalable to multi-GPU systems which has become an imperative due to capacity-boundedness of contemporary large language models [10] and the support they require for increasingly long context lengths [8]. This asserts the need for an attention mechanism that also scales well to multi-GPU scenarios.
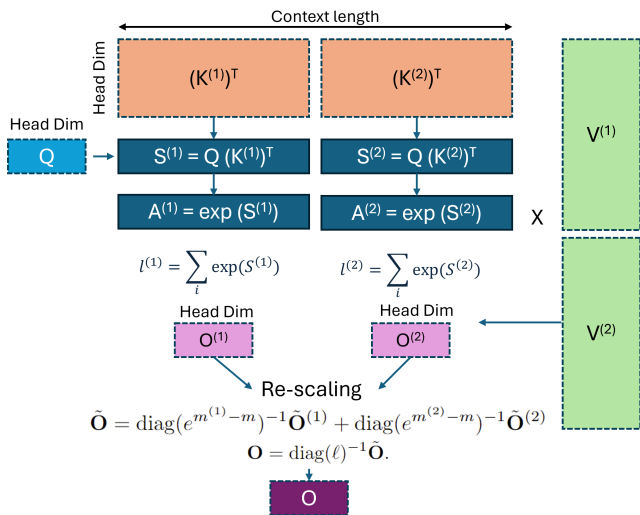
Fig. 5. Illustrative diagram showing LeanAttention's decomposition strategy with two differently sized work volumes assigned to different CTAs. The un-scaled outputs are independently computed and the re-scaled later as a reduction operation. Note that this can be generalized to any arbitrary sized work volume.

FlashDecoding utilizes fixed-split decomposition that implements this partitioning to enable parallelism along context length, yet suffers from quantization inefficiency and fix-up overheads that scale with problem size.

These challenges motivate the need for a generalized attention mechanism that works for a vast set of problem sizes (in both prefill and decode phase) and is closely aligned with the memory and compute hierarchies of modern hardware systems. We formulate this generalized attention mechanism as **LeanAttention**, which computes exact attention faster in a single fused kernel launch, has optimal quantization efficiency for all kinds of problem sizes, whilst also being scalable to multi-GPU scenarios through its support for tensor parallelism.

IV. LEANATTENTION

LeanAttention is an optimized scalable execution mechanism for computing self-attention. It provides extensive parallelism across all modes of the attention tensor, with well-balanced computation workloads to each CTA ensuring close to 100% SM occupancy and delivering a runtime speedup in attention execution.

To exploit regular matrix multiplication optimizations on the GPU for attention, namely, the state-of-the-art stream-K style of executing MatMuls, we restructure the attention operation to closely mirror the execution paradigms of a matrix multiplication partitioning.

The attention operation consists of two MatMuls joined by a softmax operation. For a stream-k style partitioning of these attention MatMuls on the GPU, we define the softmax re-scaling operation in attention as a final reduction operation to replace the addition operation that behaves as reduction in a basic MatMul.

Firstly, we identify the smallest optimal lean block of attention computation, termed as *LeanTile* (subsection IV-B), which can be mapped on the hardware resources in a flexible style

akin to "stream-k" decomposition [27] of matrix multiplications (subsection IV-C). Multiple such LeanTile's belonging to either one or many attention outputs will constitute a workload assigned to a CTA. By the nature of Stream-K's equalized load balancing strategy, each CTA will compute equal number of LeanTile's.

We identify that the associative property of softmax re-scaling enables us to treat it as a reduction operation along the context-length dimension of the attention operation (described in subsection IV-A). This softmax re-scaling reduction will help us consolidate partial attention outputs computed from varying number of LeanTile iterations on different CTAs.

In the following subsections, we first outline the identification of softmax re-scaling as a reduction operation, followed by a conceptualization of a *LeanTile* as a unit granularity in a CTA block and the stream-K style mapping within these CTAs, followed by an explanation of the overall execution flow of LeanAttention. Figure 6 shows an illustrative running example of the execution flow of LeanAttention.

A. Softmax Re-scaling as Reduction

FlashAttention-2 uses the online softmax technique [25] to split the attention computation for a single query block into chunks of work. Each chunk of work comprises a key block and a corresponding value block, and these chunks arrive in a sequential manner to update the attention output for the given query block. FlashAttention-2 computes online softmax for every incoming chunk, rescales the intermediate output obtained from the previous chunk and combines it with the partial output from the current chunk to get the latest updated output. However, this method of computing exact attention is constrained in its sequentiality, leading to slower computation especially in the decode phase where there a lot of key/value chunks to parse through for a given query block.

In Lean Attention, we propose computation of partial attention outputs resulting from these chunks of work concurrently on different hardware units, and ensuring that we have a well-balanced distribution of work across all hardware units through a Stream-K style decomposition of attention (discussed later in subsection IV-C). This decomposition results in splits of work for a given SM that are not always equal in size, i.e., the key/value tensors are not split in equal chunks of work (unlike FlashDecoding [4]). For example, in Figure 1, for computing attention for a query block h_0, SM0 and SM1 receive same sized chunks of work, but SM2 receives half the amount of work for h_0 that SM0 or SM1 received.

To reduce these partial attention outputs that result from differently sized chunks, we use a softmax re-scaling operation. This requires us to identify the softmax's associativity property that allows it to reduce chunks of different sizes: application of softmax re-scaling as a reduction operator will give the same attention output, *regardless of the way the work might be split, whether in equal-sized chunks or arbitrary differently sized chunks*.

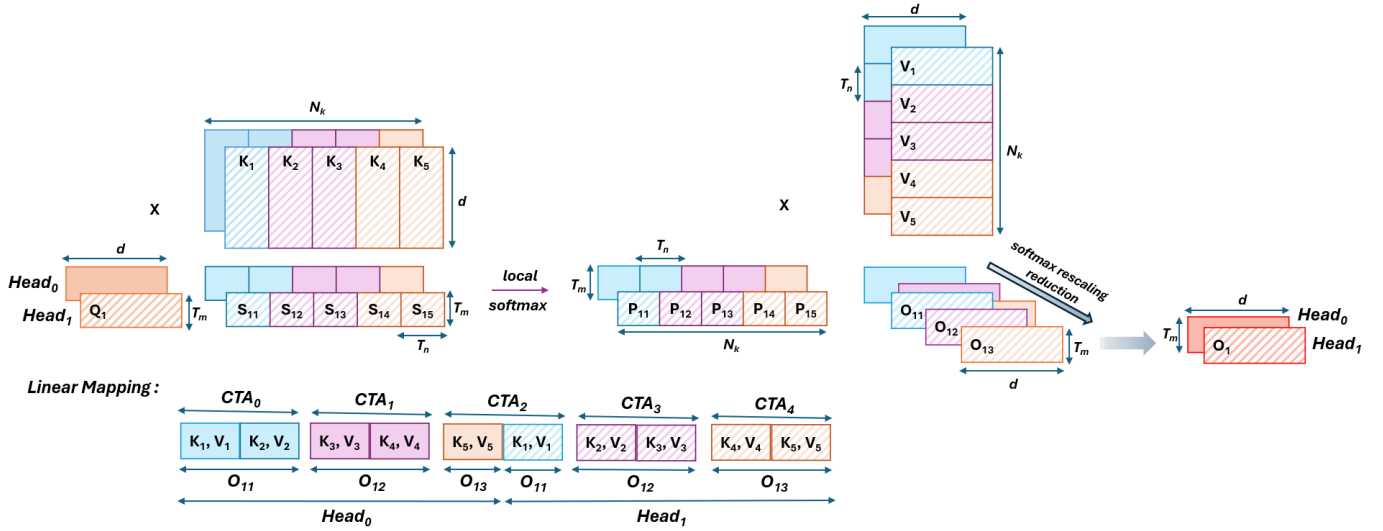Without loss of generality, we describe this process of reduction to obtain one row vector of the attention matrix

6

Fig. 6. Lean Attention for 2 heads in Decode Phase. Each head's output tile is decomposed into 5 Lean Tiles.

$\mathbf{S}$, of the form $\begin{bmatrix} \mathbf{S}^{(x)} & \mathbf{S}^{(y)} \end{bmatrix}$ consisting of some unequal length vectors $\mathbf{S}^{(x)}, \mathbf{S}^{(y)}$ where $\mathbf{S}^{(x)} \in \mathbb{R}^{1 \times B_c^{(x)}}$ and $\mathbf{S}^{(y)} \in \mathbb{R}^{1 \times B_c^{(y)}}$, where 1 is the query length and $B_c^{(x)}$ and $B_c^{(y)}$ are the unequal key/value lengths. The vectors $\mathbf{S}^{(x)}$ and $\mathbf{S}^{(y)}$ were computed from $\mathbf{Q} \times (\mathbf{K}^{(x)})^T$ and $\mathbf{Q} \times (\mathbf{K}^{(y)})^T$ as shown in Figure 5. Note that, to generalize this procedure for chunks of any size, the context length of $\mathbf{K}^{(x)}$ and $\mathbf{K}^{(y)}$ are $B_c^{(x)}$ and $B_c^{(y)}$ and are not necessarily equal.

The attention computation is split into two parts. The first part involves calculation of an "un-scaled" version of $\mathbf{O}^{(i)}$ (where i is either x or y) along with statistics $m^{(i)}$ and $\ell^{(i)}$:

$$\mathbf{S}^{(i)} = \mathbf{Q}(\mathbf{K}^{(i)})^T \in \mathbb{R}^{1 \times B_c^{(i)}}$$
$$m^{(i)} = \mathrm{rowmax}(\mathbf{S}^{(i)}) \in \mathbb{R}^{1x1}$$
$$\ell^{(i)} = \mathrm{rowsum}(e^{\mathbf{S}^{(i)} - m^{(i)}}) \in \mathbb{R}^{1x1}$$
$$A^{(i)} = exp(\mathbf{S}^{(i)} - m^{(i)}) \in \mathbb{R}^{1 \times B_c^{(i)}}$$
$$\tilde{\mathbf{O}}^{(i)} = A^{(i)}\mathbf{V}^{(i)} \in \mathbb{R}^{1 \times d}$$

Softmax Re-scaling Operation. The second part involves re-scaling the "un-scaled" outputs $\mathbf{O}^{(i)}$ using the statistics $\ell^{(i)}$. We define the softmax re-scaling operation $f(x, y)$ for two intermediate outputs $\mathbf{O}^{(x)}$ and $\mathbf{O}^{(y)}$ as follows:

$$m^{(x,y)} = \max(m^{(x)}, m^{(y)})$$
$$\ell^{(x,y)} = e^{m^{(x)} - m^{(x,y)}}\ell^{(x)} + e^{m^{(y)} - m^{(x,y)}}\ell^{(y)}$$
$$f(x,y) = \mathrm{diag}(e^{m^{(x)} - m^{(x,y)}})\tilde{\mathbf{O}}^{(x)} + \mathrm{diag}(e^{m^{(y)} - m^{(x,y)}})\tilde{\mathbf{O}}^{(y)}$$
$$f(x,y) = \tilde{\mathbf{O}}^{(x,y)}$$
$$\mathbf{O}^{(x,y)} = \mathrm{diag}(\ell^{(x,y)})^{-1}f(x,y)$$

Proof of Associativity The associative nature of softmax re-scaling $f(x, y)$ allows us to reduce intermediate outputs produced from key/value vectors of different lengths in Lean Attention. We shall prove that

$f(f(x,y),z) = f(x, f(y,z)) = f(x,y,z)$, where: $f(x,y) = \tilde{\mathbf{O}}^{(x,y)}$, $f(y,z) = \tilde{\mathbf{O}}^{(y,z)}$ and $f(x,y,z) = \tilde{\mathbf{O}}^{(x,y,z)}$. The key vectors $K^{(x)}$, $K^{(y)}$ and $K^{(z)}$ and value vectors $V^{(x)}$, $V^{(y)}$ and $V^{(z)}$ are of lengths $B_c^{(x)}$, $B_c^{(y)}$ and $B_c^{(z)}$ that are not necessarily equal.

Proving that $f(f(x,y),z) = f(x,y,z)$:

$$f(x,y) = \tilde{\mathbf{O}}^{(x,y)}$$

$$\begin{aligned}
m^{((x,y),z)} &= \max(m^{(x,y)}, m^{(z)}) \\
&= \max(\max(m^{(x)}, m^{(y)}), m^{(z)}) \\
&= \max(m^{(x)}, m^{(y)}, m^{(z)}) \\
&= m^{(x,y,z)}
\end{aligned}$$

$$\begin{aligned}
\ell^{((x,y),z)} &= e^{m^{(x,y)} - m^{((x,y),z)}}\ell^{(x,y)} + e^{m^{(z)} - m^{((x,y),z)}}\ell^{(z)} \\
&= e^{m^{(x,y)} - m^{(x,y,z)}}\ell^{(x,y)} + e^{m^{(z)} - m^{(x,y,z)}}\ell^{(z)} \\
&= e^{m^{(x,y)} - m^{(x,y,z)}} \\
&\times (e^{m^{(x)} - m^{(x,y)}}\ell^{(x)} + e^{m^{(y)} - m^{(x,y)}}\ell^{(y)}) \\
&+ e^{m^{(z)} - m^{(x,y,z)}}\ell^{(z)} \\
&= e^{m^{(x)} - m^{(x,y,z)}}\ell^{(x)} + e^{m^{(y)} - m^{(x,y,z)}}\ell^{(y)} \\
&+ e^{m^{(z)} - m^{(x,y,z)}}\ell^{(z)} \\
&= \ell^{(x,y,z)}
\end{aligned}$$

$$\begin{aligned}
f(f(x,y),z) &= \mathrm{diag}(e^{m^{(x,y)} - m^{((x,y),z)}})\tilde{\mathbf{O}}^{(x,y)} \\
&+ \mathrm{diag}(e^{m^{(z)} - m^{((x,y),z)}})\tilde{\mathbf{O}}^{(z)} \\
&= \mathrm{diag}(e^{m^{(x,y)} - m^{(x,y,z)}})\tilde{\mathbf{O}}^{(x,y)} \\
&+ \mathrm{diag}(e^{m^{(z)} - m^{(x,y,z)}})\tilde{\mathbf{O}}^{(z)} \\
&= \mathrm{diag}(e^{m^{(x,y)} - m^{(x,y,z)}})
\end{aligned}$$

$$
\begin{aligned}
&\times (\text{diag}(e^{m^{(x)}-m^{(x,y)}})\tilde{\mathbf{O}}^{(x)} \\
&+ \text{diag}(e^{m^{(y)}-m^{(x,y)}})\tilde{\mathbf{O}}^{(y)}) \\
&+ \text{diag}(e^{m^{(z)}-m^{(x,y,z)}})\tilde{\mathbf{O}}^{(z)} \\
&= \text{diag}(e^{m^{(x)}-m^{(x,y,z)}})\tilde{\mathbf{O}}^{(x)} \\
&+ \text{diag}(e^{m^{(y)}-m^{(x,y,z)}})\tilde{\mathbf{O}}^{(y)} \\
&+ \text{diag}(e^{m^{(z)}-m^{(x,y,z)}})\tilde{\mathbf{O}}^{(z)} \\
&= \tilde{\mathbf{O}}^{(x,y,z)} = f(x,y,z)
\end{aligned}
$$

Therefore, $\ell^{((x,y),z)} = \ell^{(x,y,z)}$ and $f(f(x,y),z) = f(x,y,z)$. For brevity, we omit the proof of $f(x,f(y,z)) = f(x,y,z)$, but it can deduced in a similar manner.

This associativity of softmax re-scaling is leveraged in Lean Attention to concurrently calculate the "partial" outputs produced from unequally sized contexts and then "reduce" them to obtain the final exact attention. The overall flow is described in Algorithm 2.

Algorithm. Similar to FlashAttention-2, we first partition the attention score matrix into tiles of dimensions $T_m \times T_n$. This corresponds to the nested for loops in Line 6 and 7 of Algorithm 2. The query activation matrix and the attention output matrix is partitioned into C_m tiles of sizes $T_m \times d$ each and the key and value activation matrices are partitioned into C_n tiles of sizes $T_n \times d$, making the attention matrix S a grid of $C_m C_n$ tiles, as seen in Line 4 and 5 (see Figure 6). We calculate the $query \times key$ MatMul (Line 12) followed by softmax to give the attention score matrix (Line 14), and then computing the $attn\_score \times value$ MatMul to give the partial attention output matrix as seen in Line 15. Finally, we extract the softmax re-scaling out of the inner loop of FlashAttention-2 and treat it as a reduction operation to accumulate the partial output tiles (Line 20-24).

B. LeanTile

To enables us to efficiently distribute the work of computing the partial output tensors, we define a block as a *LeanTile*. A single lean tile iteration computes "local attention" across a subset of tokens along the N_k dimension. Thus, a LeanTile takes in a query, key, and value tensors and computes the local attention to generate the un-scaled attention outputs.

Algorithm 3 depicts the subroutine for computing the partial attention outputs for that tile. This LeanTile() subroutine is called when computing each partial output tile in a CTA launched in Lean Attention as will be discussed later (Algorithm 4).

To efficiently split attention into smaller tiles, it is necessary to identify the smallest tile size capable of achieving highest compute efficiency. After extensively sweeping through various sizes for a *LeanTile*, we found that a tile size granularity of 256 and 128 tokens along the N_k dimension to be the most optimal for a head size of 64 and 128 respectively for FP16→32 problems while experimenting on an A100 GPU [11], [20].

Algorithm 2 Lean Attention - Basic Algorithm

1: **Require:** Load matrices $Q \in R^{N_q \times d}$ and $K, V \in R^{N_k \times d}$ into GMEM.
2: **Require:** Initialize Matrix O to $(0)_{N_q \times d} \in R^{N_q \times d}$ in GMEM.
3: Set block sizes T_m and T_n.
4: Partition Q, O as $Q_i, O_i \in R^{T_m \times d}$ where $i \in (1, C_m)$.
5: Partition K, V as $K_j, V_j \in R^{T_n \times d}$ where $j \in (1, C_n)$.
6: **for** $i = 1$ **to** C_m **do**
7:     **for** $j = 1$ **to** C_n **do**
8:         Load Q_i, K_j, V_j from GMEM to SMEM of an SM.
9:         Initialize O_{ij} to $(0)_{T_m \times d} \in R^{T_m \times d}$ in SMEM.
10:        Initialize m_{ij} to $(-\infty)_{T_m \times 1}$ and l_{ij} to $(0)_{T_m \times 1} \in R^{T_m \times 1}$ in SMEM.
11:        Compute on-chip:
12:        $S_{ij} = Q_i K_j^T$ where $S_{ij} \in R^{T_m \times T_n}$
13:        $m_{ij} = rowmax(S_{ij})$ where $m_{ij} \in R^{T_m \times 1}$
14:        $P_{ij} = exp(S_{ij} - m_{ij})$ where $P_{ij} \in R^{T_m \times T_n}$
15:        $O_{ij} = P_{ij} V_j$ where $O_{ij} \in R^{T_m \times d}$
16:        $l_{ij} = rowsum(P_{ij})$ where $l_{ij} \in R^{T_m \times 1}$
17:     **end for**
18: **end for**
19: **for** $i = 1$ **to** C_m **do**
20:     **for** $j = 1$ **to** $C_n - 1$ **do**
21:        $m_i^{new} = max(m_i, m_{ij})$
22:        $l_i^{new} = e^{m_i - m_i^{new}} l_i + e^{m_{ij} - m_i^{new}} l_{ij}$
23:        $O_i^{new} = e^{m_i - m_i^{new}} O_i + e^{m_{ij} - m_i^{new}} O_{ij}$
24:        Update $m_i = m_i^{new}, l_i = l_i^{new}$
25:     **end for**
26:     Compute $O_i = diag(l_i)^{-1} O_i$ and write to GMEM.
27:     Compute logexpsum $L_i = m_i + log(l_i)$ and write to GMEM.
28: **end for**

This optimal size can similarly be identified for other head dimensions and hardware architectures.

C. Decomposition and Mapping of LeanTiles

Finally, Lean Attention uses a stream-K [27] style decomposition and mapping of these LeanTiles to deliver efficient execution of attention.

Stream-K Decomposition. Stream-K is a parallel decomposition technique for dense matrix-matrix multiplication on GPUs. Stream-k partitioning addresses the inefficiencies in fixed-split by dividing the total workload (MAC operations) equally to all the CTAs using a pre-determined optimal tile size for dense matrix-matrix multiplications. It does this by rolling out the inner mode iterations of all output tiles and appending them along the inner mode to form a linear mapping. With the given grid size, it divides this total work into buckets demarcated appropriately such that each CTA has equal amount of iterations to perform. This grid size is determined by heuristics that sweep through all possible grid sizes and find the most optimal one which enables extensive parallelism and optimal wave

Algorithm 3 LeanTile() for a sequence of lean tile iterations

1: **function** LeanTile(tile_idx, iter_begin, iter_end)
2:   **_shared_** $O_{acc}[T_m, d]$
3:   **_shared_** $Q_f[T_m, d]$
4:   **_shared_** $K_f[T_n, d]$
5:   **_shared_** $V_f[T_n, d]$
6:   **_shared_** $m[T_m, 1]$
7:   **_shared_** $l[T_m, 1]$
8:   Initialize O_{acc} to $(0)_{T_m \times d} \in R^{T_m \times d}$ in SMEM.
9:   Initialize m to $(-\infty)_{T_m \times 1}$ and l to $(0)_{T_m \times 1} \in R^{T_m \times 1}$ in SMEM.
10:   $mm = T_m \times$ (tile_idx / 1)
11:   $nn = d \times$ (tile_idx % 1)
12:   Perform lean tile iterations for this output tile.
13:   **for** $iter = iter\_begin$ **to** $iter\_end$ **do**
14:     $kk = iter \times T_n$
15:     load fragments from GMEM to SMEM
16:     $Q_f = LoadFragment(Q, mm, nn)$
17:     $K_f = LoadFragment(K, nn, kk)$
18:     $V_f = LoadFragment(V, nn, kk)$
19:     Compute on chip:
20:     $S_f = Q_f K_f$ where $S_f \in R^{T_m \times T_n}$
21:     $m^{new} = max(m, rowmax(S_f))$
22:     $P_f = exp(S_f - m^{new})$ where $P_f \in R^{T_m \times T_n}$
23:     $l^{new} = e^{m-m^{new}} l + rowsum(P_f)$
24:     $O_{acc} = P_f V_f + diag(e^{m-m^{new}})O_{acc}$
25:     $l = l^{new}, m = m^{new}$
26:   **end for**
27:   **return** O_{acc}, l, m
28: **end function**

quantization that compensate well for any overhead that comes from reduction of the partial outputs.

LeanAttention extends Stream-K style of linear mapping of iterations, LeanAttention rolls out LeanTile iterations in a similar fashion, assigning equal number of N_k token iterations to each lean attention CTA as shown in Figure 6. Each CTAs range of N_k iterations is mapped contiguously into the batch size →heads →context length linearization, crossing the head and query boundary as it may. Should a given CTA's starting and/or ending iterations not coincide with the head boundary, it must consolidate its partial outputs with those of the other CTA(s) also convering that tile. In our implementation of LeanAttention, each output attention tensor is computed by the CTA that performed the tile's $N_k = 0$ token (called as a host block). Before it can do so, however, it must accumulate the un-scaled output tensors from other CTAs in temporary global storage, as shown in Figure 1. The negligible synchronization overhead of original stream-K implementation also extend to LeanAttention, thus leading to near 100% occupancy of SMs (not tensor core utilization) during the execution of a single CTA. Note that the temporary global storage overhead is minimal in the case of decode-phase where the output tensors are of dimensions $1 \times head\_dim$, where $head\_dim$ is typically in the range of 64 to 256.

Further, since we distribute the overall attention problem into optimal LeanTiles, we achieve a near 100% quantization efficiency irrespective of problem size (context length). This cohesive implementation of parallel computation and reduction happens in a single kernel launch in LeanAttention, avoiding the reduction kernel launch overheads that FlashDecoding suffers from. A difference in Stream-K decomposition in Lean Attention is in the reduction or "fix-up" phase. While Stream-K for MatMuls has addition as its reductive operation, Lean Attention has softmax rescaling and accumulation as it's reductive operation.

Naturally, some CTA's will be computing LeanTile iterations of more than one output tile. In such cases, stream-K's equalized partitioning makes lean attention more adept for problem sizes which would not occupy the hardware well if executed using its counterparts, FlashAttention-2 and FlashDecoding. To enable such a smooth transition between tiles, the input tensor view is also different in LeanAttention compared to FlashAttention-2. This requires a constant stride moving between different heads as we transition from a LeanTile of a head to another requiring query, key, and value tensors be of the shape $(batch\_size, heads, query/ctx\_length, head\_dim)$ compared to FlashAttention-2's requirement of $(batch\_size, query/ctx\_length, heads, head\_dim)$.

With this design of execution, LeanAttention behaves as a versatile attention partitioning mechanism that generalizes to FlashAttention-2 in the case where the number of output tiles is equal to grid size, and generalizes to FlashDecoding when grid size is an even multiple of number of output tiles. Finally, for all other cases (most common) LeanAttention efficiently distributes the work across the compute resources available in the system. Thus, LeanAttention will either always perform better or the same as FlashAttention-2 and FlashDecoding.

D. Execution Flow

Algorithm 4 depicts a StreamK style execution of Lean Attention. For a fixed grid size G, CTA's are launched and given equal amount of lean tiles to work with (Line 7). Each CTA block computes LeanTile() iterations for every distinct output tile that comes under its boundaries (Line 16).

The unique reduction phase of LeanAttention characterized by it's softmax rescaling and output tile accumulation is performed by the host CTA block. A host CTA (Line 17) is the CTA responsible for computing the first ever LeanTile for a given output tile, and it behaves as the consumer tile during parallel reduction of partial tiles. Similarly, a finishing CTA block is the block which computes the last ever LeanTile for a given output tile (line 18).

All non-host blocks will share their partials through global memory and signal their arrival (Line 20-23). On the other hand, a host block which is a non-finishing block (Line 24-25) needs to wait for other contributing peer CTA blocks to signal their completion (Line 28) and then proceed to carry out the reduction (Lines 29-35).

9

Algorithm 4 Lean Attention - StreamK Style Execution

1: **_shared_** $O[T_m, d]$
2: **_shared_** $m[T_m, 1]$
3: **_shared_** $l[T_m, 1]$
4: Number of output tiles: $C_m = \lceil N_q/T_m \rceil$
5: Number of iterations for each output tile: $C_n = \lceil N_k/T_n \rceil$

6: Total number of iterations: $I = C_m C_n$
7: Number of iterations per CTA: $I_G = I/G$
8: **fork** CTA$_g$ in G **do**
9:   cta_start = g I$_G$ and cta_end = cta_start + I$_G$
10:   **for** iter = cta_start **to** cta_end **do**
11:     Index of current output tile: tile_idx = iter / C_n
12:     tile_iter = tile_idx $\times C_n$
13:     tile_iter_end = tile_iter + C_n
14:     local_iter = iter - tile_iter
15:     local_iter_end = min(tile_iter_end, cta_end) - tile_iter
16:     O, m, l = LeanTile(tile_idx, local_iter, local_iter_end)
17:     host-block if: iter == tile_iter
18:     finishing-block if: cta_end $>=$ tile_iter_end
19:     **if** !(host-block) **then**
20:       StorePartials(Op[g], O)
21:       StorePartials(mp[g], m)
22:       StorePartials(lp[g], l)
23:       Signal(flags[g])
24:     **else**
25:       **if** !(finishing-block) **then**
26:         last_cta = tile_iter_end / C_n
27:         **for** cta = (g + 1) **to** last_cta **do**
28:           Wait(flags[cta])
29:           m_{cta} = LoadPartials(mp[cta])
30:           l_{cta} = LoadPartials(lp[cta])
31:           O_{cta} = LoadPartials(Op[cta])
32:           $m^{new} = max(m_{cta}, m)$
33:           $l^{new} = e^{m_{cta}-m^{new}}l_{cta} + e^{m-m^{new}}l$
34:           $O^{new} = e^{m_{cta}-m_i^{new}}O_{cta} + e^{m-m_i^{new}}O$
35:           Update $m = m_i^{new}, l = l_i^{new}$
36:         **end for**
37:       **end if**
38:       Write $O = diag(l)^{-1}O$ to GMEM.
39:       Write $L = m + log(l)$ to GMEM.
40:     **end if**
41:     iter = tile_iter_end
42:   **end for**
43: **join**

A host block that is also a finishing block completes all the LeanTile iterations for its output tile in a single CTA and so can directly store its results from LeanTile() in global memory (Line 38-39) without any reduction.

V. EVALUATION METHODOLOGY

Implementation. We implement LeanAttention using the CUTE abstractions [1]–[3] provided by Nvidia's CUT-LASS library [6]. For comparative measurements we utilize FlashAttention-2's implementation as it is available on their Github repository [5][2]. For the end-to-end inference results we use OPT models as available in the HuggingFace Transfomers repository [33] and modify them to allow execution via LeanAttention wherever necessary.

System. We benchmark the attention mechanisms on Nvidia-A100-80GB-GPU [11] system with up to 8 GPUs. We measure runtime using a single GPU as well as 8xGPUs for larger models and context lengths. A single A100 GPU consists of 108 streaming multiprocessors (SMs) with an 80GB HBM global memory.

Multi-GPU Tensor Parallelism. We utilize Tensor Parallelism for the multi-GPU measurements to reflect the large model executions. Since FlashAttention-2 does not support Tensor Parallelism, we scale the implementation to the total number of SMs available in the system.

Attention Mechanism. In addition to vanilla FlashAttention-2 as it is discussed in the original paper [13], we also benchmark FlashDecoding for comparison against LeanAttention. For the rest of the paper we refer to vanilla FlashAttention-2 as FA2, FlashDecoding as FD, and LeanAttention as LA.

VI. EVALUATION RESULTS

In this section, we evaluate the impact of LeanAttention (LA) at the attention operation-level as well as end-to-end inference performance.

A. Benchmarking Attention - Decode-Phase

We benchmark the runtime of just the attention operation using the different mechanisms at varying context lengths, number of attention heads, head dimensions (64:default and 128), and inference batch sizes on a single Nvidia A100-80GB GPU.

Increasing Context Length. Figure 7(a) shows the speedup of different attention mechanisms for a model with 56 attention heads with a single inference batch. LA delivers more than 2x speedup compared to FA2 for context lengths larger than 8k, reaching up to 2.46x speedup as the context lengths grows to 512k tokens. FD reaches a decent 2.06x speedup compared to vanilla FA2 but shows negligible increase in speedup as context lengths grow beyond 32k.

Increasing Attention Heads. Figure 7(b) shows the speedup delivered by LA compared to FA2 for models with an increasing number of heads. LA delivers a speedup of 12.5x at small model sizes working on a context length of 64k tokens. As the the number of heads in the model increases, FA2 starts leveraging the SMs in the GPUs more equally, however, LA still delivers a 2.15x improvement over FA2 at 64 heads. This shows that LA is able to scale well for both a small and large number of heads. FD scales better than vanilla FA2 but is not able to optimally utilize all compute capabilities.

Effect of Batching. Figure 7(c) shows the performance improvement of LA at varying batch sizes. As expected, we
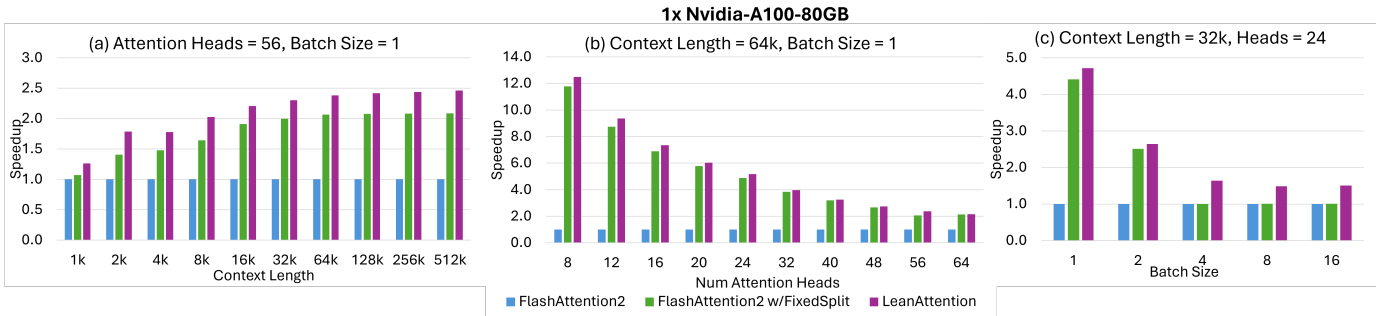
[2]Version 2.5.6

10

Fig. 7. Speedup of LA compared to state-of-the-art Attention execution mechanisms at different context lengths, batch sizes and attention heads with head dimension = 64 on a single Nvidia-A100-80GB GPU.
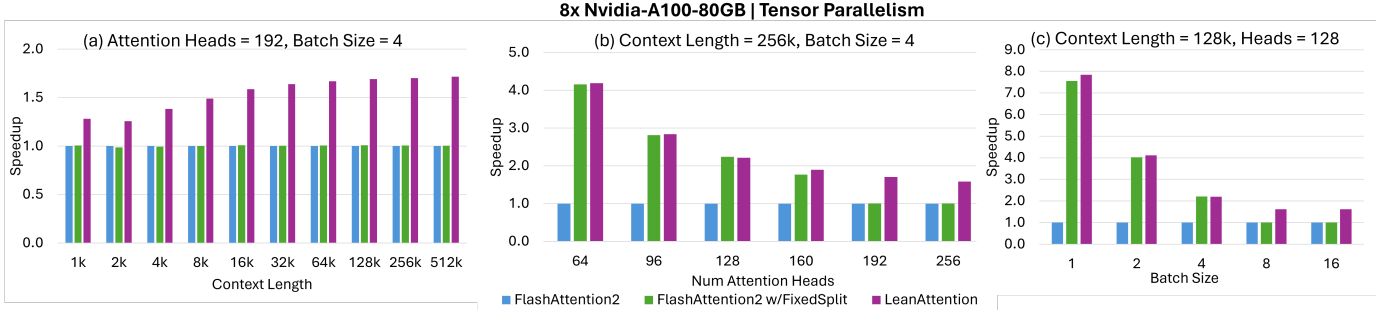

Fig. 8. Speedup of LA compared to state-of-the-art Attention execution mechanisms at different context lengths, batch sizes, attention heads, etc., on an 8x Nvidia-A100-80GB system.
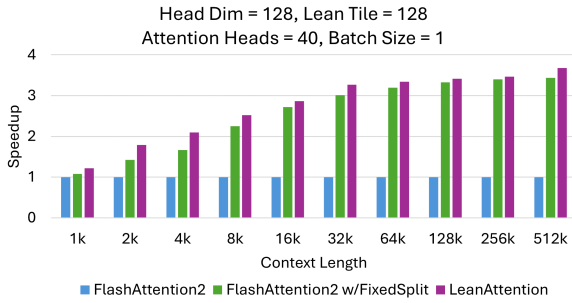

Fig. 9. Speedup offered by LA at a head dimension of 128. Note that the optimal LeanTile size of head dimension of 128 is 128 token wide in constrast to 256 tokens for head dimension = 64
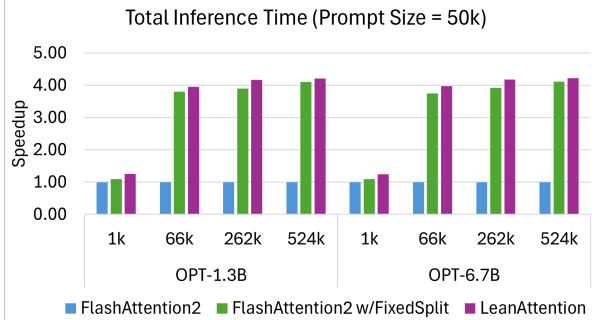

Fig. 10. End-to-End Speedup of LA compared to other execution mechanisms at different context lengths, batch size = 1

observe that LA outperforms FA2 and FD by 4.71x and 1.06x respectively at single batches. Even at higher batch sizes where FA2 has enough heads to provide all SMs with enough work, LA's efficient work allocation policy leads to outperforming both FA2 and FD by 1.5x.

Overall, we benchmarked the system over 500 samples with varying number of batch sizes, context lengths, and attention heads and observed an average 2.6x (Max:8.33x for 16 heads 512k context, Min:1.1x for 24 heads 1k context) speedup over FA2 and 1.27x (Max:1.71x for 24 heads 512k context, Min:0.99x for 24 heads 1k context) speedup over FD.

Multi-GPU Execution Repeating a similar benchmarking process on an 8xA100 GPU system, we vary the context lengths from 1k to 512k, with 192 attention heads at a batch size of 4 as shown in Figure 8(a). LeanAttention reaches a speedup of over 1.70x at longer context lengths (over 64k), while still giving a speedup of over 1.28x for smaller context lengths. This is because parallelizing over just the batch and heads (total heads = $192 \times 4 = 768$) here does not give sufficient work to each SM (total SMs = $8 \times 108 = 864$). Moreover, FD behaves exactly the same as the vanilla FA2, since there are not enough extra SMs available for splitting by 2 ($768 \times 2 = 1536 > 864$). LA is able to optimally distribute the work across the available SMs in the system.

To observe this effect in greater detail, we evaluate across a varying number of attention heads in Figure 8(b) with a context length of 256k and batch size of 4. We observe speedup of 4.18x at a smaller number of heads (64). We can see that FA2 is using more SMs as we increase the number of attention heads, but not as efficiently as LA. We can also clearly see that FD resorts to vanilla execution when we increase the number of

11

heads from 128 to 192. LA, on the other hand scales smoothly as we increase the number of heads, showcasing its hardware-aware scalable execution algorithm. With 128 heads, as we increase the batch size from 1 to 16, we can see that LA outperforms FA2 variants at lower batches delivering 7.8x speedup at single batch. This demonstrates LA's capability of delivering performance in scenarios where batching is not trivial.

Effect of Head Dimension. Figure 9 shows the speedup offered by LA for models with a head dimension of 128. We utilize a 128-token wide LeanTile for decomposition of each problem instead of 256. We observe a similar trend in performance, where LA delivers a speedup of 3.67x compared to FA2 at 512k context length. Even at smaller context lengths of 1k tokens, we observed an improved performance of 1.2x.

To summarize, LA not only outperforms FA2 at lower batch sizes, long context lengths, but also delivers better performance at higher batch sizes, and with higher number of attention heads. This is mainly due to the lean decomposition of the problem on the hardware compute resource. For cases where there are enough parallelizable dimensions, LA automatically generalizes to FA2-like execution. We can thus treat FA2's execution algorithm as a special case of LA, which occurs depending on the optimal grid size that LeanAttention chooses depending on the hardware resources and LeanTile dimensions.

B. End-to-End Inference Performance

We measure the end-to-end inference runtime using OPT models of 1.3 Billion and 6.7 Billion (32 heads each) parameters as shown in Figure 10 at increasing output token lengths with a prompt size of 50k tokens (total context = prompt tokens + tokens generated so far). This includes the prefill-stage latency as well as the total runtime of decode-phase. LeanAttention offers a 1.26x speedup with OPT 1.3B as compared to FlashAttention2 for first 1k output tokens. However, the LA offers a higher speedup as the output tokens increase beyond 64k delivering an average of 4.0x speedup compared to FA2 and 1.06x speedup over FD. Note that these numbers reflect the performance observed with 32 heads in the previous section. As we note, the inference-level runtime improvement delivered by LA will change heavily on the number of heads, total context length, batch size, etc.

VII. CONCLUSION

The attention mechanism in transformer-based language models is a slow and memory hungry process. State-of-the-art optimization mechanisms, such as FlashAttention-2, have cleverly addressed this challenge; however, they fail to adapt to the computationally distinct phases of inference. We observe that FlashAttention-2 fails at parallelization across the context-length dimension of operation during the decode phase of inference, thus resulting in low occupancy of the underlying hardware. As state-of-the-art models continue to push the limits on supporting increasingly long context lengths, optimization techniques that optimally parallelize over this dimension will become increasingly important.

To address this challenge, we propose *LeanAttention*, a scalable, exact-attention execution mechanism, specifically designed to optimize the decode phase of generative transformer models, though general enough to optimize the prefill-phase as well. LeanAttention utilizes the associative property of online softmax calculation as a reductive property and extrapolates the state-of-the-art "stream-K" matrix decomposition technique to the attention mechanism. This allows us to efficiently parallelize and execute the attention mechanism for the decode phase of inference. Our measurements show that LeanAttention delivers an average speedup of 2.6x over FlashAttention-2 and offers up to 8.33x speedup for 512k context sizes compared to FlashAttention-2. Notably, in a multi-GPU execution scenario with a large number of attention heads, the speedup realized by LeanAttention continues to increase as context length increases, providing 1.7x speedup over FlashAttention-2 (and FlashDecoding) at 512k context length. LeanAttention delivers near 100% processor occupancy, enabling the efficient scaling of next generation LLMs that leverage large context lengths.

REFERENCES

[1] "Cute layouts." https://github.com/NVIDIA/cutlass/blob/main/media/docs/cute/01_layout.md, [Accessed 19-04-2024].

[2] "Cute tensors." https://github.com/NVIDIA/cutlass/blob/main/media/docs/cute/03_tensor.md., [Accessed 19-04-2024].

[3] "Cute's support for matrix multiply-accumulate instructions." https://github.com/NVIDIA/cutlass/blob/main/media/docs/cute/0t_mma_atom.md, [Accessed 19-04-2024].

[4] "Flashdecoding: Stanford CRFM — crfm.stanford.edu," https://crfm.stanford.edu/2023/10/12/flashdecoding.html, [Accessed 22-04-2024].

[5] "GitHub - Dao-AILab/flash-attention: Fast and memory-efficient exact attention — github.com," https://github.com/Dao-AILab/flash-attention, [Accessed 19-04-2024].

[6] "GitHub - NVIDIA/cutlass: CUDA Templates for Linear Algebra Subroutines — github.com," https://github.com/NVIDIA/cutlass, [Accessed 01-04-2024].

[7] "Introducing ChatGPT — openai.com," https://openai.com/blog/chatgpt, [Accessed 01-04-2024].

[8] "Introducing the next generation of claude," https://www.anthropic.com/news/claude-3-family, [Accessed 19-04-2024].

[9] W. Brandon, A. Nrusimha, K. Qian, Z. Ankner, T. Jin, Z. Song, and J. Ragan-Kelley, "Striped attention: Faster ring attention for causal transformers," 2023.

[10] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, "Language models are few-shot learners," *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.

[11] J. Choquette, E. Lee, R. Krashinsky, V. Balan, and B. Khailany, "3.2 the a100 datacenter gpu and ampere architecture," in *2021 IEEE International Solid-State Circuits Conference (ISSCC)*, vol. 64. IEEE, 2021, pp. 48–50.

[12] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann *et al.*, "Palm: Scaling language modeling with pathways," *Journal of Machine Learning Research*, vol. 24, no. 240, pp. 1–113, 2023.

[13] T. Dao, "Flashattention-2: Faster attention with better parallelism and work partitioning," 2023.

[14] T. Dao, D. Y. Fu, S. Ermon, A. Rudra, and C. Ré, "Flashattention: Fast and memory-efficient exact attention with io-awareness," 2022.

[15] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," 2019.

[16] Y. Fu, R. Panda, X. Niu, X. Yue, H. Hajishirzi, Y. Kim, and H. Peng, "Data engineering for scaling language models to 128k context," *arXiv preprint arXiv:2402.10171*, 2024.

[17] K. Hong, G. Dai, J. Xu, Q. Mao, X. Li, J. Liu, K. Chen, Y. Dong, and Y. Wang, "Flashdecoding++: Faster large language model inference on gpus," 2024.

[18] G. Inc., "An important next step on our AI journey — blog.google," https://blog.google/technology/ai/bard-google-ai-search-updates/, [Accessed 31-03-2024].

[19] A. Ivanov, N. Dryden, T. Ben-Nun, S. Li, and T. Hoefler, "Data movement is all you need: A case study on optimizing transformers," *Proceedings of Machine Learning and Systems*, vol. 3, pp. 711–732, 2021.

[20] Z. Jia and P. Van Sandt, "Dissecting the ampere gpu architecture via microbenchmarking," in *GPU Technology Conference*, 2021.

[21] S. Kim, C. Hooper, T. Wattanawong, M. Kang, R. Yan, H. Genc, G. Dinh, Q. Huang, K. Keutzer, M. W. Mahoney *et al.*, "Full stack optimization of transformer inference: a survey," *arXiv preprint arXiv:2302.14017*, 2023.

[22] T. Li, G. Zhang, Q. D. Do, X. Yue, and W. Chen, "Long-context llms struggle with long in-context learning," 2024.

[23] Y. Li, S. Bubeck, R. Eldan, A. D. Giorno, S. Gunasekar, and Y. T. Lee, "Textbooks are all you need ii: phi-1.5 technical report," 2023.

[24] H. Liu, M. Zaharia, and P. Abbeel, "Ring attention with blockwise transformers for near-infinite context," 2023.

[25] M. Milakov and N. Gimelshein, "Online normalizer calculation for softmax," 2018.

[26] OpenAI, J. Achiam, and S. Adler, "Gpt-4 technical report," 2024.

[27] M. Osama, D. Merrill, C. Cecka, M. Garland, and J. D. Owens, "Stream-k: Work-centric parallel decomposition for dense matrix-matrix multiplication on the gpu," 2023.

[28] R. Pope, S. Douglas, A. Chowdhery, J. Devlin, J. Bradbury, J. Heek, K. Xiao, S. Agrawal, and J. Dean, "Efficiently scaling transformer inference," *Proceedings of Machine Learning and Systems*, vol. 5, 2023.

[29] J. Spataro and M. Inc., "Introducing Microsoft 365 Copilot – your copilot for work - The Official Microsoft Blog — blogs.microsoft.com," https://blogs.microsoft.com/blog/2023/03/16/introducing-microsoft-365-copilot/-your-copilot-for-work/, [Accessed 31-03-2024].

[30] H. Touvron and L. Martin, "Llama 2: Open foundation and fine-tuned chat models," 2023.

[31] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," 2023.

[32] S. Williams, A. Waterman, and D. Patterson, "Roofline: an insightful visual performance model for multicore architectures," *Communications of the ACM*, vol. 52, no. 4, pp. 65–76, 2009.

[33] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. L. Scao, S. Gugger, M. Drame, Q. Lhoest, and A. M. Rush, "Huggingface's transformers: State-of-the-art natural language processing," 2020.

[34] G.-I. Yu, J. S. Jeong, G.-W. Kim, S. Kim, and B.-G. Chun, "Orca: A distributed serving system for {Transformer-Based} generative models," in *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*, 2022, pp. 521–538.

[35] G.-I. Yu, J. S. Jeong, G.-W. Kim, S. Kim, and B.-G. Chun, "Orca: A distributed serving system for Transformer-Based generative models," in *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*. Carlsbad, CA: USENIX Association, Jul. 2022, pp. 521–538. [Online]. Available: https://www.usenix.org/conference/osdi22/presentation/yu

[36] P. Zhang, Z. Liu, S. Xiao, N. Shao, Q. Ye, and Z. Dou, "Soaring from 4k to 400k: Extending llm's context with activation beacon," *arXiv preprint arXiv:2401.03462*, 2024.

[37] S. Zhang, S. Roller, N. Goyal, M. Artetxe, M. Chen, S. Chen, C. Dewan, M. Diab, X. Li, X. V. Lin *et al.*, "Opt: Open pre-trained transformer language models," *arXiv preprint arXiv:2205.01068*, 2022.