

---

# Generalized Neighborhood Attention: Multi-dimensional Sparse Attention at the Speed of Light

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

Many sparse attention mechanisms such as Neighborhood Attention have typically failed to consistently deliver speedup over the self attention baseline. This is largely due to the level of complexity in attention infrastructure, and the rapid evolution of AI hardware architecture. At the same time, many state-of-the-art foundational models, particularly in computer vision, are heavily bound by attention, and need reliable sparsity to escape the  $\mathcal{O}(n^2)$  complexity. In this paper, we study a class of promising sparse attention mechanisms that focus on locality, and aim to develop a better analytical model of their performance improvements. We first introduce Generalized Neighborhood Attention (GNA), which can describe sliding window, strided sliding window, and blocked attention. We then consider possible design choices in implementing these approaches, and create a simulator that can provide much more realistic speedup upper bounds for any given setting. Finally, we implement GNA on top of a state-of-the-art fused multi-headed attention (FMHA) kernel designed for the NVIDIA Blackwell architecture in CUTLASS. Our implementation can fully realize the **maximum speedup theoretically possible** in many perfectly block-sparse cases, and achieves an effective utilization of 1.3 petaFLOPs/second in FP16. In addition, we plug various GNA configurations into off-the-shelf generative models, such as Cosmos-7B, HunyuanVideo, and FLUX, and show that it can deliver 28% to 46% end-to-end speedup on B200 without any fine-tuning. We will open source our simulator and Blackwell kernels directly through the *NATTEN* project.

## 1 Introduction

Fast sparse attention has been long sought-after [33, 37, 9, 5, 53, 43, 35], but rarely without complications. Infrastructure continues to be a challenge for attention in general, as implementations of attention rarely come close to adequately utilizing the computational power of modern GPUs, at least compared to dense linear algebra primitives, such as generalized matrix-matrix multiplications (GEMMs), which typically utilize around 80% of peak FLOPs/second. The most successful example to date is Flash Attention 3 [39], utilizing up to 75% of the peak FLOPs/second in the NVIDIA Hopper architecture with half precision. However, lower precision still trails behind, and with every new architecture, and changes in the programming model, new challenges in infrastructure arise. Sparse approaches that require changes to the core attention implementation have therefore lagged behind. One example is approaches that specifically target sparsity in attention weights, such as sliding window [33, 37, 5, 21] and blocked attention [43, 31], which can only accelerate computation by utilizing block-sparsity. Block-sparsity is primarily comprised of skipping tiles of computation that are fully masked in some pre-defined attention mask. Implementing block-sparsity can be non-trivial,

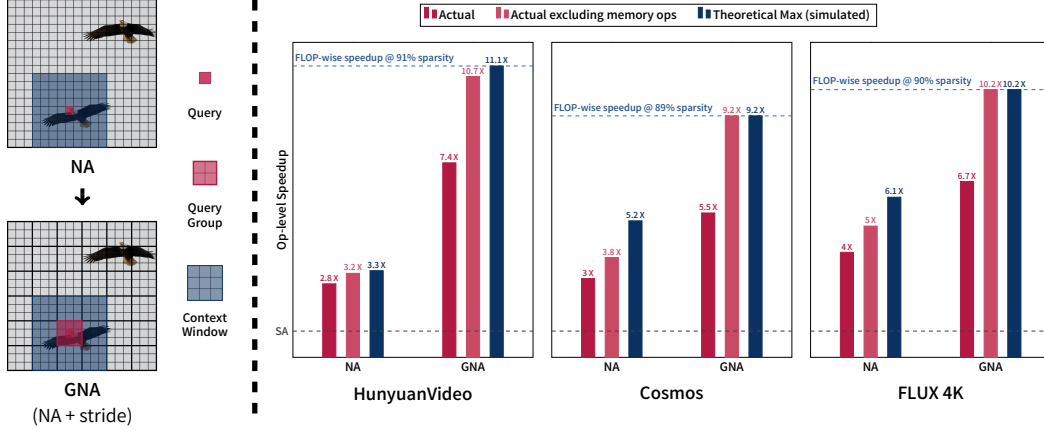


Figure 1: **Generalized Neighborhood Attention (GNA)** adds “stride” to NA, which forces multiple queries to share their context, increasing computation density, while maintaining sparsity, leading to speedups more proportional to savings in FLOPs. Our Blackwell kernel can as a result realize the **maximum speedup theoretically possible** in some cases, with respect to both FLOPs and simulation.

and at times comes with significant overhead that can undo performance gains, assuming the base implementation is already the state of the art. As a result, these approaches usually leave some performance on the table, with lower utilization than dense attention.

Despite these challenges, sliding window and blocked attention have been successfully built into large language models [25, 10, 3], where block-sparsity and fine-grained masking are concerned with a single-dimensional token layout. This makes their implementation significantly simpler compared to applications such as images and videos, which step into 2-D and 3-D token layouts. As a direct result of this, implementations focused on multi-dimensional token layouts typically fall behind even further. One such example is Fused Neighborhood Attention (FNA) [19], which implements multi-dimensional sliding window attention, but suffers a considerable performance gap due to overheads that worsen with higher dimensionality, namely software predication and fine-grained masking. At the same time, FNA targets the NVIDIA Ampere architecture, which results in even larger gap compared to the state-of-the-art on newer architectures, such as Hopper. While some successful implementations of multi-dimensional local attention have recently appeared [54], they are not as flexible in terms of pattern and parameters, offering only fully block-sparse masks, and are tied to very specific use cases. Presently a variety of frameworks that provide linear algebra and tensor layout primitives exist, such as CuTe and CULTASS [41], Triton [42], and compilers specifically designed for block-sparse attention, such as Flex Attention [14]. Despite the amount of impact this has had on the advancement of research, wildly different implementations of different approaches make it increasingly difficult to analyze the exact performance differences in these multi-dimensional sparse attention methods, and identify their root causes.

To that end, we propose Generalized Neighborhood Attention (GNA), an extension to Neighborhood Attention (NA) [21], itself a sliding window approach, which aims to unify many existing approaches under one definition. GNA can offer a tradeoff space between efficiency, and quality and inductive biases (i.e. translational equivariance). Approaches implementable with GNA can be classified into sliding window [33, 37, 21], strided sliding window [43, 54], and blocked [43, 31] attention. We revisit the problem of sliding window vs strided sliding window and blocked attention discussed in HaloNet [43], but with focus on speedup as opposed to memory operations. An illustration of different GNA patterns is presented in Fig. 2. We further create a simulation tool, which can compute the upper-bound speedup achievable by any approach defined under GNA, under different use cases, and implementation design choices. Through abstracting away implementation details, and defining speedup with tile/block granularity as opposed to FLOP/FMA granularity, we can finally fairly compare these different approaches analytically. Finally, we implement GNA on top of CUTLASS’s attention kernel for the NVIDIA Blackwell architecture, which is one of the best-performing choices available today, and show that we can successfully realize up to 100% of the FLOP-wise efficiency in many GNA configurations.

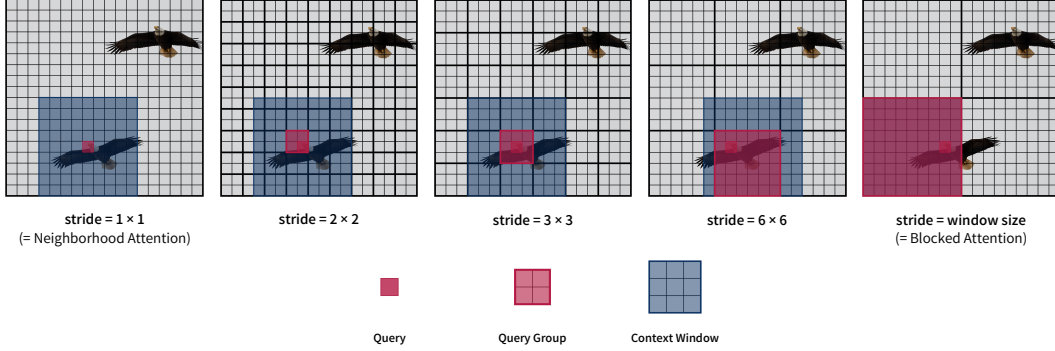


Figure 2: GNA introduces “delay steps” in the sliding window through the new **stride** parameter. Stride 1 is standard NA, and stride = window size is blocked attention (a.k.a. Window Self Attention).

Although attention itself typically does not come close to standard GEMMs in terms of Speed-of-Light (SoL) performance<sup>1</sup>, given the negligible performance overhead of our implementation, we believe our methodology can be a recipe for sparse and local attention methods to finally catch up to the performance of standard FMHA implementations. We identify variants of GNA that can exhibit perfect speedup, and introduce both them, and standard Neighborhood Attention into generative and foundation models, such as Cosmos [1], HunyuanVideo [27], and Flux [28]. We show that using our Blackwell kernel, which we have directly integrated into the *NATTEN* project, we can accelerate these models without fine-tuning up to 63% end-to-end. Beyond improving and accelerating *NATTEN*, and supporting more sparse attention patterns, we hope to provide a useful set of tools, such as our simulator, to help researchers find the most useful configurations for their models, and easily ablate over a useful set of parameters, instead of being lost at the outrageously large parameter space that GNA has to offer. All of our work will be open-sourced.

## 2 Related Works

Attention’s quadratic complexity is one of the key motivations behind research into sparse attention. Many approaches exist, such as structured sparsity [8], low-rank approximations [46, 7], and locality [33, 7, 37, 9, 5]. In this paper, we specifically focus on local methods, which are among the most successful [31, 26, 29, 23, 11, 16] due to the fact that they require little to no change to the Transformer [44] architecture, as well as spatial locality being a bias that exists in nature. Past [25] and recent [10, 3] large language models also adopt local methods into their architecture in addition to standard self attention. Some of the most popular local approaches are sliding window [33, 37, 21], dilated sliding window [5, 20], strided sliding window [43, 54], and blocked attention [43, 31]. Sliding window approaches have often been categorized as inefficient, for various reasons, while blocked attention is known for efficiency as well as being easy to implement.

However, these arguments have been made with varying implementation design choices, constraints, and use cases, and rarely analytically discuss the actual tradeoffs and inefficiencies when it comes to these approaches. For instance, until recently, the only choice for efficiently implementing multi-dimensional sliding window attention was *NATTEN*, but it has rarely caught up to state-of-the-art implementations such as Flash Attention [12, 39]. While fused kernels are available in *NATTEN*, their underlying kernel, as well as certain limitations of the implementation undo efficiency gains from sparsity in many cases. The main challenge in multi-dimensional local attention is what we call **the curse of multi-dimensionality**, which forces efficient implementations to either waste more compute, or introduce additional logic for multi-dimensional tiling to preserve spatial locality. We describe this in more detail in Appendix A.3.

Given that there are many approaches to local sparse attention, and each with various design choices in their implementation, and the various properties and performance levels of different hardware architectures, we aim to disambiguate their differences and similarities. We do so by introducing an

<sup>1</sup><https://docs.nvidia.com/nsight-compute/ProfilingGuide/index.html#roofline-charts>




Figure 3: Different strides and their analytical speedup according to *NATTENS*im, with window size  $18 \times 24 \times 24$  on HunyuanVideo use case. Simulation assumes Q tile shape  $4 \times 8 \times 8$  and KV tile shape  $2 \times 8 \times 8$ . NA (stride 1) is limited by a 3.3x speedup, while some larger strides can improve, and eventually reach block-sparsity, resulting in a speedup of 11.1x, equivalent to FLOP-wise speedup.

108 approach that unifies the above local approaches, analytically studying its performance implications  
 109 under various settings, and implementing it by learning from the strengths and shortcomings of prior  
 110 approaches [19]. We refer readers to Appendix A for a more detailed background and related works  
 111 section.

### 112 3 Methodology

113 In this section, we introduce Generalized Neighborhood Attention (GNA), which is aimed at providing  
 114 a tradeoff space between accuracy and translational equivariance and efficiency. We then describe  
 115 design choices in implementations of such methods, which is one of our motivations for creating  
 116 our analytical tool, *NATTENS*im. We finally move on to our implementation for the Blackwell  
 117 architecture. This implementation is based on one of the best-performing FMHA kernels available for  
 118 the Blackwell architecture.

#### 119 3.1 Generalized Neighborhood Attention

120 Neighborhood Attention (NA) [21], and more generally, sliding window attention patterns are fine-  
 121 grained masks that allow each query token to have a unique context window of fixed size, which  
 122 is determined based on its coordinate. The effect of this is similar to that of convolution, where  
 123 a filter of a fixed size is applied to sliding window regions, and contracted to a single point in the  
 124 output. Previous works have extended NA to support dilated windows [20], mimicking dilation in  
 125 convolution, as well as causal masking [19], which can be useful for video applications, where an  
 126 atomic spatio-temporal attention mask may be required.

127 In this work, we start with allowing even values for window size in NA. NA was defined on odd  
 128 values for window size only in order to center queries in the window. However, this can be limiting,  
 129 as we later show that window sizes that are multiples of tile sizes (usually powers of two) are more  
 130 efficient. We first split **window size** into two a left side and right side window size. In the odd  
 131 case (standard NA), both are the floor of window size divided by two. For even-sized windows,  
 132 we can choose either side to be larger by 1. We further add a fourth parameter to NA, which we  
 133 call “**stride**”. As in convolution, stride adds a delay step to the sliding windows in NA. Stride 1 is  
 134 standard NA, moving the sliding window for every query token, while stride 2 moves the sliding  
 135 window every two query tokens. Another way of viewing this is that stride **groups queries together**,  
 136 and forces them to share the same neighborhood (context window). Stride can take any positive  
 137 integer value less than or equal to window size. We present a visualization of different strides for  
 138 a fixed window size in Fig. 2. Strided NA exhibits similar behavior as blocked local attention in  
 139 HaloNet [43], with their “block size” parameter introducing the same delay step effect into sliding  
 140 window attention [33, 37] as stride does in neighborhood attention [21]. Another interesting property  
 141 in both is that when stride is equal to window size, the pattern will be equivalent to (fully) blocked



attention, also known as Window Self Attention (WSA) [31]. In some cases, Strided NA can also implement STA [54], but we note this is only guaranteed when certain assumptions are made with respect to the underlying implementation. We will clarify this further in Secs. 3.2 and 3.3.

Stride is motivated by the fact that *sliding window attention does not achieve speedup proportional to FLOP-wise sparsity*. This is because sliding window attention is a vector-matrix multiplication<sup>2</sup> problem commonly implemented with matrix-matrix multiplies (GEMMs) [19] and fine-grained masking. Masking ensures correctness, but it also wastes work (FLOPs) already done. Stride aligns the boundaries of their windows together, and bridges the gap between sliding window [33, 37, 21], and blocked approaches [43, 31, 54]. Blocked approaches can mainly perform dense computation, and often do not require fine-grained masking of attention weights, therefore waste fewer FLOPs. We therefore aim to answer questions such as: how much work is exactly wasted, what strides are more efficient, and by how much. In order to do so, we created a tool for simulating the behavior of common implementations, and use it to compute savings in *tiles* of work.

### 3.2 Analytical tool for GNA

Stride extends an already vast parameter space. This, along with the different design choices and configurations in implementation, specifically that of FMHA kernels, makes it difficult to find useful parameters for end-users that offer the best tradeoff. We therefore created an analytical tool called *NATTENS*im, which can shed light on exactly that. *NATTENS*im computes the number of context (KV) tiles visited by an implementation, taking into account various design choices such as dynamic vs static KV tiling, 1-D vs multi-dimensional tiling, and of course tile shapes for query (Q) and context (KV) tensors. We describe how *NATTENS*im works and design choices it considers in Appendix B.1. With *NATTENS*im, we can achieve the following:

1. Compare design choices, pick the one best trading off implementation difficulty and speedup.
2. Find perfectly block-sparse cases, where no fine-grained masking is required, and simulated speedup matches FLOP-wise speedup, as illustrated in Fig. 3. Under dynamic KV tiling, any setting in which  $T_{KV}$  evenly divides window size, and  $T_Q$  evenly divides stride achieves this. Under static KV tiling some simulation may be required to find those points.
3. Predict end-to-end speedup upper bounds for any given model under a specific set of parameters and design choices.


### 3.3 Implementation for the Blackwell architecture

We start off with the CUTLASS FMHA kernel for Blackwell [41], which can achieve up to 1.2 petaFLOPs/s with FP16. While our implementation specifically focuses on the Blackwell architecture, our design choices are primarily architecture-agnostic. We followed the design of the original FNA kernel [19], but instead rely on a small memory operation, dubbed token permutation, instead of fusing multi-dimensional tiling into the kernel like FNA. In the current implementation, we implement token permutation naively through PyTorch copy primitives, which utilize only 1/8th of the memory bandwidth on the B200. We foresee further possible optimizations to this approach. Additionally, many use cases can perform token permutation only once in the beginning of the model, and undo the permutation in the end, which minimizes its effect on the end-to-end workload. We present more details of our implementation, and reasons for the design choices made in Appendix B.2.

## 4 Experiments


In order to evaluate GNA’s applicability and performance, we carefully selected applications with multi-dimensional layouts of tokens that are suitable for sparse attention. Our primary criterion for choosing applications is whether or not at least 40% of their end-to-end workload is self attention specifically. Our final candidates are Cosmos [1] (World Foundation Model), HunyuanVideo [27] (video generation), and FLUX [28] (image generation). We additionally note that FLUX only meets the criterion when generating resolutions higher than 2K. We present the workload distribution for those models in Tab. 1.

<sup>2</sup>Generalized Matrix-Vector Multiply (GEMV) in BLAS.



(a) Cosmos use case with window size  $16 \times 24 \times 16$ .

(b) FLUX use case with window size  $80 \times 80$ .



(c) HunyuanVideo use case with window size  $18 \times 24 \times 24$ .

Figure 4: Op-level (attention only) speedups from GNA on generative model use cases. Analytical speedup is according to *NATTENS*im, and actual speedups are measured by running on B200. In perfectly block-sparse settings, our kernel can approximately or fully match analytical speedup, but can be limited by the naive implementation of the memory operation (token permute).

#### 190 4.1 GNA Performance

191 We first consider cases under which these applications introduce roughly 90% of sparsity into  
 192 attention, and run their problem shapes through *NATTENS*im, sweeping over all possible stride  
 193 values. For each use case, we selected Q and KV tile shapes ( $T_Q$  and  $T_{KV}$ ) according to the shape  
 194 of the token layout (feature map size). We chose window sizes that are evenly divisible by the KV  
 195 tile shape ( $T_{KV}$ ), as this can result in perfectly block-sparse forms of GNA, which waste no FLOPs  
 196 and can potentially achieve perfect speedup. A similar observation was made in STA [54], with the  
 197 exception that STA does not differentiate between  $T_Q$  and  $T_{KV}$ , and that their step size parameter  
 198 is fixed. We also prune simulator results by removing increased strides that do not result in better  
 199 speedups. This is helpful, since larger strides trade off translational equivariance and potentially  
 200 quality for potentially better performance, and if the latter is not realized under some setting, that  
 201 configuration is unlikely to have any advantage over others.

202 We present one instance of results obtained through *NATTENS*im in Fig. 3, which is on the  
 203 HunyuanVideo use case with approximately 90% sparsity. Unlike most others, in this particular  
 204 use case we have many choices with varying upper-bound speedups, any of which may offer the  
 205 best accuracy-efficiency trade-off if further fine-tuned or trained with. We also report the actual  
 206 performance of our Blackwell kernel on the three models with 90% sparsity, and compare against the  
 207 analytical upper bound reported by *NATTENS*im in Fig. 4. Two observations can immediately be

Table 1: Workload distribution with respect to self attention in Cosmos-7B, HunyuanVideo, and FLUX.1-dev. Measurements were done on a single B200 and without any additional performance optimizations, and using the original FP16/BF16 precision.

Use case	% SA in DiT	% DiT in E2E workload	% SA in E2E workload
Cosmos-7B	58.7%	> 99%	58.7%
HunyuanVideo	65.4%	92.8%	60.7%
FLUX.1-dev (4K)	56.8%	91.2%	51.8%

Table 2: FLUX end-to-end speedups from GNA under different settings. Analytical speedups based on FLOPs and *NATTENS*im, and actual speedups from our Blackwell kernel are reported.

Window Size	Stride	# SA Steps	E2E Speedup ↑		Actual
			Analytical FLOP-wise	<i>NATTENS</i> im	
80 × 80	1 × 1	0	1.88×	1.76×	1.65×
80 × 80	16 × 1	0	1.88×	1.84×	1.72×
80 × 80	16 × 16	0	1.88×	1.88×	1.82×
80 × 80	1 × 1	9	1.46×	1.42×	1.37×
80 × 80	16 × 1	9	1.46×	1.45×	1.40×
80 × 80	16 × 16	9	1.46×	1.46×	1.45×

Table 3: Cosmos-7B end-to-end speedups from GNA under different settings. Analytical speedups based on FLOPs and *NATTENS*im, and actual speedups from our Blackwell kernel are reported.

Window Size	Stride	# SA Steps	E2E Speedup $\uparrow$		Actual
			Analytical FLOP-wise	$\mathcal{N}$ ASim	
56% sparsity					
$16 \times 32 \times 48$	$1 \times 1 \times 1$	0	1.50x	1.35x	1.18x
$16 \times 32 \times 48$	$1 \times 8 \times 1$	0	1.50x	1.42x	1.20x
$16 \times 32 \times 48$	$1 \times 1 \times 16$	0	1.50x	1.42x	1.21x
$16 \times 32 \times 48$	$1 \times 8 \times 16$	0	1.50x	1.50x	1.46x
$16 \times 32 \times 48$	$1 \times 1 \times 1$	12	1.28x	1.21x	1.11x
$16 \times 32 \times 48$	$1 \times 8 \times 1$	12	1.28x	1.24x	1.12x
$16 \times 32 \times 48$	$1 \times 1 \times 16$	12	1.28x	1.24x	1.13x
$16 \times 32 \times 48$	$1 \times 8 \times 16$	12	1.28x	1.28x	1.26x
89% sparsity					
$16 \times 24 \times 16$	$1 \times 1 \times 1$	0	2.10x	1.90x	1.76x
$16 \times 24 \times 16$	$1 \times 8 \times 1$	0	2.10x	1.96x	1.79x
$16 \times 24 \times 16$	$1 \times 1 \times 16$	0	2.10x	2.05x	1.88x
$16 \times 24 \times 16$	$1 \times 8 \times 16$	0	2.10x	2.10x	2.05x
$16 \times 24 \times 16$	$1 \times 1 \times 1$	12	1.52x	1.45x	1.40x
$16 \times 24 \times 16$	$1 \times 8 \times 1$	12	1.52x	1.48x	1.40x
$16 \times 24 \times 16$	$1 \times 1 \times 16$	12	1.52x	1.51x	1.44x
$16 \times 24 \times 16$	$1 \times 8 \times 16$	12	1.52x	1.52x	1.50x

Table 4: HunyuanVideo end-to-end speedups from GNA under different settings. Analytical speedups based on FLOPs and *NATTENS*im, and actual speedups from our Blackwell kernel are reported.

Window Size	Stride	# SA Steps	E2E Speedup $\uparrow$		Actual
			Analytical FLOP-wise	$\mathcal{N}$ /ASim	
58% sparsity					
30 $\times$ 40 $\times$ 40	1 $\times$ 1 $\times$ 1	0	1.55 $\times$	1.21 $\times$	1.15 $\times$
30 $\times$ 40 $\times$ 40	1 $\times$ 1 $\times$ 8	0	1.55 $\times$	1.44 $\times$	1.26 $\times$
30 $\times$ 40 $\times$ 40	1 $\times$ 32 $\times$ 8	0	1.55 $\times$	1.55 $\times$	1.53 $\times$
30 $\times$ 40 $\times$ 40	1 $\times$ 1 $\times$ 1	15	1.33 $\times$	1.14 $\times$	1.08 $\times$
30 $\times$ 40 $\times$ 40	1 $\times$ 1 $\times$ 8	15	1.33 $\times$	1.27 $\times$	1.15 $\times$
30 $\times$ 40 $\times$ 40	1 $\times$ 32 $\times$ 8	15	1.33 $\times$	1.33 $\times$	1.30 $\times$
91% sparsity					
18 $\times$ 24 $\times$ 24	1 $\times$ 1 $\times$ 1	0	2.23 $\times$	1.73 $\times$	<b>1.73<math>\times</math></b>
18 $\times$ 24 $\times$ 24	2 $\times$ 1 $\times$ 1	0	2.23 $\times$	1.78 $\times$	1.77 $\times$
18 $\times$ 24 $\times$ 24	1 $\times$ 1 $\times$ 8	0	2.23 $\times$	1.99 $\times$	1.95 $\times$
18 $\times$ 24 $\times$ 24	2 $\times$ 1 $\times$ 8	0	2.23 $\times$	2.02 $\times$	1.98 $\times$
18 $\times$ 24 $\times$ 24	1 $\times$ 8 $\times$ 8	0	2.23 $\times$	2.17 $\times$	2.09 $\times$
18 $\times$ 24 $\times$ 24	2 $\times$ 8 $\times$ 8	0	2.23 $\times$	2.20 $\times$	2.11 $\times$
18 $\times$ 24 $\times$ 24	16 $\times$ 8 $\times$ 8	0	2.23 $\times$	2.23 $\times$	<b>2.23<math>\times</math></b>
18 $\times$ 24 $\times$ 24	1 $\times$ 1 $\times$ 1	15	1.63 $\times$	1.42 $\times$	<b>1.42<math>\times</math></b>
18 $\times$ 24 $\times$ 24	2 $\times$ 1 $\times$ 1	15	1.63 $\times$	1.44 $\times$	<b>1.44<math>\times</math></b>
18 $\times$ 24 $\times$ 24	1 $\times$ 1 $\times$ 8	15	1.63 $\times$	1.53 $\times$	1.51 $\times$
18 $\times$ 24 $\times$ 24	2 $\times$ 1 $\times$ 8	15	1.63 $\times$	1.55 $\times$	1.54 $\times$
18 $\times$ 24 $\times$ 24	1 $\times$ 8 $\times$ 8	15	1.63 $\times$	1.61 $\times$	1.58 $\times$
18 $\times$ 24 $\times$ 24	2 $\times$ 8 $\times$ 8	15	1.63 $\times$	1.62 $\times$	1.59 $\times$
18 $\times$ 24 $\times$ 24	16 $\times$ 8 $\times$ 8	15	1.63 $\times$	1.63 $\times$	<b>1.63<math>\times</math></b>

made: 1. In perfectly block-sparse cases, our kernel either fully or almost fully realizes the analytical speedup, which is also the FLOP-wise speedup (maximum achievable). This comes at no surprise, since the use of the fine-grained masking in the softmax stage is conditioned on whether or not it is required. 2. In the case of Hunyuan, even the standard Neighborhood Attention comes very close to realizing its full analytical speedup. However, the cost of fine-grained masking can still bear some non-negligible overhead in cases that are not perfectly block-sparse. We report end-to-end measures of speedup in Tabs. 2 to 4. In addition, we present qualitative results, as well as some quantitative benchmarks on these use cases in the next two subsections.

## 4.2 Video Generation

We conduct experiments on the aforementioned video generation models: Cosmos-7B [1] and HunyuanVideo [27]. Both use an isotropic Transformer architecture (DiT [34] and MMDiT [17] respectively) with full self attention. In both cases, we generate 5 seconds of video at 720p. In the case of Cosmos-7B, the token layout (DiT feature map shape) is  $16 \times 44 \times 80$ , for which we use GNA with a window size of  $16 \times 32 \times 48$ , which is approximately 56.4% sparsity. We found that going beyond this level of sparsity without further training/fine-tuning we cannot maintain visual quality. We tried both the best-performing GNA configuration (stride  $1 \times 8 \times 16$ ), and standard NA (stride  $1 \times 1 \times 1$ ). Since the initial diffusion steps have a significant impact on the overall structure of the generated video, we retain self attention for the first 12 diffusion steps and use GNA for the

Table 5: HunyuanVideo VBench performance across different GNA configurations. Following STA [54], these configurations do not have any self attention steps. We report analytical speedups based on FLOPs, as well as on  $\mathcal{N}ATTENSim$ . We also report actual speedup based on GNA runtimes, using our Blackwell kernel. We do not report STA speedups since their implementation is limited to the Hopper architecture, and cannot be run on the Blackwell architecture. Note that with a speedup of 2.23 $\times$  in the 91% sparsity case, **GNA achieves the maximum speedup theoretically possible** for this level of sparsity (with respect to reduction in FLOPs).

Method	Window Size	Stride	Attention Sparsity	VBench			Runtime ↓ on B200 (seconds)	End-to-End Speedup ↑		
				Total ↑	Quality ↑	Semantic ↑		FLOP-wise	$\mathcal{N}ASim$	Actual
SA (baseline)	-	-	0.0%	83.08%	85.01%	75.35%	628			
<i>As reported in STA [54]</i>										
Tiled NATTEN (= NA)	30 × 41 × 41	-	58.3%	82.69%	84.61%	75.00%	-	1.55×	-	-
Swin (= WSA)	30 × 40 × 40	-	58.3%	77.53%	78.84%	72.28%	-	1.55×	-	-
STA	30 × 40 × 40	-	58.3%	82.46%	84.63%	73.83%	-	1.55×	-	-
STA	18 × 24 × 24	-	91.0%	80.58%	81.47%	77.03%	-	2.23×	-	-
STA w/ training	18 × 24 × 24	-	91.0%	82.62%	84.76%	74.05%	-	2.23×	-	-
<i>Ours</i>										
GNA (= NA)	30 × 40 × 40	1 × 1 × 1	58.3%	83.24%	84.70%	77.42%	546	1.55×	1.21×	1.15×
GNA (= WSA)	30 × 40 × 40	30 × 40 × 40	58.3%	82.25%	83.11%	78.83%	491	1.55×	1.44×	1.28×
GNA	30 × 40 × 40	1 × 32 × 8	58.3%	83.40%	84.77%	77.92%	411	1.55×	1.55×	1.53×
GNA (= NA)	18 × 24 × 24	1 × 1 × 1	91.0%	82.36%	83.03%	79.68%	359	2.23×	1.73×	<b>1.73×</b>
GNA (= WSA)	18 × 24 × 24	18 × 24 × 24	91.0%	80.25%	80.97%	77.38%	314	2.23×	2.06×	2.00×
GNA	18 × 24 × 24	16 × 8 × 8	91.0%	82.04%	82.62%	79.69%	277	2.23×	2.23×	<b>2.23×</b>

remaining 23 diffusion steps. Because of this, GNA’s share of the end-to-end workload decreases to approximately 39%, which in turn further limits achievable end-to-end speedup to approximately 1.64 $\times$ . We present sample frames from some of the generated videos in Fig. 6, where we observe GNA can produce videos of comparable quality to that of the self attention baseline.

In HunyuanVideo, the token layout (MMDiT feature map shape) is 30 × 48 × 80 for which we use GNA with a window size of 18 × 24 × 24, which is approximately 91% sparsity, and a stride of 16 × 8 × 8. Similar to Cosmos, we retain self attention for some of the initial diffusion steps, which in this case we set to 15, and run GNA for the remaining 35 steps. This decreases GNA’s share of the end-to-end workload to approximately 42%, which in turn limits achievable end-to-end speedup to approximately 1.72 $\times$ . We again try both the best-performing GNA configuration (stride 16 × 8 × 8), and standard NA (stride 1 × 1 × 1). We present sample frames from output videos in Fig. 7.

We additionally evaluate GNA with HunyuanVideo on VBench [22]. In this case, we follow STA [54] by applying sparsity to all 50 steps. However, we only report results without any additional training or fine tuning. In Tab. 5, NA and GNA with 58.3% sparsity achieve VBench scores of 83.24% and 83.40% respectively, both of which are comparable to the self attention baseline at 83.08%. However, they can only achieve end-to-end speedups of 1.11 $\times$  and 1.23 $\times$  respectively. When using higher sparsity of 91%, the VBench score drops moderately to 82.36% and 82.04%, but with more significant end-to-end speedups of 1.73 $\times$  and 2.27 $\times$  respectively.

### 4.3 Image Generation

Table 6: 4K image generation results from FLUX-1.dev [28, 51] on MAN-IQA [50], QualiCLIP [2], and GenEval [18]. All GNA configurations retain self attention in the first 9 (out of 28) diffusion steps, and use window size 80 × 80, which is approximately 90% sparsity. Similar to Tab. 5, we report generation runtime on the B200, analytical speedup based on both FLOPs and  $\mathcal{N}ATTENSim$ , as well as actual speedups.

Configuration	HPD ↑		Overall	Single Object	Two Objects	GenEval ↑				Runtime on B200 ↓ (seconds)	E2E Speedup ↑		
	MAN-IQA	QualiCLIP				Counting	Colors	Position	Color Attr.		FLOP-wise	$\mathcal{N}ASim$	Actual
SA (baseline)	<b>0.3718</b>	0.4235	<b>0.5750</b>	0.9594	0.6313	<b>0.5500</b>	<b>0.7793</b>	0.1450	0.3850	129			
<i>GNA with window size = 80 × 80 (90% sparsity)</i>													
s=1 × 1	0.3467	0.4243	0.5743	0.9500	<b>0.6742</b>	0.5031	0.7686	<b>0.1500</b>	0.4000	94	1.46×	1.42×	1.37×
s=16 × 1	0.3467	0.4243	0.5742	0.9625	0.6607	0.4873	0.7739	0.1480	0.4125	92	1.46×	1.45×	1.40×
s=16 × 16	0.3462	<b>0.4247</b>	<b>0.5785</b>	<b>0.9656</b>	<b>0.6717</b>	0.4969	<b>0.7819</b>	0.1400	<b>0.4150</b>	89	1.46×	1.46×	1.45×

We conduct our experiments with image generation on FLUX-1.dev, which similar to the video models uses an isotropic Transformer architecture (MMDiT [17]) with full self attention. We only study generating 4K images, which result in 256 × 256 feature maps, as smaller resolution workloads

aren't as limited by the cost of self attention. For generating 4K images, we use adapters from URAE [51], as FLUX-1.dev does not natively support 4K image generation. We experiment with a window size of  $80 \times 80$  (90.2% sparsity), and strides  $1 \times 1$ ,  $16 \times 1$ ,  $16 \times 16$ , which were the optimal choices from *NATTENS*im. Stride  $1 \times 1$  is standard neighborhood attention, and stride  $16 \times 16$  is perfectly block-sparse, which can realize the full 10.2 $\times$  analytical op-level speedup. For preserving quality and global structure, we follow our video generation experiments, in which we retain self attention in the first few diffusion steps, which in this case we set to the first 9 out of the 28 diffusion steps. This shrinks the share of GNA in the end-to-end workload to approximately 35%, and by extension the upper-bound end-to-end improvement to 1.46 $\times$ . However, we successfully realize almost all of that with a speedup of 1.45 $\times$  over the self attention baseline. In terms of quality, we observe very few notable differences when using GNA, even with a high stride of  $16 \times 16$ , as illustrated in Fig. 8. We also present some quantitative metrics for these configurations. Following URAE [51], we evaluate them on MAN-IQA [50] and QualiCLIP [2] by generating images with prompts from the HPDv2 [47] test dataset. We additionally evaluate text-to-image alignment using GenEval [18], which is an object-focused evaluation benchmark. This benchmark is comprised of various categories, each responsible for measuring a different compositional property, such as color, positioning, attribute binding, object count, and the like. We report the results of all three benchmarks in Tab. 6. In summary, we observe FLUX-1.dev with GNA to be on par with the self attention baseline in terms of quality, while offering up to 45% end-to-end speedup on the B200.

## 5 Conclusion

Sliding window and blocked attention are two extremes of locality-focused sparse attention methods, with the former enjoying inductive biases such as translational equivariance [33, 37, 43, 21] and arguably potential for higher quality and expressiveness, and the latter potentially more efficient and trivial to implement.

In this paper, we extend the neighborhood attention family of patterns, which were already flexible in terms of choice for window size, dilation, and causal masking, and add a new “stride” parameter. Just as in convolution, this parameter adds a delay step to the sliding window effect, allowing it to implement many existing sparse attention methods. Those include, but are not limited to, blocked local attention [43], blocked / window self attention [31], and sliding tile attention [54], in addition to the existing sliding window [33, 37] and neighborhood attention [21] methods. While the concept of a delay step in sliding window attention is by no means new [43], we revisit it for a different reason compared to the original work. While Vaswani et al. [43] were concerned with the cost of explicit memory operations, the focus of this work is maximizing speedup from these methods to the point of being proportional to the level of sparsity (i.e. 10 $\times$  speedup from 90% sparsity).

We created an analytical model for GNA, called *NATTENS*im, which can simulate tiling behavior under various design choices, and compute the number of KV tiles visited per query tile, through which we compute more reliable upper-bound speedups for different configurations under GNA. These measures help quantitatively compare the upper-bound performance of different approaches without being biased by differences in implementation. In addition, we find that many GNA configurations exist that are perfectly block-sparse, under which speedup proportional to FLOPs is possible, and they do not necessarily fit the definition of either blocked attention or sliding tile attention. We further implement GNA on top of CUTLASS FMHA for the Blackwell architecture, which can achieve an effective 1.2 petaFLOPs/s with FP16 and 1.7 petaFLOPs/s with FP8, and show that specifically in the case of perfectly block-sparse configurations it can fully realize the analytical speedup computed by *NATTENS*im. We also highlight 3 potential applications for GNA, all of which we confirm spend the majority of their workload on self attention, and report end-to-end speedups close to or matching the expected FLOP-wise speedup. On Cosmos-7B [1] and with 56% sparsity introduced into 23 of 35 diffusion steps, we achieve a speedup of 1.26 $\times$ , with FLOP-wise speedup being 1.28 $\times$ . On HunyuanVideo [27], and with 91% sparsity introduced into 35 of 50 diffusion steps, we fully realize the FLOP-wise speedup of 1.63 $\times$ . On FLUX-1.dev [28], and with 90% sparsity introduced into 19 of 28 diffusion steps, we achieve a speedup of 1.45 $\times$ , with FLOP-wise speedup being 1.46 $\times$ . All three of the aforementioned configurations can still produce visually acceptable outputs, without any further training or fine-tuning. We hope that Generalized Neighborhood Attention can serve as a recipe for Speed-of-Light local attention beyond the use cases and hardware studied in this paper.

## References

- [1] Niket Agarwal, Arslan Ali, Maciej Bala, Yogesh Balaji, Erik Barker, Tiffany Cai, Prithvijit Chattopadhyay, Yongxin Chen, Yin Cui, Yifan Ding, et al. Cosmos world foundation model platform for physical ai. *arXiv preprint arXiv:2501.03575*, 2025.
- [2] Lorenzo Agnolucci, Leonardo Galteri, and Marco Bertini. Quality-aware image-text alignment for real-world image quality assessment. *arXiv preprint arXiv:2403.11176*, 5(6), 2024.
- [3] Meta AI. The llama 4 herd: The beginning of a new era of natively multimodal ai innovation, 4 2025. Accessed: 2025-04-16.
- [4] Joshua Ainslie, James Lee-Thorp, Michiel de Jong, Yury Zemlyanskiy, Federico Lebron, and Sumit Sanghai. Gqa: Training generalized multi-query transformer models from multi-head checkpoints. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 2023.
- [5] Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020.
- [6] Daniel Bolya, Cheng-Yang Fu, Xiaoliang Dai, Peizhao Zhang, Christoph Feichtenhofer, and Judy Hoffman. Token merging: Your vit but faster. In *ICLR*, 2023.
- [7] Beidi Chen, Tri Dao, Eric Winsor, Zhao Song, Atri Rudra, and Christopher Ré. Scatterbrain: Unifying sparse and low-rank attention. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [8] Zhaodong Chen, Zheng Qu, Yuying Quan, Liu Liu, Yufei Ding, and Yuan Xie. Dynamic n: M fine-grained structured sparse attention mechanism. In *Proceedings of the 28th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming*, 2023.
- [9] Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*, 2019.
- [10] Team Cohere, Arash Ahmadian, Marwan Ahmed, Jay Alammari, Yazeed Alnumay, Sophia Althammer, Arkady Arkhangorodsky, Viraat Aryabumi, Dennis Aumiller, Raphaël Avalos, et al. Command a: An enterprise-ready large language model. *arXiv preprint arXiv:2504.00698*, 2025.
- [11] Katherine Crowson, Stefan Andreas Baumann, Alex Birch, Tanishq Mathew Abraham, Daniel Z Kaplan, and Enrico Shippole. Scalable high-resolution pixel-space image synthesis with hourglass diffusion transformers. In *International Conference on Machine Learning (ICML)*, 2024.
- [12] Tri Dao. Flashattention-2: Faster attention with better parallelism and work partitioning. In *International Conference on Learning Representations (ICLR)*, 2023.
- [13] Tri Dao, Daniel Y Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. Flashattention: Fast and memory-efficient exact attention with io-awareness. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [14] Juechu Dong, Boyuan Feng, Driss Guessous, Yanbo Liang, and Horace He. Flex attention: A programming model for generating optimized attention kernels. *arXiv preprint arXiv:2412.05496*, 2024.
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Machine Learning (ICML)*, 2020.
- [16] Haoxing Du, Lyna Kim, Joan Creus-Costa, Jack Michaels, Anuj Shetty, Todd Hutchinson, Christopher Riedel, and John Dean. Weathermesh-3: Fast and accurate operational global weather forecasting. *arXiv preprint arXiv:2503.22235*, 2025.
- [17] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *International Conference on Machine Learning (ICML)*, 2024.
- [18] Dhruva Ghosh, Hannaneh Hajishirzi, and Ludwig Schmidt. Geneval: An object-focused framework for evaluating text-to-image alignment. *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- [19] Ali Hassani, Wen-Mei Hwu, and Humphrey Shi. Faster neighborhood attention: Reducing the  $O(n^2)$  cost of self attention at the threadblock level. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- [20] Ali Hassani and Humphrey Shi. Dilated neighborhood attention transformer. *arXiv preprint arXiv:2209.15001*, 2022.
- [21] Ali Hassani, Steven Walton, Jiachen Li, Shen Li, and Humphrey Shi. Neighborhood attention transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [22] Ziqi Huang, Yinan He, Jiashuo Yu, Fan Zhang, Chenyang Si, Yuming Jiang, Yuanhan Zhang, Tianxing Wu, Qingyang Jin, Nattapol Chanpaisit, Yaohui Wang, Xinyuan Chen, Limin Wang, Dahua Lin, Yu Qiao, and Ziwei Liu. VBench: Comprehensive benchmark suite for video generative models. In *CVPR*, 2024.
- [23] Jitesh Jain, Jiachen Li, Mang Tik Chiu, Ali Hassani, Nikita Orlov, and Humphrey Shi. Oneformer: One transformer to rule universal image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [24] Hanhwi Jang, Joonsung Kim, Jae-Eon Jo, Jaewon Lee, and Jangwoo Kim. Mnnfast: A fast and scalable system architecture for memory-augmented neural networks. In *Proceedings of the 46th International Symposium on Computer Architecture*, 2019.



- [25] Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. Mistral 7b. *arXiv preprint arXiv:2310.06825*, 2023.
- [26] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023.
- [27] Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, et al. Hunyuanvideo: A systematic framework for large video generative models. *arXiv preprint arXiv:2412.03603*, 2024.
- [28] Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024.
- [29] Yanghao Li, Hanzi Mao, Ross Girshick, and Kaiming He. Exploring plain vision transformer backbones for object detection. In *European Conference on Computer Vision (ECCV)*, 2022.
- [30] Aixin Liu, Bei Feng, Bin Wang, Bingxuan Wang, Bo Liu, Chenggang Zhao, Chengqi Deng, Chong Ruan, Damai Dai, Daya Guo, et al. Deepseek-v2: A strong, economical, and efficient mixture-of-experts language model. *arXiv preprint arXiv:2405.04434*, 2024.
- [31] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10012–10022, 2021.
- [32] Maxim Milakov and Natalia Gimelshein. Online normalizer calculation for softmax. *arXiv preprint arXiv:1805.02867*, 2018.
- [33] Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. In *International Conference on Machine Learning (ICML)*, 2018.
- [34] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023.
- [35] Jiezhong Qiu, Hao Ma, Omer Levy, Wen-tau Yih, Sinong Wang, and Jie Tang. Blockwise self-attention for long document understanding. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, 2020.
- [36] Markus N Rabe and Charles Staats. Self-attention does not need  $O(n^2)$  memory. *arXiv preprint arXiv:2112.05682*, 2021.
- [37] Prajit Ramachandran, Niki Parmar, Ashish Vaswani, Irwan Bello, Anselm Levskaya, and Jon Shlens. Stand-alone self-attention in vision models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [38] Aurko Roy, Mohammad Saffar, Ashish Vaswani, and David Grangier. Efficient content-based sparse attention with routing transformers. *Transactions of the Association for Computational Linguistics*, 2021.
- [39] Jay Shah, Ganesh Bikshandi, Ying Zhang, Vijay Thakkar, Pradeep Ramani, and Tri Dao. Flashattention-3: Fast and accurate attention with asynchrony and low-precision. *arXiv preprint arXiv:2407.08608*, 2024.
- [40] Noam Shazeer. Fast transformer decoding: One write-head is all you need. *arXiv preprint arXiv:1911.02150*, 2019.
- [41] Vijay Thakkar, Pradeep Ramani, Cris Cecka, Aniket Shivam, Honghao Lu, Ethan Yan, Jack Kosaian, Mark Hoemmen, Haicheng Wu, Andrew Kerr, Andrew Nicely, Duane Merrill, Dustyn Blasig, Fengqi Qiao, Piotr Majcher, Paul Springer, Markus Hohnerbach, Jin Wang, and Manish Gupta. Cutlass.
- [42] Philippe Tillet, Hsiang-Tsung Kung, and David Cox. Triton: an intermediate language and compiler for tiled neural network computations. In *Proceedings of the 3rd ACM SIGPLAN International Workshop on Machine Learning and Programming Languages*, 2019.
- [43] Ashish Vaswani, Prajit Ramachandran, Aravind Srinivas, Niki Parmar, Blake Hechtman, and Jonathon Shlens. Scaling local self-attention for parameter efficient visual backbones. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [45] Steven Walton, Ali Hassani, Xingqian Xu, Zhangyang Wang, and Humphrey Shi. Stylenat: Giving each head a new perspective. *arXiv preprint arXiv:2211.05770*, 2022.
- [46] Sinong Wang, Belinda Z Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020.
- [47] Xiaoshi Wu, Yiming Hao, Keqiang Sun, Yixiong Chen, Feng Zhu, Rui Zhao, and Hongsheng Li. Human preference score v2: A solid benchmark for evaluating human preferences of text-to-image synthesis. *arXiv preprint arXiv:2306.09341*, 2023.
- [48] Haocheng Xi, Shuo Yang, Yilong Zhao, Chenfeng Xu, Muyang Li, Xiuyu Li, Yujun Lin, Han Cai, Jintao Zhang, Dacheng Li, et al. Sparse videogen: Accelerating video diffusion transformers with spatial-temporal sparsity. *arXiv preprint arXiv:2502.01776*, 2025.
- [49] Ruyi Xu, Guangxuan Xiao, Haofeng Huang, Junxian Guo, and Song Han. Xattention: Block sparse attention with antidiagonal scoring. *arXiv preprint arXiv:2503.16428*, 2025.

- 427 [50] Sidi Yang, Tianhe Wu, Shuwei Shi, Shanshan Lao, Yuan Gong, Mingdeng Cao, Jiahao Wang, and  
 428 Yujiu Yang. Maniqa: Multi-dimension attention network for no-reference image quality assessment. In  
 429 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1191–1200,  
 430 2022.
- 431 [51] Ruonan Yu, Songhua Liu, Zhenxiong Tan, and Xinchao Wang. Ultra-resolution adaptation with ease. *arXiv*  
 432 *preprint arXiv:2503.16322*, 2025.
- 433 [52] Jingyang Yuan, Huazuo Gao, Damai Dai, Junyu Luo, Liang Zhao, Zhengyan Zhang, Zhenda Xie, YX Wei,  
 434 Lean Wang, Zhiping Xiao, et al. Native sparse attention: Hardware-aligned and natively trainable sparse  
 435 attention. *arXiv preprint arXiv:2502.11089*, 2025.
- 436 [53] Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon,  
 437 Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. Big bird: Transformers for longer sequences. In  
 438 *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- 439 [54] Peiyuan Zhang, Yongqi Chen, Runlong Su, Hangliang Ding, Ion Stoica, Zhenghong Liu, and Hao Zhang.  
 440 Fast video generation with sliding tile attention. *arXiv preprint arXiv:2502.04507*, 2025.

## A Extended Related Works

Attention has been widely called out for being a computationally intensive operation, as a result of its quadratic complexity with respect to number of tokens. This has often been one of the motivations behind research into sparse forms of attention. Dot-product attention, the most widely used form of attention, consists primarily of two matrix multiplications, which means it can enjoy many forms of sparsity which exist for matrix multiplies, such as structured sparsity [8], and low-rank approximations [46, 7]. Sparsity can also be introduced into attention by choosing coarser targets of sparsity, which is sometimes application-specific. For example, some approaches are designed specifically for introducing sparsity into LLM inference workloads, where KV caching is a necessity, and therefore calls for some form of sparsity or compression [40, 4, 30]. Approaches such as Token Merging [6] attempt to directly reduce the number of both query and context tokens together, and have been shown to be effective in certain vision tasks.

Sparse Attention approaches can also be classified into static and dynamic approaches. Dynamic approaches [38, 6, 8] can be more effective without further fine-tuning, whereas static approaches [33, 7, 37, 9, 5] are more likely to achieve better speedup [25, 10, 31]. Some approaches can be classified as hybrids [49, 6, 48], where statistics gathered at runtime guide re-layout of computation, and then use static methods such as block-sparsity [35]. Some may even use entirely dense computations [31, 6]. Some, such as Native Sparse Attention [52], combine multiple sparse approaches, both static and dynamic.

In this paper, we specifically focus on static methods, where the target of sparsity is the attention weight matrix itself, and therefore the token coordinate space determines whether or not a region / weight is masked. Local attention is the most prominent example, but variants of local attention where global context is introduced through non-contiguous context windows [5, 20] also fall into this category. The success of local approaches [31, 26, 29, 23, 11, 16] can be attributed to the fact that they require little to no change to the Transformer [44] architecture, as well as spatial locality being a bias that exists in nature. Past [25] and recent [10, 3] large language models also adopt local methods into their architecture in addition to standard self attention.

### A.1 (Locally) Sparse Attention

Over the years, many proposed introducing locality into attention, some for the inductive biases, some as a means to achieve subquadratic complexity, and some for both. In this section, we summarize these approaches into three categories, which are described below.

**Sliding Window Attention.** Some of the earliest works proposing this were Image Transformer [33] and Stand-Alone Self-Attention (SASA) [37], where they specifically designed a form of attention sharing some of the inductive biases present in convolution. The sliding window pattern itself is quite similar to, and sometimes inspired by, convolution. However, they did not become widely adopted at first, particularly in vision, due to a lack of infrastructure. Ramachandran et al. [37] stated that while their model achieved superior accuracy to a comparable CNN, it required more resources and time to be trained. The same concept, but in a single-dimensional causal domain, was later used in language, in works such as Longformer [5], and extended to a “dilated” form, which can introduce global context with the same complexity as standard local attention. Years later, neighborhood attention (NA) [21] revisited this concept for hierarchical vision transformers. The key difference between SASA and NA is in the handling of corner cases in the coordinate space. SASA employed a zero-padded feature map, similar to padding in convolution. NA “reflects” the window back, so that every query is guaranteed to attend to a fixed number of context tokens, regardless of its position in the coordinate space and window size. This behavior also guarantees that NA numerically matches self attention when window size is equal to input size. Similar to Longformer, but again in the context of vision, a dilated form of NA was later introduced [20]. The combination of standard and dilated NA in a hierarchical vision transformer surpassed the original model in accuracy across various tasks [20, 45], without incurring any additional cost.

**Strided Sliding Window Attention.** Also referred to as blocked local attention [43], this approach effectively introduces a delay step into sliding window attention. This was originally motivated by the fact that typical implementations of sliding window attention [33, 37] required explicit memory operations that extract sliding windows from context, which can quickly undo any savings

in computation, in addition to growing the memory footprint. This form of sliding window attention has an extreme case, called Blocked Attention, where stride is as large as the window size, which results in non-overlapping local windows. More recently, this approach was revisited in Sliding Tile Attention (STA) [54], where the objective was achieving perfectly block-sparse masks as a means to minimize masked, and therefore wasted FLOPs in sliding window methods [33, 37, 21]. We further clarify this approach in Appendix A.3.

**Blocked Attention.** This approach effectively partitions the query and context set, and performs self attention on each partition independently (and in parallel) to the rest. In addition to HaloNet [43], Window Self Attention (WSA) from Swin Transformer [31] is another instance of blocked attention. Blocked attention is easy to implement, and embarrassingly parallel along blocks, but this comes at the cost of no cross-block interactions. This can be remedied by introducing global attention, or varying window (block) sizes across layers (Swin’s Shifted WSA), or introducing convolution. This approach has been adopted in many vision models such as Segment Anything [26], and ViTDet [29].

## A.2 (Sparse) Attention Infrastructure

Sliding window attention, specifically in the context of vision, was commonly considered inefficient [43, 31], but that was predicated on the assumption that any form of transformation to context tokens (keys and values) must be explicit in global memory. One of the key contributions of Neighborhood Attention Transformer [21] was a set of naive CUDA kernels that simply computed the vector-matrix multiplication problem without such explicit copies in global memory. These kernels were packaged as a PyTorch extension, called *NATTEN*. However, around the same time as the initial release, implementation of attention in general was about to undergo a massive change.

Until 2022 [36, 13], most implementations of dot-product attention, especially those not tied to specific inference use cases, were implemented with two matrix multiplications, with the softmax operator in between. In most deep learning frameworks, the former typically targets General Matrix-Matrix Multiplication (GEMM) routines in powerful dense linear algebra packages, i.e. cuBLAS and CUTLASS [41] for NVIDIA GPUs, and these routines typically offer great performance, usually up to 80% of the peak FLOPs/second utilization. The issue with this approach however is that the size of the intermediary matrix, the attention weight matrix, grows quadratically, resulting in a quadratic memory footprint, and by extension quadratic number of memory operations in both GEMMs. As a result, this implementation becomes heavily limited by memory bandwidth, which on most modern GPUs is orders of magnitude smaller than computation power (FLOPs/second). Flash Attention [13] showed that by fusing the two GEMMs and softmax into the same kernel, and utilizing online softmax [32, 24], we can continuously accumulate context in SRAM without ever fully realizing the  $\mathcal{O}(n^2)$  weight matrix. This approach is also referred to as Fused Multi-headed Attention (FMHA). Flash Attention 2 [12] later improved upon the original, and is the state-of-the-art for the NVIDIA Ampere architecture. Flash Attention 3 [39] extended the approach to the NVIDIA Hopper architecture by using the new programming model and hardware features through CuTe and CUTLASS [41].

Due to the fact that self attention is the baseline for all sparse attention, and that baseline is improved significantly with FMHA, sparse attention approaches have had to follow suit. Mistral-7B [25] was one of the first models to use 1-D sliding window attention by directly implementing it as a block-sparse mask in state-of-the-art implementations, such as Flash Attention 2 [12]. More recently, language models such as Command A [10] have also adopted this approach, and use sliding window attention together with global self attention.

However, implementations such as *NATTEN* faced additional challenges, due to the additional burden of dealing with multi-dimensional layouts of tokens, which makes efficient block-sparsity for them non-trivial.

## A.3 Curse of Multi-Dimensionality

Studies involving sparsity in attention weights have one key, and often undiscussed, difference between LLM-focused applications and vision applications, which is the multi-dimensional layout of visual data (images, videos). We consider this an additional burden, as design choices available for efficient implementations are limited, and often with noticeable overhead. To illustrate this issue,

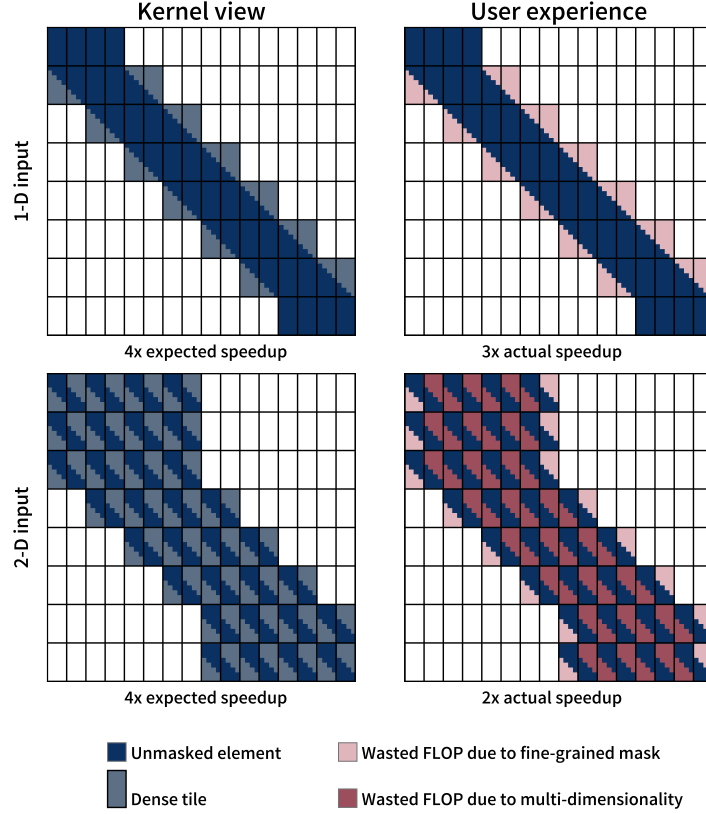


Figure 5: Curse of **multi-dimensionality**: single-dimensional tiling opens up sparsity in multi-dimensional layouts of tokens to more wasted computation (FLOPs that are still computed but masked prior to softmax). While many fine-grained attention masks, even 1-D sliding window attention (top), can still have some FLOPs masked due to the fact that the vector-matrix multiplies are packed into matrix-matrix multiplies, masked FLOPs due to multi-dimensionality can be much more significant (bottom). Note that the single-dimensional case is bi-directional and not causal for better comparison to the multi-dimensional case.

we refer readers to Fig. 5. When considering for instance a 2-dimensional layout of tokens (i.e. a feature map in an image model), if the FMHA implementation employs single-dimensional tiling over the typical layout of these tokens in memory, the potential for “wasted compute” increases, and according to the tile size.

Fused Neighborhood Attention (FNA) [19] proposed solving this by employing multi-dimensional tiling which converts the GEMMs in the FMHA into an instance of tensor contractions. This implementation naturally improved naive kernels in *NATTEN* by a significant margin, but it can also be greatly limited by the overhead of software predication as a direct result of multi-dimensional tiling. While hardware predication, through components like NVIDIA’s Tensor Memory Accelerator (TMA) <sup>3</sup>, can likely minimize this overhead, they can also impact certain design choices critical to achieving optimal speedup, such as dynamic KV tiling.

A more easily implementable alternative to multi-dimensional tiling is simulating the behavior through re-layout of tokens. One of the earliest demonstrations of this was in the implementation of FNA in Flex Attention [14]<sup>4</sup>. Flex Attention is a PyTorch API that evaluates user-defined attention masks, and compiles them into block-sparse Triton kernels. This approach adds a non-avoidable overhead from the additional memory operations required for the re-layout. It also still requires some modification

<sup>3</sup><https://docs.nvidia.com/cuda/hopper-tuning-guide/index.html#tensor-memory-accelerator>

<sup>4</sup>Dubbed “Tiled NATTEN”.

of the original attention kernel, but far fewer changes compared to fused multi-dimensional tiling. In the case of Flex Attention, this can be handled by translating 1-D coordinates into coordinates in the new layout of tokens directly in the user-defined mask. Sliding Tile Attention (STA) [54] attempts to further minimize FLOPs wasted due to fine-grained masking (see Fig. 5), and proposes defining neighborhood / sliding window attention on tiles / blocks instead of individual tokens, with the tile / block size matching that of the underlying FMHA kernel. This closely resembles strided sliding window approaches such as HaloNet [43], but is instead motivated by computational density, and therefore speedup, instead of the cost of memory operations, which are non-existent with block-sparse FMHA kernels. STA’s implementation employs re-layout of tokens from Flex Attention, instead of on-the-fly multi-dimensional tiling, and successfully outperforms FNA given the same FLOP-wise sparsity. It is noteworthy that the implementation is specific to the Hopper architecture, whereas FNA targets Ampere tensor cores. Aside from this, the key limitation of STA is that like Blocked Attention, it assumes query and context tensors are always tiled according to the same tile size. This assumption is not guaranteed to hold in practice, and even if some implementations support it, there is no guarantee that such configurations achieve optimal performance for the given use case.

## B Extended Methodology

### B.1 Analytical tool for GNA

In this appendix, we present more details on *NATTENSIM*, its use cases, and kernel design choices it considers.

**Design choice: tile sizes.** Among the most important configurations for any GEMM or GEMM-based kernel, tile sizes directly determine how the workload is divided up and scheduled among workers. In most FMHA kernels, which fuse together two GEMMs, they are forced to share most of their tile sizes, but with a permutation. In the first GEMM, query tensor (Q) is tiled along the query sequence mode with some tile size  $T_Q$ , key tensor (K) is tiled along the context sequence mode with some tile size  $T_K$ . The shared head dim mode, which is the contracting mode, is tiled by some tile size  $T_D$ . In the second GEMM, the lefthandside operand is a tile of attention weights of shape  $T_Q \times T_K$ , which means the righthandside operand, tile from value tensor (V), must match the tile size along the contracting dimension, which is now the context sequence mode. Therefore, K and V take the same tile size, which we will henceforth refer to as  $T_{KV}$ . The head dim mode of V is also typically tiled by  $T_D$ .

Tile sizes  $T_Q$  and  $T_{KV}$  therefore directly affect the number of FLOPs masked (wasted). Many factors can determine valid choices for these tile sizes, such as the amount of SRAM available, number of stages in the GEMM pipeline, layout transformations required for the operands, and of course shapes of the underlying matrix multiply-and-accumulates (MMAs), or Tensor Core instruction shapes in the case of NVIDIA GPUs. With modern architectures such as Hopper and Blackwell, and in the case of GEMMs, there are many choices for tile sizes available, which exhibit different performance levels on different problem shapes. One key limitation in works such as STA [54] is that the methodology assumes  $T_Q$  is always equivalent to  $T_{KV}$ , thus limiting the number of choices. This is while performant FMHA kernels are not guaranteed to always provide the same level of flexibility when it comes to picking tile sizes as in standard dense GEMMs.

**Design choice: Single-dimensional vs multi-dimensional tiling.** Most FMHA kernels assume a single-dimensional layout of tokens, which given attention’s permutation equivariance is logical. This however creates a challenge for cases where tokens assume a multi-dimensional layout, such as visual models where tokens represent patches of pixels in images and videos. This opens up those applications to more wasted FLOPs, as illustrated in Fig. 5. A natural fix is to tile in multiple dimensions and with respect to the original multi-dimensional layout, which essentially converts the GEMM problem into a tensor contraction (GETT<sup>5</sup>). This solution was employed by FNA [19], where  $T_Q$  and  $T_{KV}$  in the base FMHA kernel were re-interpreted as multi-dimensional tiles (i.e.  $T_Q = 64 \rightarrow T_Q = 8 \times 8$ ). One downside to this is that this can introduce a significant overhead due to additional software predication logic, and given that the kernel was designed for the Ampere architecture, hardware predication was not an option. A practical solution to this is taking multi-

<sup>5</sup>Generalized Tensor-Tensor Contraction (GETT).



dimensional tiling out of the kernel and instead implementing it as a re-layout operation. This solution was employed by Flex Attention [14] in their implementation of multi-dimensional NA masks (referred to as “Tiled NATTEN”). The only potential downside to this approach is the unavoidable fixed cost of memory operations, which is independent of the level of sparsity, and only a function of the size of the Q, K, V, and output tensors. We dub this approach **token permutation**, as it is mainly comprised of a re-layout of the token space, and is agnostic to batch, heads, and head dim.

**Design choice: KV tiling.** Tiling is typically static, and in many block-sparse FMHA kernels, static KV tiles are either visited, or skipped, according to the mask definition. However, some implementations, such as FNA [19], first slice out the region in the KV token space that would be required to be visited, and dynamically tile the region. This can save some additional computation, and result in minimal wasted FLOPs possible. On the other hand, this approach is not always realizable, especially in designs that rely on hardware predication. For instance, the Hopper Tensor Memory Accelerator (TMA) requires determining the parameters of tiled copies (tensor maps) prior to kernel launch. While on-the-fly modification/replacement of tensor maps is possible, it is not without overhead.

**Use cases.** Problem shape (layout of tokens), window size, dilation, stride, and causal masking are all user-specified, and play a role in determining computational savings in terms of tiles.

**Tiling simulation.** *NATTEN*Sim’s primary goal is to simulate how a given use case is tiled according to design choices in the implementation. Through basic operations on coordinates, and by using an exact definition of the core GNA mask, *NATTEN*Sim computes the coordinates of each KV tile visited by each Q tile. If we consider the worst case of all Q tiles (maximum number of KV tiles), we can compute a more realistic and fine-grained upper bound speedup than FLOP-wise speedup, with respect to self attention for each use case. For instance, a perfectly blocks-sparse mask and sliding window attention mask with 90% sparsity both have a FLOP-wise upper-bound speedup of  $10 \times (\frac{1}{1-90\%})$ , but the latter never get away with performing exactly  $\frac{1}{10}$  of the FLOPs unless implemented as a vector-matrix multiplication.

## B.2 Implementation for the Blackwell architecture

Our Blackwell FNA kernel closely follows the original FNA kernel [19], with the exception of taking multi-dimensional tiling outside the kernel, and instead relying on token permutation. This also forces static KV tiling instead of dynamic KV tiling in the original FNA. We chose this design for the following reasons:

1. If using the TMA for data movement, static KV tiling is the only choice.
2. Fusing multi-dimensional tiling into existing FMHA kernels can break too many assumptions made with respect to the sequence mode, even with CuTe facilitating layout transformations and interfacing with the TMA.
3. We are not leaving much performance on the table, as long as we are not limited by the memory transfer time from token permutation and reverse permutation. For example, considering use cases from Cosmos-7B [1] and HunyuanVideo [27], this would be only 6.9 and 10.5GB respectively. If we utilize even half of the 8TB/s HBM bandwidth of a single B200, this would be 1.9% and 1% of the FMHA time, which would only limit very large levels of sparsity.
4. In many cases, token permutation and reverse permutation can be done only once: permute before the first transformer layer, and reverse after the last layer. Most transformer architectures are equivariant to permutation, and this holds true for both ViT [15] and DiT [34], which are prominent in vision.
5. Fusing additional context tokens is trivial with token permutation. Certain models, such as Hunyuan [27] and Flux [28] cross attend visual tokens with text tokens. *NATTEN* has usually supported those scenarios by launching an additional FMHA kernel, and “merging” the two outputs using their logsumexp. However, we implement this feature within the same kernel, allowing some KV tokens to take a completely different layout and mask.

Since we implement multi-dimensional tiling, spatial locality allows us to define visited KV tiles as the range between the last KV tile coordinate required by the last query in the Q tile, and the first KV tile coordinate required by the first query in the Q tile. Most of the kernel remains agnostic to the multi-dimensional tiling, and only the producer and softmax warps take this into account. Producer warp simply maps iteration index to the multi-dimensional tile coordinate, and then to the static 1-D tile index, which directly interfaces with the indexing of TMA loads through CuTe. Softmax warp(s) likewise have to map the 1-D Q and KV coordinates back into the original multi-dimensional coordinates, and apply the fine-grained GNA mask. However, given the overhead of masking, we also implement predicates for perfectly block-sparse cases, and for any additional KV tokens, which can in some cases completely eliminate our implementation overhead.

As previously stated, in the current iteration, we implement token permutation as a copy operation through PyTorch directly. Problem shapes that are not evenly divisible are manually padded, output tensor is cropped after reverse token permutation, and kernel handles predication for KV padding. In future versions, we hope to also improve upon token permutation, as the current solution typically utilizes approximately 1/8th of the memory bandwidth. The kernel and the additional memory operations, and potential padding, is directly integrated into *NATTEN*, and exposed via the typical `na{1, 2, 3}d` API.

## C Future work

Our new implementation can be further optimized by predicating fine-grained masking further in the event that fully dense tiles exist in settings that are not perfectly block-sparse. Since multi-dimensional tiling preserves spatial locality, if the intersection between the neighborhoods of the first and last query in the Q tile spans entire KV tiles, fine-grained masking can be skipped. Our current implementation does this for perfectly block-sparse configurations. In addition to this, it is possible to performance-optimize the fine-grained mask logic itself. These optimizations can further close the gap between the analytical speedup expected, and the actual speedup achieved in cases that are only partially block-sparse, or not at all.

Token permutation and reverse permutation are also implemented naively with PyTorch, and barely utilize 1/8th of the B200’s peak memory bandwidth, and can be further optimized with specialized copy kernels and activation memory management. However, as noted earlier in the paper, the number of calls to these operations can be greatly reduced in certain architectures, and with certain assumptions (i.e. isotropic Transformer architecture, and no dilation in GNA).

Other extensions can include transferring this design to other SOTA implementations for earlier architectures (i.e. Flash Attention 3 [39] for Hopper and Flash Attention 2 [12] for Ampere).

## D Samples from generative models

In Figs. 6 to 8, we present samples from the three generative models we experimented with.

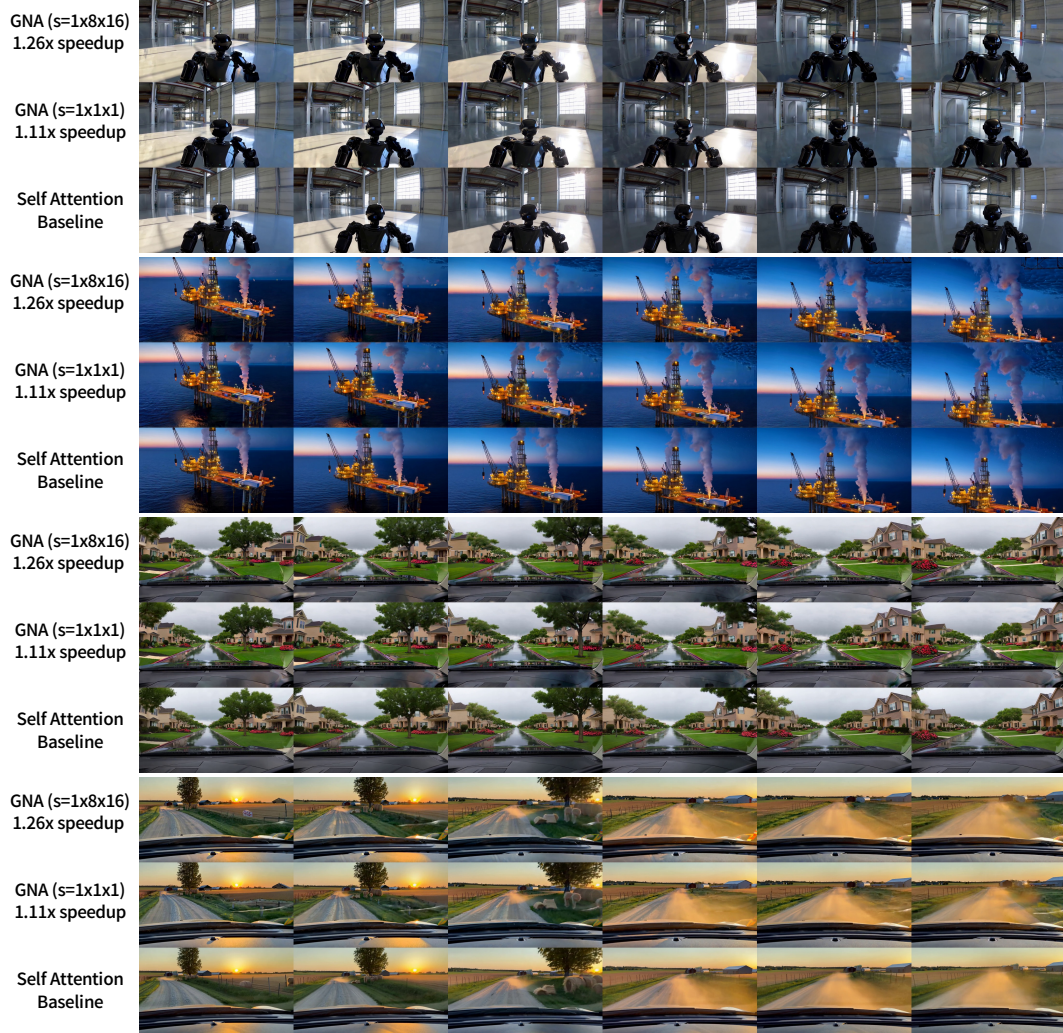


Figure 6: Sample frames from videos generated by Cosmos-7B, with GNA introduced into the last 23 of the 35 diffusion steps. Window size is  $16 \times 32 \times 48$  ( $\approx 56\%$  sparsity). Speedup limit under this setting, with the same level of sparsity, is 1.28x.




Figure 7: Sample frames from videos generated by HunyuanVideo, with GNA introduced into the last 35 of the 50 diffusion steps. Window size is  $18 \times 24 \times 24$  ( $\approx 91\%$  sparsity). Speedup limit under this setting, with the same level of sparsity, is 1.63x.






Figure 8: Images generated by ultra-resolution FLUX [28, 51] with GNA introduced into the last 19 of the 28 diffusion steps. Window size is  $80 \times 80$  ( $\approx 90\%$  sparsity). Speedup limit under this setting, with the same level of sparsity, is 1.46x.

## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope?

Answer: [Yes] .

Justification: All claims made are with regard to our implementation of Generalized Neighborhood Attention. Claims are also based on experiments presented in the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes] .

Justification: Limitations in terms of performance upper bounds, diminishing returns, and re-implementation required for different architectures are discussed in full.

### 3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA] .

Justification: [NA] .

### 4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: any reported performance numbers can be reproduced on the same hardware and with the specified versions of software, alongside our released software distribution.

### 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: All of our implementations will be open sourced. Release of anonymized code is not possible.

### 6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes] .

Justification: No training was done for this paper, as we plug our method into off-the-shelf generative models. We use their official codebase with no modification other than the integration of our method. Details such as the parameters for our method, number of diffusion steps using our method, and the like, are presented in the paper.

### 7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA] .

Justification: [NA] .

### 8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?



747 Answer: [Yes] .

748 Justification: Our focus is on one specific GPU architecture, and the requirement for

749 reproducing our results is using the same GPU model as in the paper.

750 **9. Code of ethics**

751 Question: Does the research conducted in the paper conform, in every respect, with the

752 NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

753 Answer: [Yes] .

754 Justification: -

755 **10. Broader impacts**

756 Question: Does the paper discuss both potential positive societal impacts and negative

757 societal impacts of the work performed?

758 Answer: [NA]

759 Justification: The work introduces a set of implementations, and does not of itself pose any

760 societal impacts positive or negative.

761 **11. Safeguards**

762 Question: Does the paper describe safeguards that have been put in place for responsible

763 release of data or models that have a high risk for misuse (e.g., pretrained language models,

764 image generators, or scraped datasets)?

765 Answer: [NA] .

766 Justification: [NA] .

767 **12. Licenses for existing assets**

768 Question: Are the creators or original owners of assets (e.g., code, data, models), used in

769 the paper, properly credited and are the license and terms of use explicitly mentioned and

770 properly respected?

771 Answer: [Yes] .

772 Justification: all prior arts have been cited, and all packages / open source projects and

773 software used are credited and if possible cited.

774 **13. New assets**

775 Question: Are new assets introduced in the paper well documented and is the documentation

776 provided alongside the assets?

777 Answer: [Yes] .

778 Justification: The open source project associated is well documented for all uses.

779 **14. Crowdsourcing and research with human subjects**

780 Question: For crowdsourcing experiments and research with human subjects, does the paper

781 include the full text of instructions given to participants and screenshots, if applicable, as

782 well as details about compensation (if any)?

783 Answer: [NA] .

784 Justification: [NA] .

785 **15. Institutional review board (IRB) approvals or equivalent for research with human**

786 **subjects**

787 Question: Does the paper describe potential risks incurred by study participants, whether

788 such risks were disclosed to the subjects, and whether Institutional Review Board (IRB)

789 approvals (or an equivalent approval/review based on the requirements of your country or

790 institution) were obtained?

791 Answer: [NA] .

792 Justification: [NA] .

793 **16. Declaration of LLM usage**

794 Question: Does the paper describe the usage of LLMs if it is an important, original, or  
795 non-standard component of the core methods in this research? Note that if the LLM is used  
796 only for writing, editing, or formatting purposes and does not impact the core methodology,  
797 scientific rigorousness, or originality of the research, declaration is not required.  
798 Answer: [NA] .  
799 Justification: [NA] .