

Block Cascading: Training Free Acceleration of Block-Causal Video Models

Hmrishav Bandyopadhyay^{1,2} Nikhil Pinnaraju¹ Rahim Entezari¹

Jim Scott¹ Yi-Zhe Song² Varun Jampani¹

¹Stability AI ²SketchX, University of Surrey



Figure 1. We propose Block Cascading (right), where future blocks of frames are cascaded with current ones to improve generation speed. Without requiring any fine-tuning, Block Cascading can improve inference FPS by upto 2.79x (right) in block-causal pipelines without negatively affecting quality for short, long, and interactive video generation (left). **All FPS is reported using the same python environment with Flash Attention 3 [31] on H100 GPUs. Most prompts are condensed in all figures, with full prompts available in Supplementary PDF and all videos available in Supplementary zip.**

Abstract

Block-causal video generation faces a stark speed-quality trade-off: small 1.3B models manage only 16 FPS while large 14B models crawl at 4.5 FPS, forcing users to choose between responsiveness and quality. Block Cascading significantly mitigates this trade-off through training-free parallelization. Our key insight: future video blocks do not need fully denoised current blocks to begin generation. By starting block generation with partially denoised context from predecessors, we transform sequential pipelines into parallel cascades where multiple blocks denoise simultaneously. With 5 GPUs exploiting temporal parallelism, we achieve $\sim 2\times$ acceleration across all model scales: 1.3B models accelerate from 16 to 30 FPS, 14B models from 4.5 to 12.5 FPS. Beyond inference speed, Block Cascading eliminates overhead from KV-recaching (of $\sim 200\text{ms}$) during context switches for interactive generation. Extensive evaluations validated against multiple block-causal pipelines demonstrate no significant loss in generation quality when switching from block-causal to Block Cascading pipelines for inference.

1. Introduction

Diffusion-based block-causal video generation pipelines denoise videos auto-regressively for faster, streamable gen-

eration. However, generation speed decreases significantly with model scale [16, 39]. Small-scale diffusion models (1.3B parameters) barely reach 16 FPS [14], while high-quality large-scale models (14B parameters [1]) crawl at 4.5 FPS on H100 GPUs, making them impractical for any interactive use. This speed-quality trade-off forces users to choose between responsive but lower-quality small models, or impressive but unusably slow large models. We revisit this trade-off, arguing that beyond model complexity it stems from rigid scheduling with strict causal dependencies. To overcome this, we introduce parallelised video generation with Block Cascading, unlocking faster inference with multiple GPUs and smarter scheduling.

The key intuition driving our work is that future frames need not rely on fully denoised past frames to begin generation. Current block-causal models like CausVid [51], Self-Forcing [14], and LongLive [47] enforce strict sequential dependencies: block B^{i+1} must wait for block B^i to completely denoise from $t=1000$ to $t=0$ before starting. This rigid ordering allows the denoising of only one block at a time in a causal manner, limiting opportunities to parallelise generation with multiple GPUs. We demonstrate that this strict dependency is overly conservative, especially since many of these block-causal models are adapted from bidirectional pre-trained networks [39] that naturally accommo-

date noisy context. Empirically, we observe that conditioning B^{i+1} on partially denoised outputs from B^i at intermediate timesteps (e.g. $t=750$) can yield visual quality comparable to using fully denoised blocks.

To build on this, we propose Block Cascading to overlap denoising of multiple blocks for enabling parallel generation across GPUs. Instead of sequential processing (B^1 completes, then B^2 starts, then B^3 starts), we cascade block generation: when B^1 reaches an intermediate timestep, we use its features as noisy context to start denoising B^2 . Similarly, B^3 begins denoising once B^2 reaches this intermediate checkpoint. This transforms the generation pipeline from sequential to parallel where multiple blocks denoise simultaneously at different stages, significantly reducing inference time. In effect, this reduces the waiting time for future blocks, which are already partially generated when the current block is decoded and viewed. Further, by cascading current and future blocks together, we can leverage bidirectional attention amongst these blocks to potentially improve generation quality.

Pre-generation of future blocks helps exploit their noisy context during prompt switches for interactive video generation. In block-causal pipelines, context switches have to be accompanied by expensive KV-recaching, where past KV is re-cached using new prompts, causing 200ms+ latency spikes which break interactivity. With Block Cascading, we instead switch context by simply changing the text prompt as future blocks demonstrate a gradual adaptation to new context based on their current noise level. This allows for seamless context switching without blocky artifacts or abrupt changes in scenes and subjects.

Finally, we note that our multi-GPU pipeline exploits temporal parallelism: each GPU processes a different temporal segment of the video. This approach is fundamentally distinct from data parallelism (processing independent samples) or model parallelism (splitting layers). However, sub-linear scaling (2.79× FPS on 5 GPUs with a 14B model) is expected due to VAE decoding after generation and overhead from KV communication across GPUs for self-attention. Optimizations targeting these factors, such as linear attention, quantization, or smaller VAEs can be applied on top of block-cascaded models for additional gains. Moreover, VAE decoding can be moved to a separate GPU to overlap decoding with next-block denoising. We do not discuss these optimizations in detail as they are not unique to block cascading.

Our contributions are: (i) We identify that strict causal dependencies in video generation are unnecessary, as noisy intermediate states provide sufficient context for generation; (ii) We propose Block Cascading, a training-free parallelization strategy that works universally across different block-causal models; (iii) We achieve practical inference speed with 5 GPUs: 30 FPS for 1.3B models and 12.5 FPS

for 14B models, while maintaining quality comparable to baselines; (iv) We demonstrate prompt switching without latency spikes in interactive applications, previously impossible with causal models. Across short [14], long [47], and interactive [47] video generation, and even across model scales [1], Block Cascading consistently delivers significant speed-up (average $\sim 2\times$ with 5 GPUs, 10% on single GPU) without any fine-tuning.

2. Related works

2.1. Timestep Distillation

Timestep Distillation [4, 17, 49, 50, 57] of diffusion models [12] compresses their diffusion ODE trajectory [34] for speeding up iterative denoising from noise to data. This setup typically involves using a multi-step teacher model to train a few-step student by comparing their outputs along different points of the trajectory. Progressive distillation [17, 30] directly compares teacher and student outputs for same inputs in short skips of ODE trajectory. Advanced approaches like DMD [49, 50] and SiD [58] use scores from teacher and student models to align the distributions of multi-step teacher and few-step student. Recent works can perform aggressive ODE compression [6], making it possible to generate in as few as one step [4, 50], generally by supplementing distribution matching with adversarial techniques [4, 17] and progressive distillation [50] ideas. We perform few-step inference with timestep-distilled models, to reduce the complexity of Block Cascading inference.

2.2. Auto-Regressive Video Generation

Early Autoregressive models [37, 46] have been predominantly used in conjunction with adversarial training [33, 37, 42] for sequential frame-by-frame video generation. With the advancement of diffusion-based generative models for image generation [10, 12, 29], recent works in video generation [7, 16, 22, 39, 40, 48] have adopted diffusion models to generate all video frames at once with full bidirectional attention across frames. These bidirectional models [16, 39, 48] can generate high quality videos but are slow and bounded by quadratic attention scaling [41] with respect to number of generated frames.

To improve speed in video generation pipelines, recent works [18, 23, 54–56] attempt to reduce number of sampling steps to directly reduce generation latency through timestep distillation. In parallel, several works re-introduce auto-regressive generative modelling of videos [11, 14, 19, 21, 44, 47, 51] by distilling bidirectional models [14, 51] or training auto-regressive pipelines from scratch [5, 36]. Fast, causal video generative pipelines set the stage for interactive and controllable generation [32, 47, 52], letting users inject control and action prompts that can influence future frames. In this work, we relax causality restrictions in distilled auto-regressive models to speed up inference by parallel process-

ing of future blocks.

2.3. KV Caching

Autoregressive generation of videos [14, 47] can be accelerated during inference with KV caching, where Key-Value (KV) pairs of previous frames are retained for efficient attention computation [38]. Since information from previous frames provides auto-regressive models with context for current frame generation, KV caching [28] helps prevent recomputing these features by simply storing them. KV caching was originally introduced in LLMs [2, 9, 13, 20, 38, 45] and has been since adapted to video generation [11, 14, 47, 51] to improve generation efficiency. While KV caching reduces KV computation for attention, the attention unit itself scales at a quadratic rate [41] with KV size, requiring eviction of older cache elements time to time to support longer generation. Evicting salient features, however, can cause issues like drifting [53] where the model forgets key context from previous time-frames. To reduce attention complexity without drifting, both LLMs [35, 45] and video generation approaches [14, 21, 47] use some form of KV sink [35], where key features from earlier frames are retained. In this paper, we share KV features from parallel blocks to compute self-attention, reducing the requirement of an external KV cache during inference.

3. Methodology

3.1. Background and Overview

Diffusion based multi-step generative models [16, 39, 48] have predominantly used bidirectional attention for high quality video generation at the cost of extra compute and generation time. Recent few-step video generation works like CausVid [51] and Self-Forcing [14] reduce generation time by distilling multi-step to few-step diffusion and by modelling auto-regressive generation on top of bidirectional models [39].

Adapting bidirectional pipelines like Wan2.1[39] for auto-regressive causal generation is non-trivial as it requires training to accommodate for the fundamental distribution shift where current frames x^i are modelled through previous frames $x^{<i}$ only. This is usually in the form of a regression-based objective including points on the teacher and student trajectories where the student is a causal network. This student is further trained as a few-step model with a timestep-distillation objective to catch up to the multi-step bidirectional teacher. CausVid trains from Wan2.1 [39] by matching distributions of the Wan2.1 teacher (p_{real}) and the causal few-step student (p_{fake}) as:

$$D_{\text{KL}}(p_{\text{fake}}||p_{\text{real}}) = -\mathbb{E}_{x \sim p_{\text{fake}}} \left(\log \frac{p_{\text{real}}(x)}{p_{\text{fake}}(x)} \right) \quad (1)$$

for samples x generated by the student. This KL divergence being intractable, the DMD objective allows representing

the gradient of KL as a difference of scores [25, 34, 43, 58] from the teacher and the student model as:

$$\nabla_{\theta} \mathcal{L}_{\text{DMD}} = -\mathbb{E}_{x \sim p_{\text{fake}}} \left((s_{\text{real}}(x) - s_{\text{fake}}(x)) \frac{dG_{\theta}}{d\theta} \right) \quad (2)$$

where G_{θ} corresponds to the student network, and s_{fake} and s_{real} are the teacher and student scoring functions respectively. In practice, we use the teacher network for s_{real} and a proxy s_{fake} trained to adapt to the changing distribution of the student model G_{θ} . To provide the proxy model enough room for catching up to the generator, it is trained for N iterations at every generator update.

Block-causal video generation: To balance quality and speed, CausVid [51], Self-Forcing [14] and other recent works [21, 47] employ a block-causal (rather than strictly causal) pipeline. In here, latent frames are temporally bunched into blocks $B_t^i = \{f_{t_n}^k : k \in K\}$, which are processed in a causal (i.e. $B^0 \rightarrow B^1 \rightarrow B^2 \dots$) fashion. Internally each block with it's group of frames uses bidirectional attention to generate video segments. During inference, blocks are rolled out sequentially, and denoised from $t=1000$ to $t=0$.

CausVid [51] and Self-Forcing [14] introduce efficient causal video generation through KV caching [26], where future blocks can depend on pre-computed Key-Value (KV) features of previous blocks during self-attention. This prevents recomputing KV features for obtaining context from previous blocks which have been denoised. To generate long videos and improve interactive control in generated frames, LongLive [47] introduces KV-recaching where stored KV is recalculated for new prompts. This form of KV caching where future blocks depend on clean KV from previous blocks requires future blocks to wait for current block to finish denoising. Further in the event of a KV-recache, all previous cache has to be discarded and recomputed - causing a drop in inference FPS.

In this work, we propose to remove this *stall* by denoising noisy future blocks B_t^{i+1} with noisy context in current block B_{t-1}^i . In other words, we *speculatively* roll forward with a *noisy cache* consisting of partially denoised KV features to get started with video generation in future blocks.

3.2. Noisy Caching

For Block-wise Causal generation of videos, blocks $B_t^i = \{f_t^j, f_t^{j+1}, f_t^{j+2}\}$ with frames f^* and $t > 0$ are denoised using previously denoised blocks $B_0^{<i}$. In a multi-step diffusion environment, the block B^i has to wait till all previous blocks $B_0^{<i}$ are denoised to $t=0$ for using their *clean* KV features in block-causal self-attention. However, block-causal distilled models [14, 47, 51] are often trained from bidirectional teachers [39] that use noisy features from all frames with bidirectional attention. This allows us to exploit the underlying bidirectional pre-training of causal dis-




Figure 2. **Block Cascading Inference:** Our inference pipeline reduces dependency on clean, denoised blocks for future block generation. For example in here (center), we denoise block $B_{t_3}^1 = \{f^3, f^4, f^5\}$ and $B_{t_2}^0 = \{f^3, f^4, f^5\}$ jointly using bidirectional attention instead of waiting for $B_{t_c}^0$ to denoise $B_{t_3}^1$ (left) [14, 47, 51]. By reducing dependency on previous blocks, we can free up the inference pipeline and allow parallel processing of multiple blocks (right) to improve generation speed significantly.

tiled models to kickstart frame generation in future blocks even when current block is not completely denoised.

Specifically, we start denoising block B_T^i at $t=T$ using noisy cache from blocks $\{B_{T-k}^{i-k}\}_{k=1}^i$ instead of waiting for blocks $B^{<i}$ to finish denoising. Again, at $t=t_0$, the previous blocks $B^{<i}$ have already finished denoising, but we retain context for $B_{t_0}^i$, through another forward pass, using $B_{t_c}^{i-1}$ where $t_c=0$ corresponds to the cache timestep. This is similar to how previous works [14, 51] construct KV cache from denoised blocks with an extra forward pass. Forward pass for block $B_{t_0}^i$ thus requires attending to KV from $\{B_{t_c}^{i-k}\}_{k=1}^i$ in addition to its own for self-attention.

In practice however, it is not computationally feasible to record noisy KV for all previous blocks $B^{<i}$ on a $N=50$ step diffusion model like Wan2.1[39], even if just for inference. Working with a limited window size of M blocks, this has a complexity of $O(N \times M \times M)$, with Dot-Product Attention complexity of $O(M)$ for KV size of M and a fixed number of frames per block. Since complexity linearly scales with number of steps, we can directly reduce the complexity of this operation using distillation techniques [4, 49, 50] that reduce inference timesteps to improve feasibility. We use 4 steps to denoise from $T=1000$ (i.e. $\{t_3, t_2, t_1, t_0\}$) and an additional step $t_c=0$ to improve video consistency across different blocks with cleaner cache from previous frames.

Naive implementation of this 4 step semi-causal denoising brings us to another sequential trap where we denoise the current block $B_t^i \rightarrow B_{t-1}^i$ and then start denoising the next block B_{t-1}^{i+1} using KV from B_{t-1}^i sequentially. This kickstarts future block denoising but slows down current block, resulting in slow inference. Concurrent work Rolling-Forcing [21] proposes to stitch the blocks together and jointly denoise all frames, following previous research

on robust training for causal video models [3]. Fundamentally different to Rolling Forcing [21] and Diffusion Forcing [3], we propose denoising of each block *individually* and *parallelly* through mini-batching, following individual block denoising in CausVid [51] and Self-Forcing [14]. We construct mini-batches where previous and current blocks are batched together with shared attention. (see Fig. 2).

3.3. Block Cascading

Mini-batching multiple blocks together can overburden a single GPU and reduce generation FPS (see Fig. 5a). To reduce the computation burden of our pipeline on a single GPU, we process batches on different GPUs and share KV features for computing self-attention. Using a fixed window size of M , we batch together blocks that depend on each other for denoising and can share KV during denoising. For denoising each mini-batch, we share KV features of each block with other blocks being denoised at each self attention layer. This is done by creating a global *shared KV pool* which can be referred to by all blocks, enabling continuity in generated frames. By nature, this KV sharing mimics rolling KV cache in block-causal generation pipelines [14, 51] that allows generation of long videos without running out of memory with fixed window sizes. Maximum parallelisation can be attained with a window-size equal to number of denoising steps (with t_c) where each block depends on previous block’s cleaner timestep. Increasing window size beyond this requires maintaining an add on KV cache from previous blocks (similar to [14]). While smaller window sizes are possible by restricting attention window during self-attention, we find they can impact generation as later blocks do not get access to clean latents at $t=t_c$. A better way to obtain smaller windows is through reduced parallelism (see Sec. 4.2).

We note that generating long videos with a fixed win-

down often results in drifting artifacts [53]. Following recent works [47], we reduce drifting in long video generation by using the first block as a ‘sink’ token, that stays in an external KV cache.

Algorithm 1 Multi-GPU Inference with Block Cascading

Require: Generator G_θ , Timesteps $\mathcal{T} = \{t_3, t_2, t_1, t_0, t_c\}$
Require: Total Frames M , Block Size S , GPUs $\{g_1, \dots, g_G\}$, Scheduler ψ , VAE Decoder $\mathcal{D}$

- 1: **Block computation:**
- 2: $B \leftarrow \lceil M/S \rceil$ ▷ Number of blocks
- 3: **Initialize:**
- 4: $\mathbf{X} \sim \mathcal{N}(0, I)$ ▷ Noisy latents, size M
- 5: $\mathbf{O} \leftarrow \emptyset$ ▷ Output latents
- 6: $b \leftarrow 0$ ▷ Primary block index
- 7: $\tau \leftarrow 0$ ▷ Timestep index in $\mathcal{T}$
- 8: $d \leftarrow 0$ ▷ Cascade depth
- 9: **while** $d \geq 0$ **do**
- 10: $\mathcal{B} \leftarrow \bigcup_{i=0}^{\min(d, G-1)} \{(\mathbf{x}_{b+i}, \mathcal{T}[\tau-i])\}$ on g_i ▷ Create Batch
- 11: $\{\hat{\mathbf{x}}_0^{g_0}, \dots, \hat{\mathbf{x}}_d^{g_d}\} \leftarrow G_\theta(\mathcal{B})$
- 12: **for each** $\hat{\mathbf{x}}_j^{g_j}$ **do**
- 13: **if** $\mathcal{T}[\tau-j] > t_0$ **then**
- 14: $\epsilon \sim \mathcal{N}(0, I)$
- 15: $\mathbf{X}[b+j] \leftarrow \Psi(\hat{\mathbf{x}}_j^{g_j}, \epsilon, \mathcal{T}[\tau-j+1])$
- 16: **else if** $\mathcal{T}[\tau-j] = t_0$ **then**
- 17: $\mathbf{O}[b+j] \leftarrow \hat{\mathbf{x}}_j^{g_j}$ ▷ For streaming, decode here
- 18: **end if**
- 19: **end for**
- 20: **if** block b completed t_c **then**
- 21: $b \leftarrow b+1$; $d \leftarrow d-1$ **if** $b+d \geq B$
- 22: **else**
- 23: $d \leftarrow \min(d+1, G-1)$; $\tau \leftarrow \tau+1$
- 24: **end if**
- 25: **end while**
- 26: **return** $\mathcal{D}(\mathbf{O})$ ▷ VAE decode




Figure 3. **Expensive KV-recaching:** KV-recaching can result in FPS drop in interactive generation as previous generated and stored cache has to be recached using new context. We skip KV-recaching as our KV is auto-recomputed using new context.

3.4. Training Free Adaptation

Naive causal denoising with Block Cascading can cause artifacts, specifically in fine-grained features like colour and

pattern on a surf-board in Fig. 6. These artifacts are noticeable in video models as frame-jumps, where fine-grained details are not aligned across frames. We find that artifacts are prominent particularly in early frames where the model does not have enough context from clean latents (*i.e.* $B_{t_c}^{<i}$). For example, during denoising B^1 using KV features from B^0 , B^1 observes clean features from $B_{t_c}^0$ only during it’s last time-step, during denoising $B_{t_0}^1$. This creates occasional distortion from context mismatch between previously generated context $B_{t_0}^1$ and new cache context $B_{t_c}^0$. Consequently, we find these distortions missing in ablative experiments with lower parallelism (see Sec. 4.2) where previously generated context (*e.g.* $B_{t_2}^1$) is still too noisy for conflict. To fix these artifacts and generate smooth videos, we exploit bidirectional pre-training and use full bidirectional attention across the current cascade. Bidirectional attention here helps smooth out inconsistencies by naturally aligning the features of both current and future frames during video generation (see Fig. 6). Since bidirectional attention uses more context to denoise every block, it can even surpass the block-causal baseline in quality at times (see Fig. 8).

3.5. Interactive Control

To allow interactive control over video generation, current block-causal denoising pipelines [47] perform ‘KV-recaching’, recomputing past KV cache with new prompts for influencing next block generation. While starting denoising at $t=1000$, current blocks are heavily influenced by clean past cache, and without KV-recaching, can completely ignore new instructions. However, KV-recaching heavily drops inference FPS as all previous blocks in the attention window have to be recached before current block can start generating (see Fig. 5c). We demonstrate that KV-recaching can be avoided with our Block Cascading pipeline since we adapt to changing context easily with new text prompts. Specifically during a context switch, future blocks are at various stages of denoising while the first block is denoised and is entering a cache state $B_{t_0}^i \rightarrow B_{t_c}^i$. Directly injecting new context here allows bidirectional attention to handle the context switch where different frames in the cascade are affected with the new context based on their current noise levels. Similar to KV-recaching, our first block is being cached with new context in $B_{t_0}^i \rightarrow B_{t_c}^i$ and provides clean KV features for all future blocks. We find that in practice, videos generated with bidirectional attention over our cascade can blend context much smoother than KV-recaching (more in Supplementary) and are preferred by users (see Fig. 8).

4. Experiments

We conduct experiments by adapting causal video models without training. Specifically, we adapt Self-Forcing [14], LongLive [47] and Krea [1], obtaining an inference FPS




Figure 4. **Qualitative comparisons:** Comparing our bidirectional inference pipeline with corresponding original inference pipelines

boost of 14.05, 12.66 and 8.07 respectively with multiple GPUs. We choose these particular models for a holistic assessment over different inference domains like short, long, and interactive video generation across different model capacities. For all inference configurations, we use a window size of 7 blocks, and specifically for LongLive, a sink size of 1 block where all blocks have 3 frames each. LongLive’s static sink can corrupt dynamic videos, so for interactive generation, we generate our samples without sink while LongLive refreshes sink during KV-recaching. For context change in interactive generation, we simply change the text prompt without incurring any artifacts, which helps us to prevent FPS drop from expensive KV recaching

(see Sec. 3.3). We compare quality of videos generated with our pipeline v/s the original inference pipelines with qualitative (Fig. 4) and quantitative (Sec. 4.4, Tab. 1) analysis and a comprehensive user study (Fig. 8).

We compare our videos with those generated with (i) *Block-causal video generation pipelines:* CausVid [51], Self-Forcing [14], LongLive [47], SkyReels-V2 [5] and MAGI-1 [36] and (ii) *Bidirectional Video generation pipelines:* Rolling Forcing [21], FastVideo [54] and Wan2.1 [39]. CausVid [51] learns a block-causal video generative model with the DMD [50] objective in Eq. (2). Self-Forcing [14] builds on CausVid, reducing mismatch between training and inference pipelines for block-causal




Figure 5. **Instantaneous FPS:** We measure FPS as time taken to denoise a particular block of 3 latent frames (12 video frames). FPS fluctuates from changing attention window sizes during block-causal denoising and number of parallel blocks in Block Cascading.

generation by simulating the inference environment during training. LongLive [47] addresses drifting by training with a permanent attention sink that retains specific context irrespective of causal window size. Separate from these models that are fine-tuned from Wan2.1 [39], SkyReels-V2 [5] and MAGI-1 [36] train block-causal pipelines from scratch. Fully bidirectional FastVideo [54] distills from Wan2.1 [39] using a sparse attention strategy. Rolling-Forcing [21] implements bidirectional attention on top of block-causal generations strategies and distills from Wan2.1 [39] using a rolling bidirectional attention window.

4.1. FPS Analysis

In Block Cascading, FPS initially starts low while blocks are being loaded in the cascade, then stabilises as the cascade is full, and finally goes up towards the end (Fig. 5b and Fig. 5a) as blocks are evicted. In contrast, causal video generation pipelines have initial blocks with smaller attention windows that are faster to generate than later blocks with larger windows. For our experiments, we want to define ‘Streaming FPS’ as the FPS when the attention window is full and new blocks are being generated and decoded with the VAE. This definition reflects the steady-state performance that a user would experience during long, auto-regressive video generation. We report Streaming FPS in Tab. 1 and in Fig. 5 as the average FPS across 8th and 9th blocks during generation of 13 blocks with a 7-block attention window. Additionally, we plot instantaneous FPS observed during generation to illustrate FPS changes (Fig. 5a and Fig. 5b) and to highlight FPS drops in interactive video generation (Fig. 5c).

4.2. Ablative Studies

For ablation experiments, we analyse the degree of parallelisation attainable with different forms of Block Cascading, using Self-Forcing [14] pre-trained checkpoint as our baseline. We define different parallelization forms based on timesteps for cascading, from $t \in \{t_3, t_2, t_1, t_0, t_c\}$, denoting $\mathcal{P}_i$ for i -timestep skip by using $B_{t_c-i}^j$ for denoising $B_{t_0}^{j+1}$. We start with baseline self-forcing as $\mathcal{P}_1$, where blocks are generated sequentially and cached features from




Figure 6. **Bidirectional Inference:** Causal attention in fully parallelised generation ($\mathcal{P}_4$) can yield artifacts. These can be fixed by using bidirectional attention with same fully parallelised pipeline.




Figure 7. **Cascading Types:** Different types of cascading leads to different levels of parallelism. We conduct a user study to analyse performance across differently cascaded pipelines. At $\mathcal{P}_1$, the pipeline reduces to block-causal pipelines like Self-Forcing [14]

$B_{t_c}^i$ are used for denoising $B_{t_0}^{i+1}$. To reduce dependency on t_c , we allow block $B_{t_0}^{i+1}$ to denoise using cache from $B_{t_3}^i$ in $\mathcal{P}_2$ which cascades blocks $B_{t_1}^{i+1}$ and $B_{t_0}^i$. Next, we move block $B_{t_0}^{i+1}$ further up (see Fig. 7), denoising with block $B_{t_2}^i$ for a two block cascade and with $B_{t_c}^i$ and $B_{t_0}^{i+2}$ for a three block cascade at $t=t_2$. Finally, we use our default configuration $\mathcal{P}_4$, which is a five block cascade (e.g. with $\{B_{t_c}^0, B_{t_3}^1, B_{t_2}^2, B_{t_1}^3, B_{t_0}^4\}$). We observe from Fig. 5a and Fig. 5b that causal generation from $\mathcal{P}_4$ has the highest FPS both in single (+2.9 FPS) and multi GPU (+14.5 FPS) environments. However, causal generation can yield artifacts (see Fig. 6) in early frames with limited clean con-




Figure 8. **User Study:** We conduct a user study to analyse video quality from our cascaded bidirectional pipeline against original inference pipelines of other models under different configurations.

text. In a single GPU setup, switching to $\mathcal{P}_3$ solves this as cleaner context is available for later frames to reduce artifacts (Fig. 7). In a multi GPU environment, causal and bidirectional attention have marginal difference in FPS, and the best option is to use Bidirectional $\mathcal{P}_4$ that yields high quality with high generation FPS. For analysing video quality with different ablative environments, we perform user studies (see Sec. 4.3) observing better performance with bidirectional $\mathcal{P}_4$ and causal $\mathcal{P}_2$ and $\mathcal{P}_3$, compared to causal $\mathcal{P}_4$.

4.3. User Study

To assess degradation in generated videos, we perform user studies where we show users videos generated by Self-Forcing [14], LongLive [47], and Krea [1] inference pipelines and those generated by ours. With the same seed and prompt, our method produces videos that closely match those generated by the vanilla inference pipeline (see Fig. 4). Therefore, we do not report prompts in our study, except in interactive video analysis where prompts are essential for illustrating context changes. The instruction provided is to choose the video that has higher quality given two competing videos - with each video reviewed by three annotators. The options available to users are ‘left’, ‘right’, ‘both’ and ‘none’. We include more details of the user study in supplementary. For Self-Forcing, Krea and our ablative configurations (Sec. 4.2 and Fig. 7), we generate 418 short videos from each method with prompts sampled from VBench [15] extended prompts [14]. We evaluate LongLive on 100 long videos with prompts sampled from MovieGen [24] extended prompts and 20 interactive videos with prompts from [47] and Gemini 2.5 Pro [8]. Both long videos and interactive videos are generated to a total length of 120 latent frames (30 seconds) while short videos are generated to 21 latent frames (5 seconds). From results of this user study in Fig. 8, we note that users have a hard time distinguishing our pipeline from original pipeline when compared with Self-Forcing or LongLive. We find our pipeline is preferred in interactive video generation with LongLive and short video generation with Krea, potentially from additional context from bidirectional attention.

4.4. Quantitative Analysis

We use VBench [15] for quantitative analysis with prompts re-written using Qwen2.5-7B-Instruct [27] following other methods [14]. All samples for evaluation with VBench in Tab. 1 use re-written prompts for video generation. In addition to a brief summary in Tab. 1, we include a complete table with all VBench scores in Supplementary. From Tab. 1, inference with our denoising pipeline closely follows scores of the original pipelines without significant drop in quality. During inference, our model is parallelised across 5 GPUs, providing us an avg. boost of $\sim 2\times$ in FPS.

Table 1. **Video quality benchmark:** We report VBench [15] scores for all competitors. FPS for other methods reported with 1xH100, FPS for ours reported on 5xH100s

Model	Streaming FPS ($\uparrow$)	Total Score ($\uparrow$)	Quality Score ($\uparrow$)	Semantic Score ($\uparrow$)
Wan2.1 [†] [39]	-	0.8381	0.8471	0.8019
FastVideo [†] [54]	-	0.8313	0.8647	0.6978
SkyReels-V2 [5]*	0.49	0.8267	0.8470	0.7453
MAGI-1 [36]*	0.19	0.7918	0.8204	0.6774
CausVid [51]	16.31	0.8502	0.8555	0.8290
Rolling Forcing [21]	17.80	0.8287	0.8358	0.8003
Self-Forcing [14]	16.31	0.8440	0.8498	0.8206
Ours (w Self-Forcing ckpt)	30.36	0.8353	0.8435	0.8024
LongLive [47]	15.60	0.8335	0.8378	0.8162
Ours (w LongLive ckpt)	30.46	0.8230	0.8232	0.8223
Krea-14B [1]	4.52	0.8437	0.8471	0.8304
Ours (w Krea-14B ckpt)	12.59	0.8492	0.8532	0.8332

*Numbers taken from [14], [†] Bidirectional pipelines cannot stream.

5. Limitations

We note that window size in pre-training configuration can be a limitation for Block Cascading when full parallelism is desired for optimal FPS boost.

Empirically, we notice slightly higher drifting in some samples during inference with a 7-block window, while using a checkpoint that was trained with a 4-block window [47]. Nevertheless, the degree of parallelism can be reduced (e.g. with $\mathcal{P}_3$ or $\mathcal{P}_2$) for generation with strict small window models at some cost to inference speed. We also note sub-linear scaling with GPUs in multi-GPU environments, making Block Cascading ideal for single video generation but a poor alternative for generating large batches of videos. The latter case is better parallelised with distributed sampling.

6. Conclusion

In conclusion, Block Cascading parallelises inference in block-causal video pipelines without requiring any re-training or fine-tuning. Extensive user studies demonstrate that Block Cascading during inference does not degrade video quality in pre-trained causal pipelines. This allows us to boost streaming inference speed for 1.3B video models to 30 FPS and 14B models to 12.5 FPS using multiple GPUs, supporting downstream applications like interactive and controllable video generation. Deploying large-scale video models like Krea-14B for dynamic

generation and live-viewing bridges the gap between high quality video generation and real-time inference.

References

- [1] Krea AI. Krea realtime 14b: Real-time video generation. 2025. 1, 2, 5, 8
- [2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *NeurIPS*, 2020. 3
- [3] Boyuan Chen, Diego Martí Monsó, Yilun Du, Max Simchowitz, Russ Tedrake, and Vincent Sitzmann. Diffusion forcing: Next-token prediction meets full-sequence diffusion. *NeurIPS*, 2024. 4
- [4] Dar-Yen Chen, Hmrishav Bandyopadhyay, Kai Zou, and Yi-Zhe Song. Nitrofusion: High-fidelity single-step diffusion through dynamic adversarial training. In *CVPR*, 2024. 2, 4
- [5] Guibin Chen, Dixuan Lin, Jiangping Yang, Chunze Lin, Junchen Zhu, Mingyuan Fan, Hao Zhang, Sheng Chen, Zheng Chen, Chengcheng Ma, et al. Skyreels-v2: Infinite-length film generative model. *arXiv preprint arXiv:2504.13074*, 2025. 2, 6, 7, 8
- [6] Junsong Chen, Shuchen Xue, Yuyang Zhao, Jincheng Yu, Sayak Paul, Junyu Chen, Han Cai, Enze Xie, and Song Han. Sana-sprint: One-step diffusion with continuous-time consistency distillation. *arXiv preprint arXiv:2503.09641*, 2025. 2
- [7] Junsong Chen, Yuyang Zhao, Jincheng Yu, Ruihang Chu, Junyu Chen, Shuai Yang, Xianbang Wang, Yicheng Pan, Daquan Zhou, Huan Ling, et al. Sana-video: Efficient video generation with block linear diffusion transformer. *arXiv preprint arXiv:2509.24695*, 2025. 2
- [8] Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blstein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025. 8
- [9] Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V Le, and Ruslan Salakhutdinov. Transformer-xl: Attentive language models beyond a fixed-length context. *arXiv preprint arXiv:1901.02860*, 2019. 3
- [10] 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 *ICML*, 2024. 2
- [11] Kaifeng Gao, Jiabin Shi, Hanwang Zhang, Chunping Wang, and Jun Xiao. Vid-gpt: Introducing gpt-style autoregressive generation in video diffusion models. *arXiv preprint arXiv:2406.10981*, 2024. 2, 3
- [12] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 2
- [13] Coleman Hooper, Sehoon Kim, Hiva Mohammadzadeh, Michael W Mahoney, Yakun S Shao, Kurt Keutzer, and Amir Gholami. Kvquant: Towards 10 million context length llm inference with kv cache quantization. *NeurIPS*, 2024. 3
- [14] Xun Huang, Zhengqi Li, Guande He, Mingyuan Zhou, and Eli Shechtman. Self forcing: Bridging the train-test gap in autoregressive video diffusion. *arXiv preprint arXiv:2506.08009*, 2025. 1, 2, 3, 4, 5, 6, 7, 8
- [15] Ziqi Huang, Yanan 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. 8
- [16] 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. 1, 2, 3
- [17] Shanchuan Lin, Anran Wang, and Xiao Yang. Sdxl-lightning: Progressive adversarial diffusion distillation. *arXiv preprint arXiv:2402.13929*, 2024. 2
- [18] Shanchuan Lin, Xin Xia, Yuxi Ren, Ceyuan Yang, Xuefeng Xiao, and Lu Jiang. Diffusion adversarial post-training for one-step video generation. *arXiv preprint arXiv:2501.08316*, 2025. 2
- [19] Shanchuan Lin, Ceyuan Yang, Hao He, Jianwen Jiang, Yuxi Ren, Xin Xia, Yang Zhao, Xuefeng Xiao, and Lu Jiang. Autoregressive adversarial post-training for real-time interactive video generation. *arXiv preprint arXiv:2506.09350*, 2025. 2
- [20] Hao Liu, Matei Zaharia, and Pieter Abbeel. Ring attention with blockwise transformers for near-infinite context. *arXiv preprint arXiv:2310.01889*, 2023. 3
- [21] Kunhao Liu, Wenbo Hu, Jiale Xu, Ying Shan, and Shijian Lu. Rolling forcing: Autoregressive long video diffusion in real time. *arXiv preprint arXiv:2509.25161*, 2025. 2, 3, 4, 6, 7, 8
- [22] Guoqing Ma, Haoyang Huang, Kun Yan, Liangyu Chen, Nan Duan, Shengming Yin, Changyi Wan, Ranchen Ming, Xiaoni Song, Xing Chen, et al. Step-video-t2v technical report: The practice, challenges, and future of video foundation model. *arXiv preprint arXiv:2502.10248*, 2025. 2
- [23] Xiaofeng Mao, Zhengkai Jiang, Fu-Yun Wang, Jiangning Zhang, Hao Chen, Mingmin Chi, Yabiao Wang, and Wenhao Luo. Osv: One step is enough for high-quality image to video generation. In *CVPR*, 2025. 2
- [24] Adam Polyak, Amit Zohar, Andrew Brown, Andros Tjandra, Animesh Sinha, Ann Lee, Apoorv Vyas, Bowen Shi, Chih-Yao Ma, Ching-Yao Chuang, et al. Movie gen: A cast of media foundation models. *arXiv preprint arXiv:2410.13720*, 2024. 8
- [25] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 3
- [26] Reiner Pope, Sholto Douglas, Aakanksha Chowdhery, Jacob Devlin, James Bradbury, Jonathan Heek, Kefan Xiao, Shivan Agrawal, and Jeff Dean. Efficiently scaling transformer inference. *MLSys*, 2023. 3

- [27] A Yang Qwen, Baosong Yang, B Zhang, B Hui, B Zheng, B Yu, Chengpeng Li, D Liu, F Huang, H Wei, et al. Qwen2. 5 technical report. *arXiv preprint*, 2024. 8
- [28] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 2019. 3
- [29] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2
- [30] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512*, 2022. 2
- [31] Jay Shah, Ganesh Bikshandi, Ying Zhang, Vijay Thakkar, Pradeep Ramani, and Tri Dao. Flashattention-3: Fast and accurate attention with asynchrony and low-precision. *NeurIPS*, 2024. 1
- [32] Joonghyuk Shin, Zhengqi Li, Richard Zhang, Jun-Yan Zhu, Jaesik Park, Eli Shechtman, and Xun Huang. Motion-stream: Real-time video generation with interactive motion controls. *arXiv preprint arXiv:2511.01266*, 2025. 2
- [33] Ivan Skorokhodov, Sergey Tulyakov, and Mohamed Elhoseiny. Stylegan-v: A continuous video generator with the price, image quality and perks of stylegan2. In *CVPR*, 2022. 2
- [34] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 2, 3
- [35] Zunhai Su and Kehong Yuan. Kvsink: Understanding and enhancing the preservation of attention sinks in kv cache quantization for llms. *arXiv preprint arXiv:2508.04257*, 2025. 3
- [36] Hansi Teng, Hongyu Jia, Lei Sun, Lingzhi Li, Maolin Li, Mingqiu Tang, Shuai Han, Tianning Zhang, WQ Zhang, Weifeng Luo, et al. Magi-1: Autoregressive video generation at scale. *arXiv preprint arXiv:2505.13211*, 2025. 2, 6, 7, 8
- [37] Yu Tian, Jian Ren, Menglei Chai, Kyle Olszewski, Xi Peng, Dimitris N Metaxas, and Sergey Tulyakov. A good image generator is what you need for high-resolution video synthesis. *arXiv preprint arXiv:2104.15069*, 2021. 2
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 2017. 3
- [39] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwei Yu, Haiming Zhao, Jianxiao Yang, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025. 1, 2, 3, 4, 6, 7, 8
- [40] Jiuniu Wang, Hangjie Yuan, Dayou Chen, Yingya Zhang, Xiang Wang, and Shiwei Zhang. Modelscope text-to-video technical report. *arXiv preprint arXiv:2308.06571*, 2023. 2
- [41] Sinong Wang, Belinda Z Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020. 2, 3
- [42] Yuhan Wang, Liming Jiang, and Chen Change Loy. Styleinv: A temporal style modulated inversion network for unconditional video generation. In *ICCV*, 2023. 2
- [43] Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. In *NeurIPS*, 2023. 3
- [44] Wenming Weng, Ruoyu Feng, Yanhui Wang, Qi Dai, Chunyu Wang, Dacheng Yin, Zhiyuan Zhao, Kai Qiu, Jianmin Bao, Yuhui Yuan, et al. Art-v: Auto-regressive text-to-video generation with diffusion models. In *CVPR*, 2024. 2
- [45] Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. Efficient streaming language models with attention sinks. *arXiv preprint arXiv:2309.17453*, 2023. 3
- [46] Wilson Yan, Yunzhi Zhang, Pieter Abbeel, and Aravind Srinivas. Videogpt: Video generation using vq-vae and transformers. *arXiv preprint arXiv:2104.10157*, 2021. 2
- [47] Shuai Yang, Wei Huang, Ruihang Chu, Yicheng Xiao, Yuyang Zhao, Xianbang Wang, Muyang Li, Enze Xie, Yingcong Chen, Yao Lu, et al. Longlive: Real-time interactive long video generation. *arXiv preprint arXiv:2509.22622*, 2025. 1, 2, 3, 4, 5, 6, 7, 8
- [48] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024. 2, 3
- [49] Tianwei Yin, Michaël Gharbi, Taesung Park, Richard Zhang, Eli Shechtman, Fredo Durand, and Bill Freeman. Improved distribution matching distillation for fast image synthesis. In *NeurIPS*, 2024. 2, 4
- [50] Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Fredo Durand, William T Freeman, and Taesung Park. One-step diffusion with distribution matching distillation. In *CVPR*, 2024. 2, 4, 6
- [51] Tianwei Yin, Qiang Zhang, Richard Zhang, William T Freeman, Fredo Durand, Eli Shechtman, and Xun Huang. From slow bidirectional to fast autoregressive video diffusion models. In *CVPR*, 2025. 1, 2, 3, 4, 6, 8
- [52] Jiwen Yu, Yiran Qin, Xintao Wang, Pengfei Wan, Di Zhang, and Xihui Liu. Gamefactory: Creating new games with generative interactive videos. *arXiv preprint arXiv:2501.08325*, 2025. 2
- [53] Lvmin Zhang, Shengqu Cai, Muyang Li, Gordon Wetzstein, and Maneesh Agrawala. Frame context packing and drift prevention in next-frame-prediction video diffusion models. In *NeurIPS*, 2025. 3, 5
- [54] Peiyuan Zhang, Yongqi Chen, Haofeng Huang, Will Lin, Zhengzhong Liu, Ion Stoica, Eric Xing, and Hao Zhang. Vsa: Faster video diffusion with trainable sparse attention. *arXiv preprint arXiv:2505.13389*, 2025. 2, 6, 7, 8
- [55] Peiyuan Zhang, Yongqi Chen, Runlong Su, Hangliang Ding, Ion Stoica, Zhengzhong Liu, and Hao Zhang. Fast video generation with sliding tile attention. *arXiv preprint arXiv:2502.04507*, 2025.
- [56] Zhixing Zhang, Yanyu Li, Yushu Wu, Anil Kag, Ivan Skorokhodov, Willi Menapace, Aliaksandr Siarohin, Junli Cao, Dimitris Metaxas, Sergey Tulyakov, et al. Sf-v: Single forward video generation model. *NeurIPS*, 2024. 2

- [57] Jianbin Zheng, Minghui Hu, Zhongyi Fan, Chaoyue Wang, Changxing Ding, Dacheng Tao, and Tat-Jen Cham. Trajectory consistency distillation: Improved latent consistency distillation by semi-linear consistency function with trajectory mapping. *arXiv preprint arXiv:2402.19159*, 2024. [2](#)
- [58] Mingyuan Zhou, Huangjie Zheng, Zhendong Wang, Mingzhang Yin, and Hai Huang. Score identity distillation: Exponentially fast distillation of pretrained diffusion models for one-step generation. In *ICML*, 2024. [2](#), [3](#)