

---

# DIFFUSIONPIPE: TRAINING LARGE DIFFUSION MODELS WITH EFFICIENT PIPELINES

---

Ye Tian<sup>1\*</sup> Zhen Jia<sup>2</sup> Ziyue Luo<sup>3</sup> Yida Wang<sup>2</sup> Chuan Wu<sup>1</sup>

## ABSTRACT

Diffusion models have emerged as dominant performers for image generation. To support training large diffusion models, this paper studies pipeline parallel training of diffusion models and proposes DiffusionPipe, a synchronous pipeline training system that advocates innovative pipeline bubble filling technique, catering to structural characteristics of diffusion models. State-of-the-art diffusion models typically include trainable (the backbone) and non-trainable (e.g., frozen input encoders) parts. We first unify optimal stage partitioning and pipeline scheduling of single and multiple backbones in representative diffusion models with a dynamic programming approach. We then propose to fill the computation of non-trainable model parts into idle periods of the pipeline training of the backbones by an efficient greedy algorithm, thus achieving high training throughput. Extensive experiments show that DiffusionPipe can achieve up to 1.41x speedup over pipeline parallel methods and 1.28x speedup over data parallel training on popular diffusion models.

## 1 INTRODUCTION

Diffusion models have become the dominant choice for content generation today, including text-image synthesis (Choi et al., 2021) and video generation (Ramesh et al., 2022). Large diffusion models such as Stable Diffusion (Rombach et al., 2022), ControlNet (Zhang & Agrawala, 2023), and Imagen (Saharia et al., 2022) achieve state-of-the-art performance in various scenarios. There is a continuing trend to develop larger diffusion models by increasing the backbone size (Rombach et al., 2022; Peebles & Xie, 2022a; Bao et al., 2023; Podell et al., 2023), cascading multiple backbones to enable higher resolution image generation (Nichol et al., 2021; Peebles & Xie, 2022a; Saharia et al., 2022; Ho et al., 2022; Podell et al., 2023), and combining different transformer architectures with diffusion models (Peebles & Xie, 2022a; Zhang & Agrawala, 2023; Wu et al., 2023).

Data parallelism is adopted for distributed diffusion model training (Falcon & The PyTorch Lightning team, 2019; Bian et al., 2021; von Platen et al., 2022). For large diffusion models, this method duplicates parameters, which limits the training batch size (Rombach et al., 2022; Ho et al., 2022; Saharia et al., 2022) and device utilization, and causes significant synchronization overhead, especially when the

training scale is large (Narayanan et al., 2019).

Pipeline parallelism (Huang et al., 2019; Narayanan et al., 2019; Luo et al., 2022) has been widely adopted to train large DNN models, which partitions networks across multiple devices and pipelines micro-batch processing across model partitions, substantially alleviating memory consumption on a single device and enabling larger training batch sizes. Although pipeline parallelism is potentially useful in enabling larger diffusion model training, it has not been well explored for diffusion models and its application faces several challenges as follows:

*First*, the structural characteristics and special training procedures of diffusion models cannot be handled well by traditional pipelining methods. A diffusion model typically contains a trainable part with one or multiple backbone models (e.g., U-Net) (Rombach et al., 2022), and a non-trainable part with frozen text and image encoders, and they are usually trained with special techniques such as self-conditioning (Chen et al., 2022), which involves an additional forward computation pass on the backbone. Pipeline training involves only the trainable part, while the non-trainable part is not readily handled by existing pipeline training methods because it does not require pipelining. Self-conditioning is beyond the scope of existing pipeline systems, as they assume that there is only one forward pass.

*Second*, pipeline bubbles are often significant in synchronous pipeline training (Huang et al., 2019; Fan et al., 2021; Luo et al., 2022), which is more widely used in practice due to not altering model performance but involves

\*Work done during internship at AWS. <sup>1</sup>The University of Hong Kong, Hong Kong <sup>2</sup>Amazon Web Services, USA <sup>3</sup>The Ohio State University, USA. Correspondence to: Ye Tian <yetiansh@connect.hku.hk>.

periodic pipeline flushing. We identify a unique opportunity to fill the pipeline bubbles using the computation of non-trainable model components, to substantially improve device utilization and expedite training speed. However, there are dependencies between the trainable and non-trainable part that block pipeline bubble filling by overlapping their execution. In addition, how to partition the non-trainable part into sets of layers and insert them into the pipeline bubble is not studied.

*Third*, Non-trainable layers with extra-long execution time are common in frozen encoders (Kingma & Welling, 2013). Such layers may not fit into any pipeline bubble and block filling pipeline bubble with all subsequent layers in the non-trainable part, which cannot be solved by only partitioning the non-trainable part into sets of layers. In addition, as non-trainable layers’ execution time is discrete, it is unlikely to fully utilize idle time in individual pipeline bubble, leading to performance degradation.

In this paper, we propose *DiffusionPipe*, an efficient pipeline training system designed specifically for large diffusion models. DiffusionPipe systematically determines optimized model partitioning, stages, and replication settings while applying pipeline bubble filling techniques. These optimizations are tailored for a variety of representative diffusion models and training methods. To the best of our knowledge, we are the first to enable efficient pipeline parallel training of diffusion models. Our contributions can be summarized as follows:

- ▷ We propose a unified dynamic programming-based algorithm for optimized model partitioning, that can handle various training scenarios, e.g., models with different numbers of backbones and models trained with self-conditioning. The proposed partitioning algorithm optimizes the model partitioning scheme under various settings of number of stages and number of micro-batches, with performance comparable to state-of-the-art pipeline paradigms under traditional pipelining, and effectively handles scenarios beyond traditional pipelining and specific to diffusion models.
- ▷ We design a novel pipeline bubble filling strategy that fills the non-trainable part computation into the bubble time of the pipeline training of the backbone(s), effectively eliminating pipeline bubbles. It efficiently partitions the non-trainable components and the input data for bubble filling, and effectively addresses dependencies between the non-trainable part and the trainable part by allowing *cross-iteration* overlapping of backbone training of an iteration and non-trainable part computation of the next iteration and filling pipeline bubbles of the former with the latter.
- ▷ We effectively handles extra-long non-trainable layers which do not fit into individual pipeline bubbles, by a *partial-batch* processing design, for the non-training layer to pro-

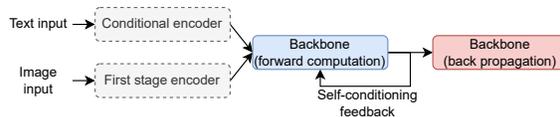


Figure 1. Training process of Stable Diffusion v2.1 (Rombach et al., 2022) and additional feedback of self-conditioning (Chen et al., 2022). Non-trainable components are marked in grey boxes.

Table 1. Ratio of the forward time of the non-trainable part to the forward and backward time of the trainable part on A100 GPU

Model / Batch size	8	16	32	64
Stable Diffusion v2.1	38%	41%	43%	44%
ControlNet v1.0	76%	81%	86%	89%

cess only a portion of a training batch. Partial-batch layer’s execution time can be precisely controlled by its input batch size, enabling it to be inserted into bubbles. In addition, partial-batch layers help eliminate the remaining idle time in pipeline bubbles after inserting non-trainable layers (processing a complete batch).

We implement DiffusionPipe and compare it to state-of-the-art data parallel training systems (Rasley et al., 2020) and ZeRO-3 (Rajbhandari et al., 2021), together with synchronous pipeline training paradigms, including SPP (Luo et al., 2022) and GPipe (Huang et al., 2019). Experimental results show that DiffusionPipe achieves up to 1.28x speedup over data parallel training and up to 1.41x speedup over existing pipeline parallel methods on representative diffusion models. We observe that DiffusionPipe achieves almost complete elimination of pipeline bubbles and effectively handles multiple training scenarios of diffusion models.

## 2 BACKGROUND AND MOTIVATION

### 2.1 Diffusion models and training

Diffusion models (Ho et al., 2020; Song et al., 2020; Chen et al., 2022; Rombach et al., 2022; Ho et al., 2022; Saharia et al., 2022; Podell et al., 2023) are generative models that learn to reverse the diffusion process that gradually turns data into noise. They typically comprise a backbone model that performs image generation and multiple frozen encoders that encode image and conditional information, e.g., class information (Yu et al., 2015), text description (Deng et al., 2009), canny edge (Canny, 1986) and human pose (Kreiss et al., 2021), and provide it as input to the backbone. During diffusion model training, the encoders are typically fixed and executed in advance in the forward computation pass (referred to as the non-trainable part), while the backbone (the trainable part) is trained with both forward computation and backward propagation (Fig. 1). Table 1 compares the execution time of the non-trainable part and the training time (forward and backward) of the trainable part.

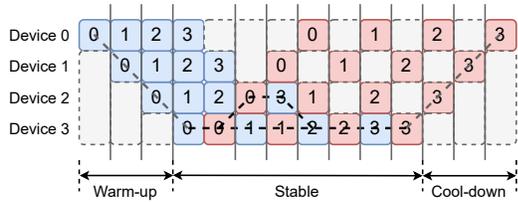


Figure 2. FIFO-1F1B schedule of a DNN. Gray blocks without numbers indicate pipeline bubbles. Potential critical paths are marked with a dashed line. Numbers indicate micro-batch index in both forward (blue) and backward (pink) steps.

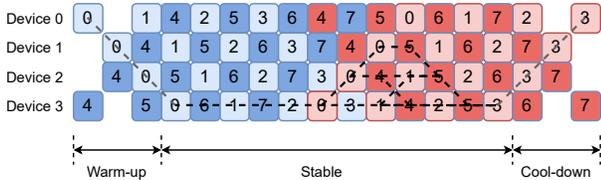


Figure 3. Bidirectional pipeline schedule of a DNN. Communication omitted. The same meaning of number and color with Fig. 2. Micro-batch 0 to 3 pipeline from device 0 to 3 (down direction), while micro-batch 4 to 7 pipeline from device 3 to device 0 (up direction).

Some diffusion models, e.g., Cascaded Diffusion Models (CDM) (Ho et al., 2022; Ramesh et al., 2022; Podell et al., 2023), involve multiple backbones of different capacities for high-resolution image generation. Multiple backbones accept the same encoder outputs, and each backbone also takes the output of the preceding backbone as input. The training of backbones in a CDM are typically independent, and each is trained on a different set of devices using the same procedure, as shown in Fig. 1.

In the current mainstream diffusion models, U-Net (Ho et al., 2020; Rombach et al., 2022) is widely used as the backbone model. Transformer models can also serve as the backbone (Peebles & Xie, 2022a; Bao et al., 2023). T5-xxl (Raffel et al., 2020), BERT (Devlin et al., 2018) and CLIP (Radford et al., 2021) text encoders are popular text encoders, while the image encoders are often variational auto-encoders (Kingma & Welling, 2013), ViT (Dosovitskiy et al., 2020) and CLIP image encoders. There are corresponding encoders (Zhang & Agrawala, 2023) for other modalities, such as canny edge and human pose.

Self-conditioning (Chen et al., 2022) has become a very popular technique for training diffusion models (Rombach et al., 2022; Saharia et al., 2022; Yuan et al., 2022), which improves the sampling quality by introducing an additional forward computation pass of the backbone (Fig. 1). The output of this forward pass is fed back to the backbone and serves as a conditional input. The fidelity of the image is then improved because each step is conditioned on the previously generated samples.

Table 2. Proportion of synchronization in training iteration time at local batch size 8 on A100 GPUs

Model / GPU count	8	16	32	64
Stable Diffusion v2.1	5.2%	19.3%	36.1%	38.1%
ControlNet v1.0	6.9%	22.7%	39.1%	40.1%

## 2.2 Pipeline parallel training, schedule and pipeline bubble

Pipeline parallel training partitions the model into stages, and each stage is deployed on a single device; the input data batch in each training iteration is divided into multiple micro-batches, which are processed through the model stages in a pipelined manner. The micro-batch execution pipelines are typically scheduled by a First-In-First-Out (FIFO) heuristic (Chen et al., 2015; Abadi et al., 2016; Sergeev & Del Balso, 2018), which executes micro-batches on model stages according to their ready order. The One-Forward-One-Backward (1F1B) schedule is widely adopted with FIFO, that alternatively executes forward computation and back propagation of micro-batches on each model stage in the stable phase of the pipeline execution (when multiple micro-batches are available to run on a model stage at the same time). As illustrated in Fig. 2, this schedule allows releasing intermediate activations and reduces peak memory usage by launching the backward computation when forward computation of the micro-batch is complete.

Chimera (Li & Hoefler, 2021) proposes bidirectional pipelining to reduce pipeline bubbles while retaining training synchronous. Chimera maintains two pipelines of micro-batch execution in different device rank orders (i.e., pipeline directions) on the same set of model stages, with the two pipeline execution schedules being symmetric along the device dimension. An example of bidirectional pipelining is shown in Fig. 3. Each micro-batch’s execution can fit perfectly into pipeline bubbles of its counterpart in the pipeline of the other direction (when the number of stages is even).

In synchronous pipeline training, pipeline bubbles generally exist in the pipeline schedule (Fig. 2). There is a barrier that gradient synchronization imposes between pipeline stages of the trainable part of diffusion models at different iterations, disabling pipeline bubbles be filled by the trainable part at different iterations. Therefore, although pipeline bubbles can be partially reduced by applying a better model partitioning and pipeline schedule, e.g., SPP (Luo et al., 2022) and Chimera (Li & Hoefler, 2021), such approaches cannot fundamentally eliminate pipeline bubbles, as they only manipulate the trainable part of the model and do not take the non-trainable part into consideration.

## 2.3 Synchronization overhead and memory consumption of data parallel training

Diffusion models are largely trained using data parallelism nowadays (Rombach et al., 2022; Ho et al., 2022; Saharia

Number of stages	Stable Diffusion v2.1				ControlNet v1.0			
	1	2	3	4	1	2	3	4
4	67.6%	51.0%	41.0%	34.3%	61.3%	44.2%	34.5%	28.4%
3	684.3%	342.2%	228.1%	171.1%	335.4%	167.7%	111.8%	83.9%
2	58.2%	41.0%	31.7%	25.8%	51.4%	34.5%	26.0%	20.9%
1	456.2%	228.1%	152.1%	114.1%	223.6%	111.8%	74.5%	55.9%
2	41.0%	25.8%	18.8%	14.8%	34.5%	20.9%	15.0%	11.7%
1	228.1%	114.1%	76.0%	57.0%	111.8%	55.9%	37.3%	28.0%

(a) Stable Diffusion v2.1

(b) ControlNet v1.0

Figure 4. Ratio of pipeline bubble time to iteration time (upper) and ratio of pipeline bubble time to non-trainable part execution time (lower) at batch size 64 using FIFO-1F1B scheduling.

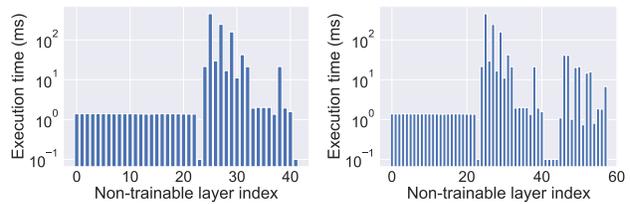
et al., 2022; Podell et al., 2023), which involves significant parameter synchronization overhead among devices and large memory consumption on each device that restricts the maximum feasible local batch size and the device utilization. For example, Stable Diffusion is trained at a *local* batch size of only 8 on each TPU-v3 (32GB) in (Rombach et al., 2022) consuming about 24.3 GB memory, which results in limited device utilization and exacerbates the synchronization portion of the training time. The synchronization overhead in Table 2 is computed as the ratio of parameter synchronization time to the end-to-end time of a training iteration. As the number of devices increases, parameter synchronization soon takes up a significant portion of the iteration time. In summary, the data parallel style of diffusion model training limits the training batch size and imposes high synchronization overhead.

## 2.4 Efficient pipeline bubble filling with non-trainable components

We profile the iteration training time of two popular diffusion models (without self-conditioning) by pipelining their backbones under different model stage and micro-batch number settings, and executing the non-trainable part using data parallelism before backbone training.

Fig. 4 shows the pipeline bubble ratios, where the iteration time is the sum of pipeline training time of the backbone and the execution time of the non-trainable part in each training iteration. Pipeline bubbles can take up to 68% of the overall training time, which is quite significant, according to the upper line in Fig. 4. In the lower line, a ratio close to 1 indicates that the pipeline bubble time can be almost completely filled by scheduling the non-trainable part in pipeline bubbles, under the respective model stage and micro-batch numbers. This observation motivates us to advocate pipeline bubble filling with the non-trainable part, and to study the detailed bubble filling strategies.

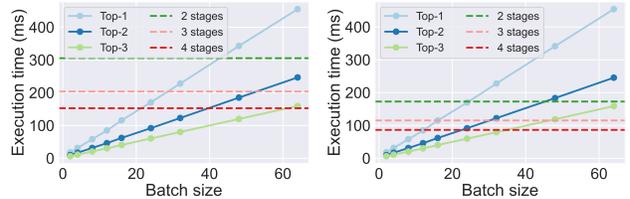
Fig. 5 shows that many non-trainable layers (indexed 0 to 21) in both models have short execution times, which belong to the frozen text encoder. Most layers (indexed 22 to 41) from the frozen image encoder take a moderate amount of time to compute (less than 30 ms). Such a distribution of



(a) Stable Diffusion v2.1

(b) ControlNet v1.0

Figure 5. Execution time of non-trainable layers at batch size 64.



(a) Stable Diffusion v2.1

(b) ControlNet v1.0

Figure 6. Execution time of top-3 non-trainable layers with longest execution time under different batch sizes, compared to longest pipeline bubble time when there are 4 micro-batches and different number of stages at batch size 64 using FIFO-1F1B scheduling.

non-trainable layers with a large proportion of short and moderately long layer execution times provides excellent opportunities for executing individual layers in pipeline bubbles ranging from 10 to 100 ms.

There are also some non-trainable layers with extra-long execution times (more than 400 ms), as shown in Fig. 5. Such layers may not fit into any pipeline bubble. Nevertheless, we observe that the layer execution time can be precisely controlled by adjusting the input batch size. Fig. 6 shows the execution times of the layers with the longest execution times at different batch sizes. When the batch size is reduced to 16, most of these non-trainable layers can fit into the longest pipeline bubble obtained by the way we identify bubbles in Fig. 2, implying that we can run such layers in pipeline bubbles by partitioning their input.

We seek to design an efficient algorithm to schedule the execution of non-trainable layers into pipeline bubbles.

## 3 SYSTEM DESIGN

Fig. 7 presents an overview of DiffusionPipe, which comprises of two modules: (1) The front-end carries out our workflow of generating an optimized pipeline training schedule for an input diffusion model, including pipeline training configurations of the backbone(s) and bubble-filling strategies of the non-trainable part; (2) The back-end is an execution engine that performs pipeline training according to the optimized pipeline schedule.

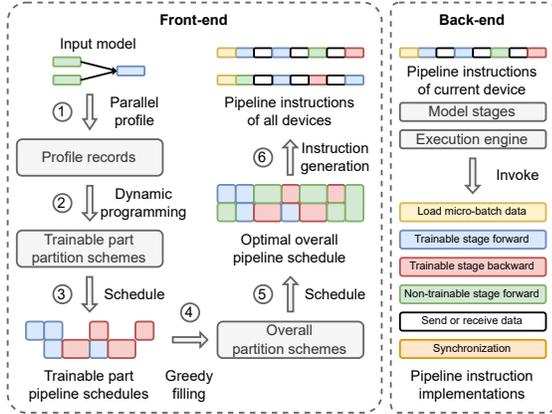


Figure 7. The architecture of DiffusionPipe

Table 3. Pipeline training hyper-parameters

Symbol	Description
$S$	Number of model stages
$M$	Number of micro-batches
$D$	Pipeline parallel group size <sup>1</sup>

### 3.1 Workflow

DiffusionPipe takes the diffusion model configuration, the training batch size, and the cluster configuration (i.e., number of machines and number of devices per machine) as inputs. DiffusionPipe first performs parallel profiling on the entire cluster to obtain the model layer execution time at different batch sizes (step 1 in Fig. 7), which is used in steps 2 to 5. Based on the input specifications, DiffusionPipe searches for pipeline training hyper-parameters as listed in Table 3. Note that DiffusionPipe supports mixed pipeline and data parallelism, as shown in Fig. 8. For each feasible hyper-parameter combination ( $S$ ,  $M$ , and  $D$ ), DiffusionPipe generates a near-optimal partitioning scheme for the trainable backbone(s) (§4, step 2), including the number of layers in each model stage and the number of devices on which each stage replicates. According to the corresponding pipeline schedule generated in step 3, DiffusionPipe further partitions the non-trainable part and fills it into pipeline bubbles (§5, step 4). Then DiffusionPipe generates the overall pipeline training schedules, and selects the optimal one with minimum iteration time (step 5). Finally, DiffusionPipe generates pipeline instructions for the back-end module according to the overall pipeline schedule (step 6).

### 3.2 Cross-iteration pipelining

For effective pipeline bubble filling that respects data dependencies between the non-trainable part and the backbone(s), DiffusionPipe advocates cross-iteration pipeline bubble fill-

<sup>1</sup>Pipeline parallel group is a minimum group of devices on which a complete set of pipeline communications is performed. In DiffusionPipe, pipeline parallel group size (i.e.,  $D$ ) = world size (i.e., number of devices in the cluster) / data parallel degree.

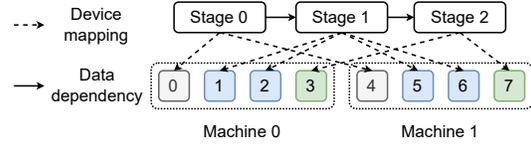


Figure 8. DiffusionPipe's data and pipeline parallelism. Devices in the same color run the same model stage.

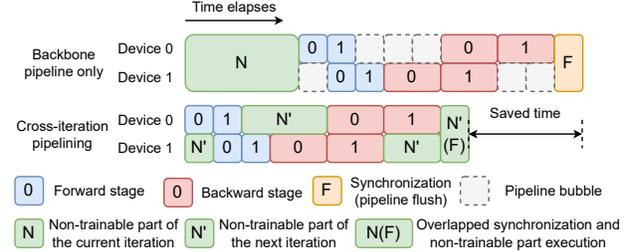


Figure 9. Cross-iteration pipelining of a diffusion model. Numbers indicate the micro-batch index of a pipeline stage.

ing, filling the bubble time of the backbone pipeline training in one iteration with the non-trainable part computation of the next iteration, as shown in Fig. 9. Non-trainable layers can be computed in a data parallel manner without pipelining, following their inter-layer data dependencies. At the end of a training iteration, the output of the non-trainable part is collected and divided into micro-batches according to the pipeline training configurations of the backbone(s). In the next iteration, these intermediate results are loaded onto the correct devices and fed as input to the pipeline training of the backbone(s). In addition, we only run the non-trainable part in the first iteration to enable such overlapping. The cross-iteration pipeline is mathematically equivalent to data parallel and synchronous pipeline training.

## 4 BACKBONE PARTITIONING

In this section, we present a unified dynamic programming approach to optimize partitioning and device assignment of the trainable part in diffusion models.

### 4.1 Single backbone

We first consider a diffusion model with a single backbone. The high-level idea is to analyze the critical path of FIFO-1F1B pipelining of the backbone and derive an upper bound on its execution time, to identify the optimal partitioning scheme that minimizes the execution time. We use the notations in Table 4.

FIFO pipeline execution can be divided into 3 phases, i.e., warm-up, stable and cool-down, as shown in Fig. 2. It launches micro-batch processing one by one in the warm-up phase and waits for all micro-batches to be completed in the cool-down phase. When we enlarge the last stage's execution time to the longest among all stages, enforcing it on the critical path, the warm-up phase contains forward

Table 4. Notations

Symbol	Definition
$L$	Number of layers in backbone model
$\mathbf{B}, B, b$	Training batch size, micro-batch size and number of samples in a partial-batch
$\mathbf{S}, s$	Set of model stages and model stage
$\mathbf{P}_l^f(B), \mathbf{P}_l^b(B)$	Forward and backward computation time of layer $l$ given batch size $B$
$\mathbf{C}_{l,l+1}^f(B), \mathbf{C}_{l+1,l}^b(B)$	Data size of communication in forward and backward pass between layers $l$ and $l + 1$ given batch size $B$
$\mathbf{R}_x, \mathbf{L}_x$	Bandwidth and latency of communication type $x$ (e.g., allreduce (ar), point-to-point (p2p))
$\mathbf{G}_l(B)$	Gradient size of layer $l$ given batch size $B$
$\mathbf{O}_l(B)$	Output size of layer $l$ given batch size $B$
$T_S(s)$	Synchronization time of stage $s$
$T_C(s)$	Compensation time of stage $s$
$T_0$	Maximum micro-batch execution time per stage or inter-stage communication time
$T_0^{S-C}$	Maximum gap between synchronization time and compensation time per stage
$T_B$	Length of a pipeline bubble (idle time)

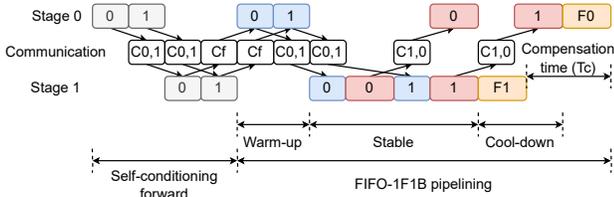


Figure 10. FIFO-1F1B scheduling of pipelining a backbone model with 2 stages, 2 micro-batches and self-conditioning. The same color and number setting with Fig. 9.  $C_{i,j}$  is communication from stage  $i$  to stage  $j$ .  $C_f$  feeds back the output of the backbone to stage 0.  $F_i$  refers to the parameter synchronization of stage  $i$ ,  $T_c$  is the compensation time of stage 1.

computation on  $S - 1$  model stages (aka *forward stages*). Similarly, the cool-down phase includes backward computation on  $S - 1$  model stages (aka *backward stages*). The stable phase of the critical path contains  $M$  forward stages and  $M$  backward stages, where  $M$  is the number of micro-batches. Therefore, there are a total of  $2(M + S - 1)$  forward and backward stages on the critical path of the FIFO-1F1B pipeline schedule in total. Considering the intermediate data communication between model stages in pipeline training, we add  $S - 1$  inter-stage communications in the forward and backward passes, respectively, which then becomes  $2(M + S - 1) + 2(S - 1)$  forward and backward stages, together with communications on the critical path.

We use  $T_0$  to denote the maximum of the time to run the forward *plus* backward computation of a micro-batch on a model stage, and the communication time between two stages, among all model stages. Then we have an upper bound  $T_0(M + 2S - 2)$  on the execution time of the critical path. We further consider the parameter synchronization

time among the micro-batches and add  $T_0^{S-C}$ , i.e., the maximum gap between  $T_S(s)$  and  $T_C(s)$  to the pipeline training time of the backbone for all stages  $s$ , where  $T_S(s)$  indicates the synchronization time of stage  $s$  and  $T_C(s)$  is used to compensate the overlapping time of parameter synchronization of stage  $s$  and computation of later stages. Fig. 10 gives an illustration. Putting the above together, an upper bound on the FIFO-1F1B pipeline execution time is:

$$T^{max} = T_0(M + 2S - 2) + T_0^{S-C} \quad (1)$$

We design a dynamic programming approach to identify the backbone partition and device assignment by minimizing  $T^{max}$ . We order the  $D$  devices in a pipeline parallel group into a chain according to their rank. Let  $W(L, S, r, D)$  denote  $T_0$  when partitioning the first  $L$  consecutive layers of the backbone into  $S$  stages, with these stages placed on devices 1 to  $D$  and the last stage  $s$  replicated on the last  $r$  devices (of the 1 to  $D$  device chain). Additionally, let  $Y(L, S, r, D)$  denote  $T_0^{S-C}$  under the same setting. The optimal partition of the backbone into  $S$  stages with the device placement of each stage can be computed by:

$$\min_{1 \leq r \leq D} \{(M + 2S - 2)W(L, S, r, D) + Y(L, S, r, D)\} \quad (2)$$

$W(L, S, r, D)$  can be decomposed into sub-problems that further partition the first  $l$  model layers into  $S - 1$  stages on the remaining  $D - r$  devices, with the last stage replicated on  $r'$  devices<sup>2</sup>. Then,  $W(L, S, r, D)$  can be computed by the maximum of  $W(l, S - 1, r', D - r)$  and the estimation of  $T_0$  by the last stage  $s$  (i.e.,  $T_0(s)$ ), and  $Y(L, S, r, D)$  can be computed in the same way, following Eqn. (3) to (8). Then we add the range in Eqn. (9) when optimizing Eqn. (2).

$$T_0(s) = \max \left\{ \sum_{l < i \leq L} \mathbf{P}_i^f \left( \frac{B}{r} \right) + \sum_{l < i \leq L} \mathbf{P}_i^b \left( \frac{B}{r} \right), \frac{\mathbf{C}_{l,l+1}^f \left( \frac{B}{r} \right) + \mathbf{C}_{l+1,l}^b \left( \frac{B}{r} \right)}{\mathbf{R}_{p2p}} + 2\mathbf{L}_{p2p} \right\} \quad (3)$$

$$T_S(s) = \sum_{l < i \leq L} \mathbf{G}_i \left( \frac{B}{r} \right) / \mathbf{R}_{ar} + \mathbf{L}_{ar} \quad (4)$$

$$T_C(s) = \sum_{l < i \leq L} \mathbf{P}_i^b \left( \frac{B}{r} \right) \quad (5)$$

$$T_0^{S-C}(s) = T_S(s) - T_C(s) \quad (6)$$

$$W(L, S, r, D) = \max \{ W(l, S - 1, r', D - r), T_0(s) \} \quad (7)$$

$$Y(L, S, r, D) = \max \{ Y(l, S - 1, r', D - r), T_0^{S-C}(s) \} \quad (8)$$

$$\forall l, r', \quad 1 \leq l \leq L - 1, \quad 1 \leq r' \leq D - r \quad (9)$$

Here  $B$  is the micro-batch size,  $\mathbf{P}_i^{\{f/b\}} \left( \frac{B}{r} \right)$  is the forward / backward computation time of layer  $i$  given local batch

<sup>2</sup>Though we support different model stages using different data parallel degrees (e.g.,  $r \neq r'$ ), we find that such cases are rare. They can result in strange bubble filling schemes (§5) and require complex implementations. In evaluation (§6), we force all stages to have the same data parallel degree.

size  $\frac{B}{r}$ .  $\mathbf{C}_{l,l+1}^f(\frac{B}{r})$  and  $\mathbf{C}_{l+1,l}^b(\frac{B}{r})$  are data sizes in forward and backward pass between layer  $l$  and  $l+1$  given local batch size  $\frac{B}{r}$ .  $\mathbf{G}_i(\frac{B}{r})$  is the gradient size of layer  $i$  given local batch size  $\frac{B}{r}$ .  $\mathbf{R}_x$  and  $\mathbf{L}_x$  are bandwidth and latency of communication type  $x$ , while  $ar$  indicates all-reduce used in synchronization and  $p2p$  indicates point-to-point communication between model stages.

Note that a sub-problem in Eqn. (8) does not know the partition scheme of all subsequent layers (with indices greater than  $l$ ) for computing the compensation time  $T_C$ . Instead, we use a lower bound of  $T_C$  in Eqn. (5), i.e., the sum of the backward computation time of all these layers on  $r$  devices.

## 4.2 Multiple backbones

For a cascaded diffusion model with multiple backbones, we advocate bidirectional pipelining to train the backbones on the same set of devices (instead of using a separate set of devices to train each backbone), in order to utilize the devices more efficiently. In particular, we leverage bidirectional pipelining (Li & Hoefler, 2021) to train multiple backbones, with each backbone pipelining in different direction. Here, we consider pipelining from low-rank device to high-rank device as the down direction and vice versa, and the corresponding pipelines are down and up pipelines.

Consider 2 backbones in a CDM. As shown in Fig. 3, the duration of the stable phase of the critical path in bidirectional pipelining differs from unidirectional pipelining while the duration of the warm-up and cool-down phases are not affected. We calculate the number of paired forward and backward stages between the down and up pipelines ( $M_{CDM}$ ), and derive an upper bound on bi-directional pipeline execution time for training two backbones:

$$T_{0,CDM} = \max\{T_{0,down}, T_{0,up}\} \quad (10)$$

$$T_{0,CDM}^{S-C} = \max\{T_{0,down}^{S-C}, T_{0,up}^{S-C}\} \quad (11)$$

$$T_{CDM}^{max} = (M_{CDM} + 2S - 2)T_{0,CDM} + T_{0,CDM}^{S-C} \quad (12)$$

Here  $T_{0,\{down/up\}}$  is the maximum of the time to perform the forward and backward computation of a micro-batch and the communication time between two stages in the down or up pipeline among all model stages. The sub-problem in the dynamic programming approach in bi-directional pipelining should decide partitioning and placement of model stages for both backbones. Let  $W(L_d, L_u, S, r, D)$  denote  $T_{0,CDM}$  when partitioning the last  $L_d$  and the first  $L_u$  consecutive layers of the down- and up-pipelined backbones, respectively, into  $S$  stages, while placing them on  $D$  devices and replicating the last stage  $s_d$  and the first stage  $s_u$  of the two backbones on the last  $r$  devices of the 1 to  $D$  device chain, and we have  $Y(L_d, L_u, S, r, D)$  similarly. The optimal partitioning of two backbones can be computed by:

$$\min_{1 \leq r \leq D} \{(M_{CDM} + 2S - 2)W(L_d, L_u, S, r, D) + Y(L_d, L_u, S, r, D)\} \quad (13)$$

In Eqn. (14) and (15) we give the definition of  $W(L_d, L_u, S, r, D)$  and  $Y(L_d, L_u, S, r, D)$ . Eqn. (16) presents the additional optimization range. Communication in the bidirectional pipelines may compete for resources, and we reasonably enlarge the communication time in Eqn. (3) by a factor of 2 (as there are two pipelining directions). Other equations are the same with §4.1.

$$W(L_d, L_u, S, r, D) = \max\{W(l_d, l_u, S - 1, r', D - r), T_0(s_d), T_0(s_u)\} \quad (14)$$

$$Y(L_d, L_u, S, r, D) = \max\{Y(l_d, l_u, S - 1, r', D - r), T_0^{S-C}(s_d), T_0^{S-C}(s_u)\} \quad (15)$$

$$\forall l_{d/u}, r', \quad 1 \leq l_{\{d/u\}} \leq L_{\{d/u\}} - 1, \\ 1 \leq r' \leq D - r \quad (16)$$

For a diffusion model with more than two backbones, we can divide the backbones into two groups, one to be pipelined in each direction. We then combine stages of the backbones in the same pipeline direction to form a larger model stage and apply our design for bi-directional pipelining accordingly.

## 4.3 Backbone(s) with self-conditioning

DiffusionPipe performs self-conditioning on the same device to eliminate unnecessary parameter storage and updating, as shown in Fig. 10. There is an additional forward pass at each model stage, and Eqn. (3) is changed to:

$$T_{0,SC}(s) = \max\{2 \sum_{l < i \leq L} \mathbf{P}_i^f(\frac{B}{r}) + \sum_{l < i \leq L} \mathbf{P}_i^b(\frac{B}{r}), \frac{2\mathbf{C}_{l,l+1}^f(\frac{B}{r}) + \mathbf{C}_{l+1,l}^b(\frac{B}{r})}{\mathbf{R}_{p2p}} + 3\mathbf{L}_{p2p}\} \quad (17)$$

In addition, the communication time for sending the output from the last stage to the first stage ( $C_f$  in Fig. 10) should be considered. We use a point-to-point transmission time as the upper bound of this feedback time:  $T_F = \mathbf{O}_L(B)/\mathbf{R}_{p2p} + \mathbf{L}_{p2p}$ , where  $\mathbf{O}_L(\frac{B}{r})$  is the output size of the last layer  $L$  at local batch size  $\frac{B}{r}$ . The upper bound on the pipeline execution time with self-conditioning is:

$$T_{SC}^{max} = (M + 2S - 2)T_{0,SC} + T_0^{S-C} + T_F \quad (18)$$

The dynamic programming formulation remains the same as in §4.1. Since self-conditioning is usually randomly activated during the training process with a certain probability  $p$  (0.5 in (Chen et al., 2022)), the formulation optimizes an expectation of  $T_{SC}^{max}$  and  $T^{max}$ .

In case self-conditioning is applied to CDMs, we can readily extend the formulation in §4.2 by counting the additional number of forward stages in the critical path.

## 5 PIPELINE BUBBLE FILLING

In DiffusionPipe, we divide the pipeline idle time along the timeline and define a pipeline bubble using a tuple (*start time, end time, idle devices*) so that a bubble contains the

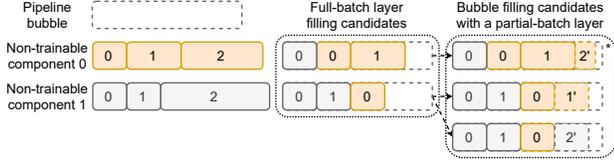


Figure 11. Full-batch bubble filling candidates and corresponding candidates with a partial-batch layer. Numbers indicate the index of the non-trainable layer. 1' and 2' denote partial-batch layers. \* marks the candidate with the longest execution time.

same number of idle devices in its time span. For example, in Fig. 2, the first pipeline bubble is in the first time slot with idle devices 1 to 3. Pipeline bubble filling is always performed under the cross-iteration style of pipelining (§3.2), regardless of the number of backbones and whether self-conditioning is applied. We also define a partial-batch layer using a tuple (*component index, layer index, number of samples in partial-batch*).

Non-trainable components in a diffusion model may have inter-dependencies (e.g., ControlNet (Zhang & Agrawala, 2023)), and layers within each component are linearly dependent. We schedule the execution of non-trainable components in pipeline bubbles following a topological order of components according to their dependencies. Especially, we fill in the pipeline bubbles sequentially in their chronological order<sup>3</sup>. To fill a pipeline bubble, we consider all of the components that are ready at the time, i.e., their dependencies are resolved. Whenever a component becomes ready, we add it to the set of ready components.

An efficient algorithm is designed to fill a pipeline bubble with ready components (Alg. 1). Its input mainly includes the bubble time  $T_B$  and the number of idle devices  $d$ , a list  $u$  containing the index of the starting layer in each currently ready component (layers of a component can be executed in multiple bubbles). It first finds candidates  $K$  containing full-batch layers of ready components to fill the current pipeline bubble (Alg. 2), whose execution completes within the bubble time. Then it adds at most one layer from a component to be executed on a partial batch in the remaining bubble time to each candidate, and finally it produces the optimal bubble filling scheme with the longest execution time (not exceeding the bubble time), as shown in Fig. 11. Note that the component layers assigned to the bubble are executed in a data parallel manner at local batch size  $\frac{B}{d}$ .

Input of Alg. 2 includes the input of Alg. 1, and the component index  $i$  that it focuses on. It finds bubble filling candidates containing full-batch layers in a recursive man-

<sup>3</sup>For bubble filling efficiency, we only identify pipeline bubbles longer than 10 ms, which is empirically greater than the cost of setting up inputs and outputs for pipeline bubble filling. Chronological order of pipeline bubbles is achieved by analyzing the pipeline schedule, which is simulated using the profiled results obtained in step 1 of Fig. 7. All proposed algorithms work offline only.

### Algorithm 1 Filling One Pipeline Bubble

**Input:** Number of ready non-trainable components  $n$ , training batch size  $B$ , pipeline bubble time  $T_B$ , number of idle devices  $d$ , indices of starting layers of components  $u$  (list, length is  $n$ ), numbers of layers of components  $L$  (list, length is  $n$ )

**Output:** Optimal bubble filling candidate  $k^*$

```

1:  $K_0, K \leftarrow \text{emptyList}(), \text{FFC}(n, B, T_B, d, u, L, 0)$ 
2: for  $k$  in  $K$ ,  $h$  in  $0, \dots, n-1$  do
3:    $b \leftarrow \text{maximum of getValidNumSamples}(B, d), \text{s.t.}, ^4$ 
4:    $T_B \geq \sum_{i \in [n], j \in [k_i]} \mathbf{P}_{i, u_i + j}^f(\frac{B}{d}) + \mathbf{P}_{h, u_h + k_h}^f(\frac{b}{d})$ 
   // Bubble time should be greater than the sum of execution
   // time of candidate  $k$  and a partial-batch layer
5:    $K_0.\text{append}((k, (h, u_h + k_h, b)))$  // Add candidate  $k$  enhanced
   // with a partial-batch layer  $(h, u_h + k_h, b)$ 
6: end for
7: return the candidate in  $K_0$  with the longest execution time
    
```

### Algorithm 2 FFC - Full-batch Layer Bubble Filling Candidates

**Input:**  $n, B, T_B, d, u, L$ , current component index  $i$

**Output:** bubble filling candidates  $K$

```

1:  $t, k_0, K \leftarrow 0, 0, \text{emptyList}()$ 
2: while  $t + \mathbf{P}_{i, u_i + k_0}^f(\frac{B}{d}) \leq T_B$  and  $u_i + k_0 < L_i$  do
3:    $t \leftarrow t + \mathbf{P}_{i, u_i + k_0}^f(\frac{B}{d})$  // Cumulative execution time
4:    $k_0 \leftarrow k_0 + 1$ 
5: end while
6: if  $i = n-1$  then
7:   return  $[[k_0]]$  // Add all  $k_0$  layers to the candidate as it is
   // the last component
8: else
9:   for  $k$  in  $k_0, \dots, 0$  do
10:     $T_B' \leftarrow T_B - \sum_{h \in [k]} \mathbf{P}_{i, u_i + h}^f(\frac{B}{d})$  // Remaining bubble
    // time after adding  $k$  layers to the candidate
11:     $K' \leftarrow \text{FFC}(n, B, T_B', d, u, L, i+1)$ 
12:     $K.\text{extend}([\text{concat}([k], k') \text{ for } k' \text{ in } K'])$ 
13:   end for
14: return  $K$ 
15: end if
    
```

ner: assuming layers from components with indices smaller than  $i$  are already considered, it adds layers from component  $i$  to the candidate. Alg. 2 first computes how many layers can be added at most from line 2 to 5, where  $\mathbf{P}_{i, u_i + k}^f(\frac{B}{d})$  is the computation time of layer  $u_i + k$  of component  $i$  given local batch size  $\frac{B}{d}$ . Then it adds different numbers of layers to the candidate (with total execution time not exceeding  $T_B$ ), and recursively calls itself to add layers from the next component  $i+1$  from line 9 to 13. Alg. 2 returns a list  $K$  containing bubble filling candidates, where each candidate is a list containing  $n$  elements ( $n$  is the number of ready components), with each element containing the indices of the layers of that component to be executed in the bubble.

Then for each bubble filling candidate  $k$  in  $K$ , an additional partial-batch layer is added to it. Especially, we identify a layer that is the subsequent layer following the scheduled

<sup>4</sup>Here  $[n] := \{0, 1, \dots, n-1\}$ ,  $[k_i] := \{0, 1, \dots, k_i-1\}$ .

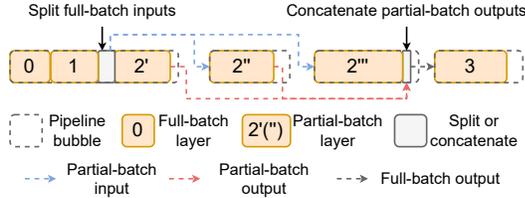


Figure 12. Input split and output concatenation of partial-batch layer’s processing among pipeline bubbles. The partial-batch layer 2 of a non-trainable component is scheduled in 3 consecutive pipeline bubbles.

full-batch layers in the candidate, as well as a partial batch to process, whose execution can occupy the longest of the remaining bubble time (line 3 to 4 of Alg. 1). We then choose among the enhanced bubble filling candidates with a partial-batch layer each, the one achieving the longest execution time to maximally utilize the idle time.

To decide the partial batch for the extra layer to process in a pipeline bubble (function `getValidNumSamples` in line 3 of Alg. 1), we follow two principles: (1) The local batch size  $b/d$  should not be too small, as otherwise the benefit of inserting a partial-batch layer will not compensate for the overhead of handling its input and output, as illustrated in Fig. 12; (2)  $b/d$  should be a *regular* value to avoid potential kernel performance degradation at unusual batch sizes. We empirically use 4, 8, 12, 16, 24, 32, 48, 64 and 96 as the local batch size candidates. If pipeline bubbles cannot completely accommodate the non-trainable part, the remaining part will be executed after pipelining completes.

Furthermore, after introducing a partial-batch layer ( $h, u_h + k_h, b$ ) in a pipeline bubble (line 5 of Alg. 1), the layer  $u_h + k_h$  of component  $h$  is the first ready layer of that component to be considered when filling the following pipeline bubbles, and it is treated as a full-batch layer on the remaining batch. In this way, this layer can be scheduled to process all or part of the remaining batch in a subsequent pipeline bubble. Fig 12 shows an example of scheduling part of the remaining batch in a subsequent (i.e., the second) pipeline bubble.

## 6 EVALUATION

We build DiffusionPipe on PyTorch 2.0.1 and CUDA 11.7 with 20k LoC in Python, and integrate it with DeepSpeed 0.8.3 to support pipeline and data parallel training. Communication operations are implemented using PyTorch’s distributed communication package and NCCL 2.17.1. Though DiffusionPipe is integrated into DeepSpeed, it is easy to migrate DiffusionPipe to other vendor frameworks. We only need to switch to the new launching method and replace the communication and optimizer implementations with corresponding implementations.

**Test-bed** We conduct our experiments on a cluster of 8

Amazon EC2 p4de.24xlarge machines, each containing 8 NVIDIA A100-80GB GPUs and 96 vCPU cores. The inter-node connection (EFA) bandwidth is 400 Gbps and the intra-node connection (NVSwitch) bandwidth is 600 GBps.

**Models** We train these models: Stable-Diffusion (Rombach et al., 2022) v2.1, ControlNet (Zhang & Agrawala, 2023) v1.0, CDM-LSUN and CDM-ImageNet (Ho et al., 2022). For CDM-LSUN, we train its 2 backbones using bi-directional pipelining. For CDM-ImageNet, we only train its second and third backbones because training all of them will exceed the GPU memory. The backbones of the same CDM are trained under the same batch size. The input configurations of all models (Table 5) are the same as in their original papers.

Table 5. Diffusion models and training configurations

Model	Input shape	Self-cond
Stable Diffusion v2.1	512x512	Enabled
ControlNet v1.0		
CDM-LSUN	64x64, 128x128	Not enabled
CDM-ImageNet	(2 image inputs)	

**Baselines** We run DeepSpeed (Rasley et al., 2020) with vanilla distributed data parallelism (DDP) and ZeRO-3 (Rajbhandari et al., 2021) as baselines for data parallel training. We use GPipe (Huang et al., 2019) and SPP (Luo et al., 2022) as baselines of pipeline parallelism, which perform the backbone only pipelining in Fig. 9. For GPipe that partitions a model into stages with equal number of layers, we evaluate it with 2 pipeline stages and 4 micro-batches. For SPP that solves a dynamic programming problem to optimize model partitioning, we perform the same hyper-parameter searching as in DiffusionPipe. When self-conditioning is enabled, we also run the extra-forward part in the way shown in Fig. 10 for pipeline parallel baselines. Bubble filling is not performed for pipeline parallel baselines.

For cascaded diffusion models, data parallel training is performed in two ways: (1) Training multiple backbones in sequential using all devices, i.e., DeepSpeed(-ZeRO-3)-S; (2) Training multiple backbones in parallel on evenly partitioned sets of devices, i.e., DeepSpeed(-ZeRO-3)-P, which is the default strategy in many CDM works. Both SPP and GPipe do not apply to CDM, because they do not support pipelining of multiple models.

**Metrics** We present the training throughput in terms of the number of samples processed per second. The throughput of DeepSpeed(-ZeRO-3)-S and DeepSpeed(-ZeRO-3)-P is computed by  $\frac{\text{total batch size of all backbones}}{\text{sum of iteration time of all backbones}}$  and sum of  $\frac{\text{batch size}}{\text{iteration time}}$  of all backbones, respectively. We also present the pipeline bubble ratio of DiffusionPipe and pipelined baselines, which is computed by  $\frac{\sum_{b \in \text{pipeline bubbles}} T_b \times d_b}{\text{iteration\_time} \times \text{total\_num\_devices}}$ , where  $T_b$  and  $d_b$  are the duration and the number of idle devices of the bubble  $b$ .

## DiffusionPipe: Training Large Diffusion Models with Efficient Pipelines

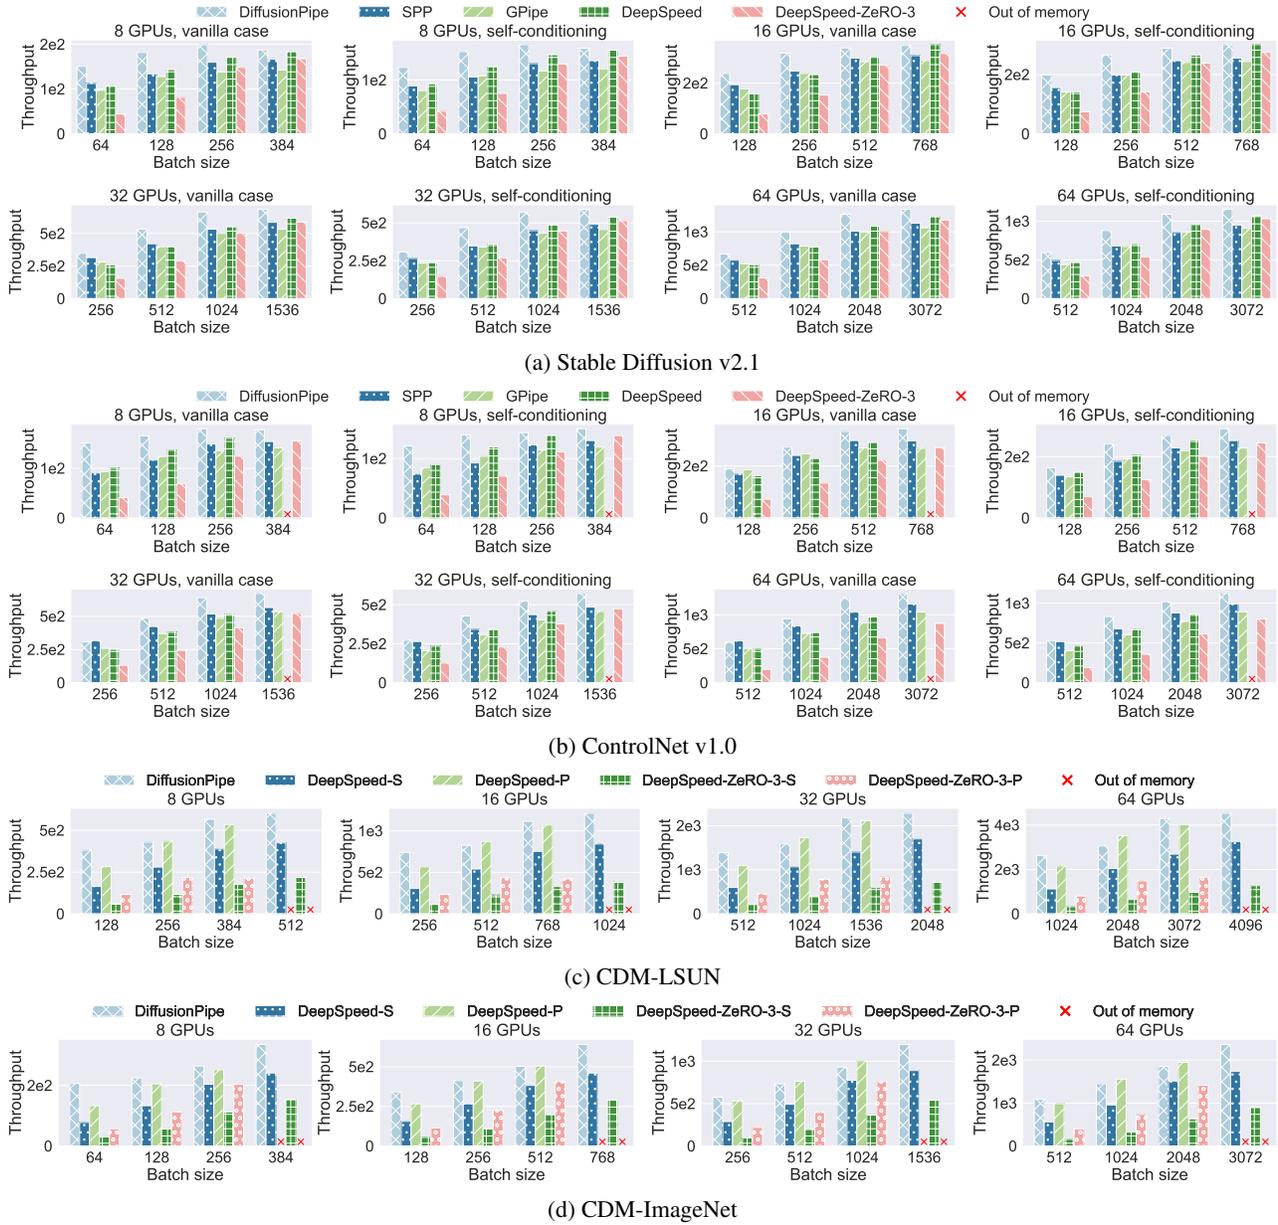


Figure 13. Training throughput (samples/second)

### 6.1 Training throughput

In Fig. 13 we present the throughput of training the diffusion models at different cluster scales and batch sizes.

For single backbone models (Fig. 13a and 13b), DiffusionPipe outperforms other pipeline systems both with and without self-conditioning, as it efficiently fills pipeline bubbles with non-trainable layer execution. When training on a machine, device utilization determines performance. DiffusionPipe can outperform data parallel baselines because both trainable model stages and the non-trainable part occupy only part of the cluster, and it processes the input batch with a larger local batch size, thus achieving better device

utilization. At batch size 256, DiffusionPipe achieves 1.44x and 1.16x speedups over GPipe and DeepSpeed, respectively, when training Stable Diffusion v2.1.

When training on multiple machines, synchronization overhead has a more significant impact on training throughput. DiffusionPipe outperforms data parallel baselines as it can mitigate the overhead in two ways: (1) Each device hosts fewer parameters in pipeline training, so less synchronization communication is required; (2) Synchronization can be overlapped with the non-trainable part, further reducing its impact on the throughput. At batch size 2048 on 64 GPUs, DiffusionPipe achieves 1.41x and 1.28x speedups over GPipe and DeepSpeed training of ControlNet v1.0.

For the cascaded diffusion models (Fig. 13c and Fig. 13d), DiffusionPipe’s throughput is comparable to DeepSpeed-P for two reasons: (1) In both CDM models, there is little non-trainable part to fill bubbles, so we cannot get speedup from the non-trainable part; (2) Backbone sizes in both CDMs are relatively close to each other, and DeepSpeed-S already achieves balanced training iteration time with respect to backbones. However, DiffusionPipe can still achieve a higher training batch size compared to DeepSpeed-P because the activation memory of micro-batches does not persist during the entire backward process.

### 6.2 Pipeline bubble ratio

In Fig. 14, we observe that DiffusionPipe can reduce the pipeline bubble ratio to less than 5% for both Stable Diffusion v2.1 and ControlNet v1.0, which is dramatically lower than other pipeline training baselines. The unfilled pipeline bubble time can be explained by: (1) The difference between the actual execution time and the profiled execution time (used to drive the bubble-filling algorithm); (2) The non-continuous execution times of non-trainable layers, which make it unlikely to perfectly fill the bubble.

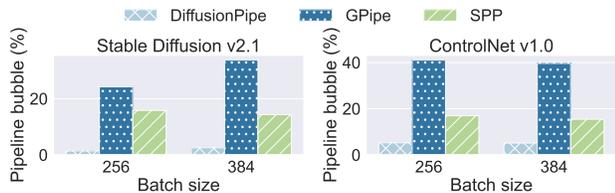


Figure 14. Pipeline bubble ratio on 8 GPUs

### 6.3 Ablation study

In Fig. 15, we evaluate the throughput of DiffusionPipe when the partial-batch layer design is disabled and when the pipeline bubble filling design is completely disabled, respectively. We observe that disabling the partial-batch layer significantly degrades throughput, and disabling bubble filling degrades it even more (by 10.9% and 17.6% for ControlNet v1.0 at batch size 256). This demonstrates that pipeline bubble filling and the partial-batch layer design can effectively improve training efficiency. We also observe that at batch size 384, disabling the partial-batch layer achieves almost the same throughput as no bubble filling, indicating that the extra-long layer in Fig. 5 blocks almost all layers during bubble filling, and validating our partial-batch design.

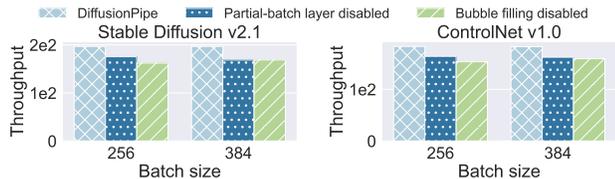


Figure 15. Ablation study on 8 GPUs (samples/second)

### 6.4 Pre-processing overhead

Pre-processing, including profiling, running the model partitioning and the pipeline bubble filling algorithm, is performed once and completes within a few minutes, which is acceptable given training usually takes much longer time.

Profiling is executed in parallel on all GPUs, and its overhead is decided by the number of GPUs. A typical profiling time of Stable Diffusion v2.1 on 2 AWS EC2 p4de.24xlarge machines at batch size 512 is 55 seconds.

Model partitioning algorithm is executed in parallel on all CPUs in the host machine, and its overhead is decided by the number of CPUs in the host, the number of trainable components and the number of layers in them. For Stable Diffusion v2.1 and ControlNet v1.0 at the same setting, the overhead is about 0.5 second.

Pipeline bubble fulling algorithm is executed on only 1 CPU, and its overhead is decided by the number of pipeline bubbles and the number of non-trainable components. For the same models at the same setting, the overhead is less than 1 second.

## 7 CONCLUSION

This paper presents DiffusionPipe, a system that automatically optimizes pipeline training for large diffusion models. Our unified partitioning algorithm for the trainable part optimizes partitioning schemes of multiple training scenarios of diffusion models. We also propose to fill pipeline bubbles with the non-trainable part of diffusion models, which achieves higher training throughput compared to pipelining only the backbone model and training in data parallel. Experimental results demonstrate that DiffusionPipe achieves speedups of up to 1.41x compared to pipeline baselines and 1.28x compared to data parallel baselines. This is accomplished by reducing the pipeline bubble to less than 5% of the training iteration time. Moreover, DiffusionPipe enables the use of larger training batch sizes in comparison to data parallel baselines. Our design of filling pipeline bubbles with non-trainable parts can extend to more applications, e.g., training or fine-tuning diffusion models with transformer backbones (Peebles & Xie, 2022b; Ma et al., 2024; Chen et al., 2023; 2024), together with multimodal models with frozen encoder components (Li et al., 2023a;b; Yu et al., 2023).

## 8 ACKNOWLEDGEMENT

We would like to thank the Program Chairs and anonymous reviewers for their valuable feedback. This work was supported by an Amazon Research Award (ARA) on AWS AI and grants from Hong Kong RGC under the contracts HKU 17208920, 17204423 and C7004-22G (CRF).

## REFERENCES

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. Tensorflow: a system for large-scale machine learning. In *12th USENIX symposium on operating systems design and implementation (OSDI 16)*, pp. 265–283, 2016.
- Bao, F., Nie, S., Xue, K., Cao, Y., Li, C., Su, H., and Zhu, J. All are worth words: A vit backbone for diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 22669–22679, 2023.
- Bian, Z., Liu, H., Wang, B., Huang, H., Li, Y., Wang, C., Cui, F., and You, Y. Colossal-ai: A unified deep learning system for large-scale parallel training. *arXiv preprint arXiv:2110.14883*, 2021.
- Canny, J. F. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8:679–698, 1986.
- Chen, J., Yu, J., Ge, C., Yao, L., Xie, E., Wu, Y., Wang, Z., Kwok, J. T., Luo, P., Lu, H., and Li, Z. Pixart- $\alpha$ : Fast training of diffusion transformer for photorealistic text-to-image synthesis. *ArXiv*, abs/2310.00426, 2023.
- Chen, J., Wu, Y., Luo, S., Xie, E., Paul, S., Luo, P., Zhao, H., and Li, Z. Pixart- $\delta$ : Fast and controllable image generation with latent consistency models. *ArXiv*, abs/2401.05252, 2024.
- Chen, T., Li, M., Li, Y., Lin, M., Wang, N., Wang, M., Xiao, T., Xu, B., Zhang, C., and Zhang, Z. Mxnet: A flexible and efficient machine learning library for heterogeneous distributed systems. *arXiv preprint arXiv:1512.01274*, 2015.
- Chen, T., Zhang, R., and Hinton, G. Analog bits: Generating discrete data using diffusion models with self-conditioning. *arXiv preprint arXiv:2208.04202*, 2022.
- Choi, J., Kim, S., Jeong, Y., Gwon, Y., and Yoon, S. Ilvr: Conditioning method for denoising diffusion probabilistic models. *arXiv preprint arXiv:2108.02938*, 2021.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255, 2009.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Falcon, W. and The PyTorch Lightning team. PyTorch Lightning, March 2019. URL <https://github.com/Lightning-AI/lightning>.
- Fan, S., Rong, Y., Meng, C., Cao, Z., Wang, S., Zheng, Z., Wu, C., Long, G., Yang, J., Xia, L., et al. Dapple: A pipelined data parallel approach for training large models. In *Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pp. 431–445, 2021.
- Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- Ho, J., Saharia, C., Chan, W., Fleet, D. J., Norouzi, M., and Salimans, T. Cascaded diffusion models for high fidelity image generation. *J. Mach. Learn. Res.*, 23(47):1–33, 2022.
- Huang, Y., Cheng, Y., Bapna, A., Firat, O., Chen, D., Chen, M., Lee, H., Ngiam, J., Le, Q. V., Wu, Y., et al. Gpipe: Efficient training of giant neural networks using pipeline parallelism. *Advances in neural information processing systems*, 32, 2019.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Kreiss, S., Bertoni, L., and Alahi, A. Openpipaf: Composite fields for semantic keypoint detection and spatio-temporal association. *IEEE Transactions on Intelligent Transportation Systems*, 23(8):13498–13511, 2021.
- Li, J., Pan, K., Ge, Z., Gao, M., Zhang, H., Ji, W., Zhang, W., Chua, T.-S., Tang, S., and Zhuang, Y. Fine-tuning multi-modal llms to follow zero-shot demonstrative instructions. 2023a.
- Li, J. Y., Liu, C., Cheng, S., Arcucci, R., and linda Qiao. Frozen language model helps ecg zero-shot learning. *arXiv preprint arXiv:2303.12311*, 2023b.
- Li, S. and Hoefler, T. Chimera: efficiently training large-scale neural networks with bidirectional pipelines. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1–14, 2021.
- Luo, Z., Yi, X., Long, G., Fan, S., Wu, C., Yang, J., and Lin, W. Efficient pipeline planning for expedited distributed dnn training. In *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*, pp. 340–349. IEEE, 2022.

- Ma, N., Goldstein, M., Albergo, M. S., Boffi, N. M., Vandeneijnden, E., and Xie, S. Sit: Exploring flow and diffusion-based generative models with scalable interpolant transformers. *ArXiv*, abs/2401.08740, 2024.
- Narayanan, D., Harlap, A., Phanishayee, A., Seshadri, V., Devanur, N. R., Ganger, G. R., Gibbons, P. B., and Zaharia, M. Pipedream: Generalized pipeline parallelism for dnn training. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles*, pp. 1–15, 2019.
- Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., and Chen, M. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021.
- Peebles, W. and Xie, S. Scalable diffusion models with transformers. *arXiv preprint arXiv:2212.09748*, 2022a.
- Peebles, W. S. and Xie, S. Scalable diffusion models with transformers. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 4172–4182, 2022b.
- Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J., Penna, J., and Rombach, R. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.
- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020.
- Rajbhandari, S., Ruwase, O., Rasley, J., Smith, S., and He, Y. Zero-infinity: Breaking the gpu memory wall for extreme scale deep learning. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1–14, 2021.
- Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., and Chen, M. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- Rasley, J., Rajbhandari, S., Ruwase, O., and He, Y. Deep-speed: System optimizations enable training deep learning models with over 100 billion parameters. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 3505–3506, 2020.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E., Ghasemipour, S. K. S., Ayan, B. K., Mahdavi, S. S., Lopes, R. G., et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022.
- Sergeev, A. and Del Balso, M. Horovod: fast and easy distributed deep learning in tensorflow. *arXiv preprint arXiv:1802.05799*, 2018.
- Song, J., Meng, C., and Ermon, S. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.
- von Platen, P., Patil, S., Lozhkov, A., Cuenca, P., Lambert, N., Rasul, K., Davaadorj, M., and Wolf, T. Diffusers: State-of-the-art diffusion models, 2022. URL <https://github.com/huggingface/diffusers>.
- Wu, C., Yin, S., Qi, W., Wang, X., Tang, Z., and Duan, N. Visual chatgpt: Talking, drawing and editing with visual foundation models. *arXiv preprint arXiv:2303.04671*, 2023.
- Yu, F., Seff, A., Zhang, Y., Song, S., Funkhouser, T., and Xiao, J. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015.
- Yu, L., Cheng, Y., Wang, Z., Kumar, V., Macherey, W., Huang, Y., Ross, D. A., Essa, I., Bisk, Y., Yang, M.-H., et al. Spae: Semantic pyramid autoencoder for multimodal generation with frozen llms. *arXiv preprint arXiv:2306.17842*, 2023.
- Yuan, H., Yuan, Z., Tan, C., Huang, F., and Huang, S. Seqdif-fuseq: Text diffusion with encoder-decoder transformers. *arXiv preprint arXiv:2212.10325*, 2022.
- Zhang, L. and Agrawala, M. Adding conditional control to text-to-image diffusion models. *arXiv preprint arXiv:2302.05543*, 2023.