

SWIFT: Expedited Failure Recovery for Large-scale DNN Training

Yuchen Zhong, Guangming Sheng, Juncheng Liu, Jinhui Yuan, and Chuan Wu, *Senior Member, IEEE*

Abstract—As the size of deep learning models gets larger and larger, training takes longer time and more resources, making fault tolerance more and more critical. Existing state-of-the-art methods like CheckFreq and Elastic Horovod need to back up a copy of the model state (i.e., parameters and optimizer states) in memory, which is costly for large models and leads to non-trivial overhead. This paper presents SWIFT, a novel recovery design for distributed deep neural network training that significantly reduces the failure recovery overhead without affecting training throughput and model accuracy. Instead of making an additional copy of the model state, SWIFT resolves the inconsistencies of the model state caused by the failure and exploits the replicas of the model state in data parallelism for failure recovery. We propose a logging-based approach when replicas are unavailable, which records intermediate data and replays the computation to recover the lost state upon a failure. The re-computation is distributed across multiple machines to accelerate failure recovery further. We also log intermediate data selectively, exploring the trade-off between recovery time and intermediate data storage overhead. Evaluations show that SWIFT significantly reduces the failure recovery time and achieves similar or better training throughput during failure-free execution compared to state-of-the-art methods without degrading final model accuracy. SWIFT can also achieve up to 1.16x speedup in total training time compared to state-of-the-art methods.

Index Terms—Distributed DNN Training; Failure Resilience

1 INTRODUCTION

LARGER and larger deep neural networks (DNNs) have recently emerged for improved model performance [1], [2], [3]. Large DNN model training jobs typically use many accelerators (e.g., GPUs) and have long-running times [4]. For example, training a GPT-3 model [3] on 1024 A100 GPUs is estimated to take more than one month [5].

Job failures are common in a GPU training cluster [4]. For example, machine crashes and network failures happen occasionally, or higher priority jobs take up resources [4]. In these cases, the distributed DNN training job is interrupted, resulting in the loss of the DNN model state (i.e., model parameters and optimizer states) and failure of the training job. Failures are more severe for large DNN model training jobs: increasing the number of machines will inevitably lead to an increased chance of failure; training large models takes days to months, making it more likely for failures to happen during the course. Recent works also echo this [6], [7].

Global checkpointing is the *de facto* method for fault tolerance in deep learning (DL) frameworks [8], [9]. The training job periodically checkpoints the entire model state. All workers restart from the latest checkpoint when the job fails. Depending on the checkpointing frequency, this often results in several hours of lost computation time [10]. CheckFreq [10] achieves more frequent checkpoints by splitting the operation into two phases: first, the model state is copied into the GPU memory, called a *snapshot*, or to the CPU memory if the GPU memory is insufficient; in the second phase, the snapshot is written to the disk asynchronously.

Elastic Horovod [11], a framework for elastic training, takes a similar approach, but without the second phase. The reason is that Elastic Horovod assumes distributed data-parallel training, where each worker maintains a replica of the model state; during failure recovery, one of the surviving workers broadcasts the snapshot to other workers, and all workers restart training from the snapshot. Taking a snapshot is necessary for Elastic Horovod to prevent a corrupted state: if a worker crashes during the parameters update, the other workers are in an awkward situation - some parameters are updated while others are not. We identify this problem as the *crash-consistency problem* (§2.3). However, as we shall see in §2.2, for large DNN models, both methods can slow down the training due to the overhead of snapshotting.

This paper studies a better failure resilience design for distributed DNN training that significantly reduces the recovery overhead without affecting training throughput and final model accuracy. One of our key observations is that many of the optimizers used for model state updates in DNN training are mathematically *invertible*. For example, stochastic gradient descent (SGD) only involves linear operators such as element-wise addition and scalar multiplication, and the inverse operators are straightforward. In case of a crash-consistency problem in distributed data-parallel training, we can restore the model states of the surviving workers to a consistent state by *undoing the update* of the updated parameters (§4). Therefore, we do not need to snapshot periodically as CheckFreq and Elastic Horovod do, reducing the overhead during failure-free execution to zero (except for periodic checkpoints). Since this approach exploits replicas of the model state in surviving workers for recovery, we name this recovery method *replication-based recovery*.

However, replicas are not always available, even with data parallelism. For instance, some prior works advocate

- Yuchen Zhong, Guangming Sheng, and Chuan Wu are with the University of Hong Kong, Hong Kong, China.
Email: {yczhong, gmsheng, cwu}@cs.hku.hk.
- Juncheng Liu and Jinhui Yuan are with OneFlow Inc., Beijing, China.
Email: {liujuncheng, yuanjinhui}@oneflow.org.

data parallelism only across multiple GPUs on the same machine to leverage high-speed intra-server interconnects such as NVLink [5], [12] to accelerate gradient synchronization. All replicas would be lost in the event of a machine failure.

We then investigate another fundamental approach for fault tolerance in distributed systems - logging, which has been widely explored in data processing systems [13], [14], [15], [16]. We introduce *logging-based recovery* (§5) for pipeline-parallel training. In pipeline parallelism, workers form a chain topology and pass intermediate activations or gradients to the successor or predecessor worker using point-to-point communication. Figure 1a illustrates One-Forward-One-Backward (1F1B) pipeline schedule [17] (§2.1). With logging, each worker locally records all outgoing data (i.e., intermediate activations and gradients) to the adjacent workers on the other machine. Upon a failure, the replacements of the failed workers retrieve the logging data and replay the computation to recover the lost state. Moreover, we spread the logging data to surviving workers to have them assist in recovery (§5.2). One notable distinction of pure global checkpointing is that surviving workers do not have to load the checkpoint and revert their training progress for recovery. This task is solely assigned to the relaunched workers on the replacement machine. Consequently, logging *limits the recovery scope from the complete computation graph across all workers to the computation graph on the failed workers*, thus reducing the recovery time compared to global checkpointing. An example is given in Figure 1b. To the best of our knowledge, we are the first to bring logging into distributed DL systems for failure resilience.


However, logging-based recovery brings unique challenges. Logging needs to be done constantly during DNN training, and the overhead in runtime and space can be prohibitive. Once a piece of logged data is missing, the original state cannot be recovered precisely. To reduce the runtime overhead of logging, we only log inter-machine communication data since failures often occur at machines rather than at individual workers on machines. Moreover, we perform logging asynchronously, storing the data in the background. We further utilize workers’ idle time (i.e., bubble time in pipeline parallelism) to do logging. This way, logging is *off the critical path* (§5.1). To control the space overhead at a manageable level, we devise an algorithm to select only a subset of machines to log intermediate data (§5.3). Selective logging trades the recovery time for space consumption.

We design and implement SWIFT, including replication-based and logging-based recovery for expedited failure recovery. Our key contributions are summarized below:


- ▷ We propose a novel mechanism called update-undo that resolves model state inconsistencies caused by the failure and enables failure recovery using replicas of the model state in data parallelism without creating additional copies.

- ▷ We propose to use the logging method to achieve expedited failure recovery in pipeline parallelism. We use asynchronous logging, logging during the bubble time, and selective logging to reduce the runtime and space overhead.

- ▷ We implement SWIFT in PyTorch and demonstrate its benefits in the distributed training of large DNN models. For replication-based recovery in training Wide-ResNet-50 [18], SWIFT reduces recovery time by 98.9%, 98.1%, and 98.1% compared to global checkpointing, CheckFreq and Elastic



(a) Pipeline parallelism with four devices (P0, ..., P3) with 1F1B schedule. The number in a block indicates the micro-batch being processed. The white blocks represent bubble time.



(b) An example showing that logging can expedite recovery after failure, where worker P1 in (a) has failed and is replaced by a new worker P1'. During normal execution, P1 needs to wait for intermediate activations or gradients from other workers. When recovering the lost state of P1, P1' can read the logged data (activations and gradients from other workers) and directly do the re-computation (without the bubble time as in normal execution), achieving much reduced recovery time than checkpoint recovery.

Fig. 1: Pipeline parallelism and logging example.

Horovod, respectively. For logging-based recovery in training BERT [1] and ViT [19], SWIFT reduces recovery time by 57.3% and 76.3% compared to global checkpointing, respectively. Using traces collected in our experiments, we show that SWIFT can achieve up to 1.16x speedup in total training time compared to state-of-the-art methods. We have open-sourced SWIFT at <https://github.com/jasperzhong/swift>.

2 BACKGROUND AND MOTIVATION

2.1 Distributed DNN Training

We focus on *synchronous* distributed DNN training, where many workers on multiple machines collectively work on the latest DNN model iteratively. Each training iteration contains a forward computation pass (to compute a loss) and a backward pass (to compute the gradients), and the gradients computed are used for the model update [20]. Synchronous training ensures better model accuracy than asynchronous training and is thus popular for large-scale DNN training [5], [20], [21], [22].

Data Parallelism is the most widely used paradigm for distributed DNN training [21], [22]. Input data is partitioned across workers. Each worker has a model replica and computes local gradients on a subset of data. Gradient synchronization is performed among workers in each iteration to ensure the consistency of model replicas.

Operator Parallelism is a solution to handle large DNNs by splitting an operator in a DNN model among multiple workers along non-batch axes [23]. Communication is needed to fetch the input data from other workers [2].

Pipeline Parallelism splits a mini-batch into smaller micro-batches and pipelines them to the DNN model stages hosted on different workers so that workers can process different micro-batches simultaneously [12], [17], [24], [25]. Point-to-point communication is performed between workers hosting

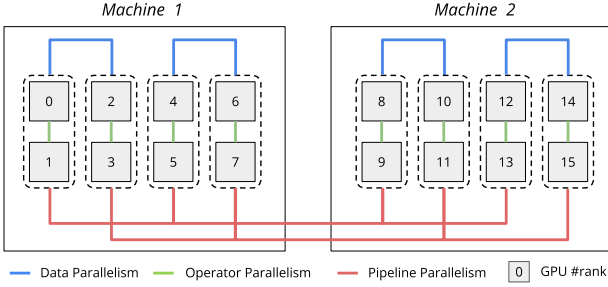


Fig. 2: A hand-optimized 3D parallelism plan in Megatron-LM, using 16 GPUs on two machines. The DNN model is split into four pipeline-parallel stages, each stage is partitioned onto two GPUs for operator parallelism and each stage has a replica. Replicas of a stage are on the same machine.

neighbor stages to transfer intermediate activations. Synchronous pipeline parallelism schedules like GPipe [24] and One-Forward-One-Backward (1F1B) [25] flush the pipeline in each iteration, i.e., worker waits for all in-flight micro-batches of the iteration to complete before moving on to the next iteration. Despite better model accuracy, pipeline flush causes worker idling (i.e., *bubbles*) in pipeline execution [5], [12], [24], [25]. For GPipe and 1F1B, the ratio of the bubble time is $(p - 1)/(m + p - 1)$, where p is the number of stages and m is the number of micro-batches [5]. For the example shown in Figure 1a, the ratio of the bubble time is $3/7$. This paper adopts 1F1B [25] because it has the same bubble time ratio but lower peak memory usage than GPipe [24]. Note that our approach is not limited to 1F1B.

Recent works combine the three parallelism paradigms, called *3D parallelism* [2], [5], [12], [26], [27]. Figure 2 shows a hand-optimized parallelism plan in Megatron-LM [2], [5], a state-of-the-art training system for transformer language models. Although data parallelism is used in this example, the replicas reside on the same machine. If one machine fails, we lose the model state on that failed machine. Furthermore, apart from the aforementioned works that primarily concentrate on homogeneous systems, there are also other studies that have explored parallel computing techniques to distribute deep learning workloads across heterogeneous systems [28], [29], [30].

2.2 Problems with Snapshotting Large Models

CheckFreq [10] and Elastic Horovod [11], state-of-the-art methods for fault tolerance, rely on the *snapshot* operation. After updating the model state of the iteration, a copy of the model state (called a snapshot) is captured in GPU memory or copied to CPU memory if the GPU cannot hold it. Snapshotting can overlap with the next iteration’s forward and backward pass. The next iteration of the update operation does not start until the snapshot operation is completed, leading to a checkpoint stall. However, DNN models have proliferated from millions to billions of parameters in recent years and become too large to fit into a single GPU [31]. It became increasingly difficult to fit a complete snapshot on a single GPU. In that case, the snapshot is copied to the CPU.

We experimentally find that snapshotting to CPU memory is costly for large models and reduces training throughput.




Fig. 3: Training throughput of Wide-ResNet-50 during failure-free execution. During the snapshots taken at iterations 30, 60, and 90, the iteration time experiences a notable increase with CheckFreq and Elastic Horovod.




Fig. 4: Crash-consistency problem in the layer-wise wait-free update. The number is the layer index. Arrows represent data dependency. The red dashed line indicates the failure.

We train an enlarged Wide-ResNet-50 [18] model, increasing the base channel size from 64 to 320. With Adam optimizer [32], the model state size is 9.8 GB. We utilize data parallelism on two machines, each equipped with 8 32 GB V100 GPUs. To implement snapshotting, we utilize CheckFreq’s code [33] and replace Elastic Horovod’s snapshot implementation. This change is necessary as Elastic Horovod does not support snapshotting to the CPU. The snapshot operation needs to copy the model state to the CPU. The optimal snapshot frequency is determined using CheckFreq’s algorithm, which suggests taking a snapshot once every 30 iterations. During training, the GPU memory consumption reaches 30.4 GB, which cannot accommodate a snapshot. Figure 3 shows that at the time of snapshots (iterations 30, 60, and 90), the iteration time is significantly longer with CheckFreq and Elastic Horovod. After the snapshots, CheckFreq’s iteration time is longer, showing that writing the snapshot to the disk also affects normal training. But global checkpointing causes a large overhead (iteration 100) since it is synchronous. Interestingly, the checkpoint stall is indeed negligible with CheckFreq, for only 0.2 milliseconds. This experiment shows that the snapshot operation can still incur non-trivial runtime overhead and slow down the training process.

2.3 Crash-consistency Problem

Most DL frameworks [8], [9], [26], [34] adopt wait-free model updates, as illustrated in Figure 4. Model state update of a DNN layer can be performed as soon as the gradient of that layer is ready. If a worker crashes during the update, the other workers are in an inconsistent state, where some layers are updated, and others are not. The problem exists not only

in data parallelism, but also in pipeline parallelism, where the DNN model is spread across multiple workers. Model state updates occur at different times due to the dependencies of the computation graph, and a worker moves on to the next training iteration once the part of the model state it hosts has been updated. Such inconsistencies can lead to a degradation of the accuracy of the final trained model [6], [35].

Elastic Horovod solves the problem with the snapshot operation. However, it can be costly for large models (§2.2). Another workaround is to wait for gradients of all layers to be ready before updating the model state at the workers (by adding a barrier before the model update). With this method, the other workers can still complete their model updates and remain consistent even if a worker fails during the update. However, this update method requires more waiting time. In §4, we propose a better solution to tackle the crash-consistency problem without snapshotting or incurring the waiting.

2.4 Logging-based Failure Recovery

The logging method logs a job’s application data at runtime and exactly replays the computation during failure recovery. Two main types of data are logged. Spark [13] and Ray [14], [36] record lineage, i.e., the computation graph. Other systems record (or just buffer) raw, intermediate data [15], [16]. For example, in *upstream backup* of stream processing systems [15], the upstream machines preserve the data in the output queue while the downstream machines are processing them.

In DNN training, the computation graph is usually fixed, and the execution time of a single operator in the computation graph is usually in the order of milliseconds. Recording the lineage of fine-grained operators adds significant overhead but does not benefit much since operators do not fail very often [8]. Therefore, we consider logging intermediate data. The parallelism paradigm determines the communication operators and, thus, the data to be logged. Data parallelism and operator parallelism use collective communications, such as all-reduce and all-gather [27]. Collective communications have complex data dependencies (e.g., many-to-many), thus complicating logging. Pipeline parallelism performs point-to-point communication, which simplifies logging. Also, the communication data volume in pipeline parallelism is much smaller than in data parallelism and operator parallelism for transformer models, and the current mainstream large models are usually transformer-based models [3], [5]. Therefore, we advocate logging for pipeline parallelism.

3 SWIFT OVERVIEW

Given a distributed DNN training job and its specific parallelism configuration, SWIFT introduces failure resilient mechanisms into the training job, targeting reduced failure recovery time without affecting the training throughput and final model accuracy. The distributed DNN training job spans a cluster of physical machines, with each machine hosting one or multiple workers. We focus on a fail-stop model [37] throughout the paper, which is more common in real-world clusters [4]. A machine in the cluster may crash during

TABLE 1: Operators used in five representative optimizers. EW = element-wise; Inv. = invertible.

Operator		SGD	Adam [32]	AdamW [38]	LAMB [22]	AMSGrad [39]
Inv.	EW add	✓	✓	✓	✓	✓
	scalar mul	✓	✓	✓	✓	✓
	EW mul		✓	✓	✓	✓
	EW sqrt		✓	✓	✓	✓
	EW div		✓	✓	✓	✓
Not Inv.	EW-max					✓
	sum				✓	

training, losing the volatile model states of workers it hosts, i.e., parameters and optimizer states, which are mainly stored on the GPUs.

SWIFT decides on a fault tolerance strategy before the training job starts. It always exploits redundancies if available (i.e., if the model state has at least one replica on another machine) because replication-based recovery achieves both low runtime and recovery overhead. When a replica is unavailable, pipeline parallelism is used, and logging is *worth doing* (§5.4), then use logging-based recovery. If none of the above conditions are met, use global checkpointing only. In any case, global checkpointing is performed periodically to ensure that the system remains on track in case of a catastrophic failure (e.g., loss of all replicas or logging data). Replication-based recovery and logging-based recovery can be combined to use, as hybrid parallelism is common for large DNN training, e.g., parts of the model use data parallelism while other parts use pipeline parallelism [27].

A machine failure can be detected by catching communication errors by workers that communicate with the failed machine. After the failure is detected, a replacement machine will be added to the training job. The surviving workers stop training and start the failure recovery procedure. Surviving workers first resolve the model state inconsistency issue (§2.3) with the *update-undo* approach (§4). Recovery is then performed for the replacement of the failed workers. For replication-based recovery, one of the surviving workers which holds the replica broadcasts the model state to the replacement workers (which uses data parallelism with the surviving worker). For logging-based recovery, replacement workers load the most recent checkpoint and then recompute the lost iterations based on the logged data until recovering up to the pre-failure iteration. We also discuss multiple failures and cascading failures in Appendix B.

For the example in Figure 2, logging-based recovery can be used since replicas are unavailable and pipeline parallelism is used across the two machines. We record data of inter-machine communication during training (§5.1): GPU 3 & 7 log the intermediate activations in the forward pass, while GPU 11 & 15 log the gradients in the backward pass.

4 UPDATE-UNDO

We propose undoing the update to address the crash-consistency problem (§2.3) to facilitate failure recovery using replicas of the model state in data parallelism without generating additional copies. Our idea is simple: if a failure occurs during the model update when some parameters at

Algorithm 1 SGD with Momentum

- 1: **Input:** learning rate sequence $\{\eta_t\}_{t=1}^T$; weight decay $\lambda > 0$; momentum parameter $0 \leq \mu \leq 1$; dampening for momentum $0 \leq \tau \leq 1$.
- 2: **Initialize:** $x_1 \in \mathbb{R}^d$; $m_0 = 0 \in \mathbb{R}^d$.
- 3: **for** $i = 1$ to T **do**
- 4: $g_t = \nabla f(x_t)$
- 5: $m_t = \mu m_{t-1} + (1 - \tau)(g_t + \lambda x_t)$
- 6: $x_{t+1} = x_t - \eta_t m_t$
- 7: **end for**

Algorithm 2 Undo SGD with Momentum

- 1: **Input:** learning rate η_t ; weight decay $\lambda > 0$; momentum parameter $0 \leq \mu \leq 1$; dampening for momentum $0 \leq \tau \leq 1$; $x_{t+1} \in \mathbb{R}^d$; $g_t \in \mathbb{R}^d$; $m_t \in \mathbb{R}^d$.
- 2: $x_t = x_{t+1} + \eta_t m_t$
- 3: $m_{t-1} = (m_t - (1 - \tau)(g_t + \lambda x_t)) / \mu$

the workers have been updated and some have not, the surviving workers will *undo* the update for the updated parameters. In addition to model parameters, optimizer states, such as the momentum, also need to be restored. Finally, all workers return to a consistent version of the model state.

We observe that many update operators of optimizers are mathematically *invertible*, i.e., for an operator f , there exists an inverse operator f^{-1} that undoes the operation of f . For example, linear operators like element-wise addition and scalar multiplication are all invertible [40]. Table 1 summarizes operators used in five representative optimizers. Algorithm 2 demonstrates how the update of SGD with momentum (Algorithm 1) can be undone (i.e., from x_{t+1} to x_t and from m_t to m_{t-1}). More examples can be found in Appendix A. If an optimizer only has linear operators, then undoing is straightforward. However, some optimizers involve non-linear operators, e.g., LAMB optimizer [22] scales the gradients with the L2 norm of the parameters. For the LAMB optimizer, we can additionally save the L2 norm (a scalar), and recover the previous model state accordingly. For AMSGrad [39], update undo is not applicable. Although the undo algorithms are mathematically correct, the recovered state may slightly differ from the original state due to floating-point errors [41]. Our experiments show that this minor error does not affect trained model accuracy (§7.2). Figure 5 gives an example of how undoing updates helps replication-based recovery.

Undoing the update does not require extra GPU memory. It only needs to cache the latest gradients g_t , a common practice in mainstream DL frameworks [8], [9], [34]. When gradients for the next iteration, g_{t+1} , are computed at the next backward pass, with synchronous training, parameters of all workers have already been updated by g_t , so that g_{t+1} can safely overwrite g_t in memory. It is only necessary to maintain one version of the gradients.

5 LOGGING-BASED RECOVERY

We propose employing the logging method to expedite failure recovery in pipeline parallelism. Our approach incorporates asynchronous logging, logging during the bubble

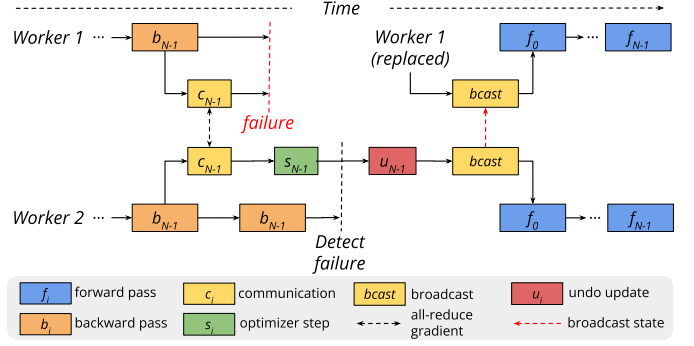


Fig. 5: Replication-based recovery with data parallelism and undo operation. Two workers train a DNN model using synchronous data parallelism. In training iteration t , worker 1 crashes and loses all the volatile states during the backward pass. At this time, worker 2 has already updated the parameters of layer $N - 1$, but not the parameters of the other layers. Worker 2 then undoes layer $N - 1$'s update to ensure consistency of its state. When the replacement of worker 1 joins the system, worker 2 sends its state to worker 1, and then both continue the training from iteration t .

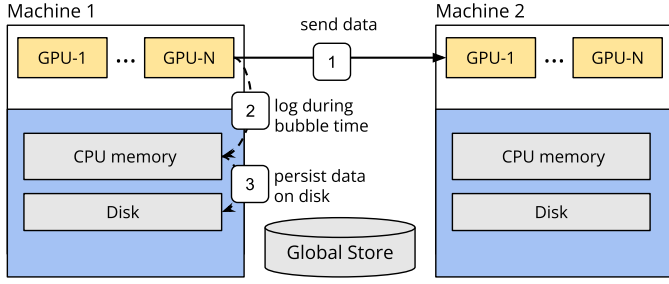
time, and selective logging to reduce both runtime and space overhead effectively.

5.1 Basic Mechanism


What data to log. For multi-GPU machines, GPUs rarely fail individually, while a machine crash is more common [4]. We hence do not record GPU-to-GPU communication within a machine but only *inter-machine communication*, reducing the logging overhead substantially. The data to log include the intermediate activations in the forward pass and the gradients in the backward pass. In addition to saving the raw tensor, we need to record some metadata, including the sender and the receiver, and the timestamp (which contains identifiers of the current training iteration and the current micro-batch being trained). The timestamp is used to determine the order of the data to replay during recovery.

How to log data. Our logging method is similar to upstream backup (§2.4), i.e., the sender rather than the receiver logs the message. This way, the intermediate data needed for failure recovery are not lost but remain on the sender machines upon a failure. A sender does not need to store the message before sending it but can log it asynchronously in the background. A queue is set up for each worker. The worker keeps pushing outgoing tensors into the queue during training, while another background thread keeps reading tensors from the queue and doing the logging. Each worker flushes the queue of uncompleted logging tasks when detecting a failure in the training job. This *asynchronous logging* significantly alleviates its impact on the training throughput.


Similar to asynchronous checkpointing (as in CheckFreq), which still incurs significant overhead (§2.2), we need to reduce the overhead of our asynchronous logging further. In synchronous pipeline-parallel training, there are many bubbles during which the worker is idle, which is an ideal time for logging (§2.1). With *logging during the bubble time*, the outgoing tensor is not logged immediately after production,



(a) Logging during runtime. It eliminates the runtime overhead of logging by implementing asynchronous logging and logging during bubble time.



(b) Logging during recovery. Surviving workers flush uncommitted data and upload logging files to global storage. The replacement machine downloads relevant logging files, loads the latest checkpoint, and replays tensors in chronological order.



(c) Logging with parallel recovery. Steps 1 to 3 are the same as in Figure 6b and are omitted. Surviving workers checkpoint their states, and load model states from replacement workers. All workers download logging files and select micro-batches for re-computation. After the recovery, the surviving workers load their checkpoints to restore their states.

Fig. 6: Logging mechanism.

but waits until a bubble occurs. The waiting time is not longer than one training iteration because a bubble always exists within an iteration. Once the logging data is transferred to the CPU, the data on the GPU can be safely removed. Therefore, logging during the bubble time will not cause substantial data accumulation on the GPU. This way, logging is *off the critical path*.

Figure 6a illustrates how logging is done during failure-free execution. After a worker (GPU) sends a tensor to a downstream worker (step 1), the data remains on the GPU for a short time. During the next bubble time, the tensor is copied asynchronously to the CPU memory (step 2). A background thread writes the logging data in CPU memory to the disk and saves them into a file (step 3).

Recovery. Figure 6b shows how failure recovery can be made using the logging data. Suppose machine 2 crashes. The

surviving upstream workers (machine 1 in this example) copy unlogged data to CPU memory (step 1) and then write them to the disk (step 2). The upstream workers then upload their logging files to global storage (step 3), e.g., HDFS [42]. Machine 2's replacement downloads the logging files that it needs from the global store to its local disk (step 4), loads the latest checkpoint, and replays previously received tensors from the logging file in the exact order of their timestamps (step 5). The step 3, 4, and 5 can be executed in a pipeline by chunking the logging file into multiple smaller files. If necessary, the surviving workers will undo the update (§4) (not shown in the figure).

The most significant difference with pure global checkpointing is that surviving workers do not need to load the checkpoint and roll back their training progress for recovery. Only relaunched workers on the replacement machine do. The recovery scope is limited to the local computation graph on the failed machine rather than the whole computation graph, thus expediting the recovery.

Garbage Collection. All earlier logging files are obsoleted after a global checkpoint, and garbage is collected because the system can then directly load the latest checkpoint. Even though the logging size increases as the number of iterations increases, the size is upper bound due to periodic global checkpointing. Therefore, the frequency of global checkpointing determines the upper bound of the logging size. We will discuss the storage overhead in §5.3.

Consistency. We replay the computations in the same order as the pre-failure execution using timestamps, using the same inputs as the pre-failure computation. Note that logging requires the computation to be *deterministic* (i.e., the same input leads to the same output). Otherwise, we would get different outputs when re-computing with the logged data. We provide the details of achieving determinism in §6.

5.2 Parallel Recovery

We utilize the surviving workers to assist in the recovery of the failed workers. Since the intermediate results of all micro-batches have been logged, we can perform data-parallel training based on the logged data to expedite the re-computation of the lost states. Specifically, all workers, including replacement workers and surviving workers, retrieve the logging files and have a copy of the computation graph on the failed machine. Each worker reads logging data of different micro-batches from the logging files and uses them as input to re-compute gradients, synchronizes gradients with other workers computing other micro-batches and then performs the model update. This way, the micro-batches are re-computed in parallel by multiple workers for the failed machine, accelerating recovery while ensuring logical equivalence to executing these micro-batches sequentially by each replacement worker. If a batch is divided into m micro-batches and we use d workers for parallel recovery, each worker is assigned with m/d micro-batches for re-computation.

An example is given in Figure 7, where two machines run a 4-stage training pipeline. Machine 2 (hosting stages 2 and 3) fails, and Machine 1 (hosting stages 0 and 1 in normal training) assists in the recovery computation of the replacement machine. Worker P0 on Machine 1 and worker




Fig. 7: Parallel recovery. The number of micro-batches is 4 ($m = 4$) and we use 2 workers ($d = 2$) in each data-parallel recovery group. Each machine has two workers. Suppose machine 2 crashes and is replaced. Worker P0 and P2 re-compute the stage-2 model and worker P1 and P3 re-compute the stage-3 model. We decompose the pipeline in Figure 1a into data-parallel 2-stage sub-pipelines.

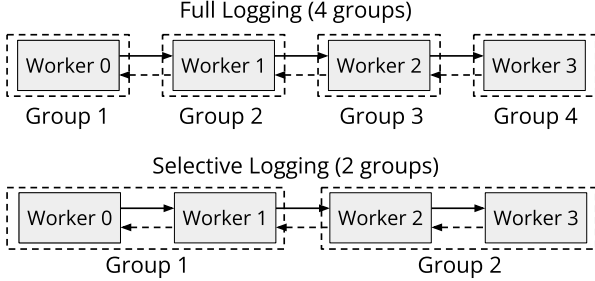


Fig. 8: Selective logging.

P2 on the replacement machine re-compute the stage-2 model in a data-parallel manner, each using two micro-batches (0, 2 and 1, 3, respectively); worker P1 and worker P3 re-compute the stage-3 model, using micro-batches 0, 2 and 1, 3, respectively. Note that extra time is needed for gradient synchronization with parallel recovery.

The parallel recovery procedures are given in Figure 6c. Similar to Figure 6b, at a surviving worker, uncommitted logging data are first flushed and uploaded to the global store (steps 1 to 3, omitted in the figure). Then, the surviving workers checkpoint their states (step 4). The replacement workers load their model states from their latest checkpoints and broadcast their states to the surviving workers (step 5). Meanwhile, all workers download logging files from the global store (step 6) and select the logging data of corresponding micro-batches for re-computation (step 7). After the recovery, the surviving workers load their checkpoints to restore their original model parameters and optimizer states (not shown in the figure).

5.3 Selective Logging

Logging all cross-machine messages may consume large storage space. We next investigate a trade-off between the storage space and the recovery time with selective logging. Our idea is to group machines and log inter-group communication but not intra-group communication, as shown in the lower section of Figure 8. We can consider the original approach as a particular case where each machine forms a group, as illustrated in the upper section of Figure 8. In this way, if one machine in a group fails, training on the entire group of machines needs to be rolled back from the latest checkpoint, as we do not record intra-group communication. As a result, and the recovery time will be longer. Thus, selective logging trades recovery time for space overhead. A simple grouping strategy is to have a balanced number of machines in each group. However, due to the often unbalanced model partition in pipeline parallelism [12], this grouping strategy is usually

suboptimal. Given a storage capacity constraint for logging data, how do we group machines to minimize the failure recovery time? Suppose we have N machines and create N groups initially. We profile the averaged per-iteration computation time $R(G_i)$ for each group G_i . For each pair of adjacent groups G_i and G_{i+1} (i.e., hosting adjacent workers in the pipeline), we obtain the transmission size per iteration $M(G_i, G_{i+1})$ between them. Then, with storage capacity limit M_{\max} , network bandwidth B (assuming homogeneous bandwidth) and checkpointing interval T (iterations), we aim at finding a group configuration $\mathcal{G} = \{G_1, \dots, G_k\}$ that minimizes the overall recovery time R :

$$\min_{\mathcal{G}} R(\mathcal{G}) \quad \text{s.t. } M(\mathcal{G}) \leq M_{\max},$$

where $M(\mathcal{G})$ denotes the overall storage space needed by the logging data. As discussed in §5.1, it is determined by the global checkpointing frequency:

$$M(\mathcal{G}) = T \cdot \sum_{G_i, G_{i+1} \in \mathcal{G}} M(G_i, G_{i+1}).$$

Suppose we merge two adjacent groups G_i and G_{i+1} , and have the following recovery time for the merged group:

$$R(G_i, G_{i+1}) = R(G_i) + R(G_{i+1}) + M(G_i, G_{i+1})/B,$$

where $M(G_i, G_{i+1})/B$ is the point-to-point communication time between the two adjacent groups. We ignore the bubble time for simplicity, and derive the change in overall recovery time R and overall space overhead M :

$$\begin{aligned} \Delta R &= R(G_i, G_{i+1}) \cdot \frac{|G_i| + |G_{i+1}|}{N} - R(G_i) \cdot \frac{|G_i|}{N} \\ &\quad - R(G_{i+1}) \cdot \frac{|G_{i+1}|}{N}, \\ \Delta M &= M(G_i, G_{i+1}) \cdot T, \end{aligned}$$

where $|G_i|$ is the number of machines in G_i . ΔR is calculated assuming that each machine has an equal failure probability. Note that ΔR is always positive. We minimize increased recovery time ΔR per unit storage space reduction when merging G_i and G_{i+1} , i.e., minimize $\Delta R/\Delta M$.

To identify the grouping of machines, we iteratively merge two adjacent groups with the smallest $\Delta R/\Delta M$, until the overall space consumption is less than M_{\max} . Note that it runs for at most $N - 1$ iterations, at which point all machines form a single group and there will be no logging. So the time complexity is at most $O(N^2)$. If parallel recovery is used, the recovery of a group G_i is parallelized by at most $\lfloor N/|G_i| \rfloor$ data-parallel groups. For simplicity, we assume it can achieve linear scalability with data parallelism. Thus, we divide the $R(G_i)$ with $\lfloor N/|G_i| \rfloor$ in calculation.

5.4 Use Case

Not all cases are suitable for logging. For example, it would be better to checkpoint a model when the logging size far exceeds the model size. Typically, the intermediate activations for CNN-based models would be massive and unsuitable for logging (even unsuitable for pipeline parallelism) [43]. We can calculate the per-iteration logging size. For transformer-based models, the intermediate activation/gradient size would be $\text{micro_batch_size} \times \text{hidden_size} \times \text{sequence_length}$ in a micro-batch [31]. Further, we can calculate the bubble time ratio according to the pipeline schedule (§2.1). Given the iteration time and PCIe bandwidth, we can determine whether the logging data can be transferred from GPU to

CPU within the bubble time. If not, then logging is not worth doing.

6 IMPLEMENTATION

We implement SWIFT in PyTorch 1.9.0 [9] with NCCL 2.7.6 [44], using 2.6k LoC in Python. We also add about 400 lines of C++ code for PyTorch and NCCL.

Failure detection. We launch a background thread on each worker that uses NCCL’s `ncclCommGetAsyncError()` function to keep polling whether a communication failure has occurred. If a failure occurs, the worker first sets a failure flag to true in a global key-value store and then aborts its own NCCL communicators. The global key-value store is co-located with the master machine (rank 0). Other workers’ background threads also poll this flag from the global key-value store, and if a worker finds that the flag is set to true, it will abort its own NCCL communicators.

Update-undo. In data parallelism, we insert a CUDA event after the all-reduce operation for each tensor and query whether it has been completed before updating the gradient’s corresponding parameter. If it does not complete, it waits until the all-reduce operation completes. After it completes, the CUDA kernels for the corresponding parameter are launched to update the parameter and optimizer states, and the parameter is marked as updated. Note that even if there is a failure at this point, we need to let these kernels finish executing. Upon a failure, surviving workers undo the update of parameters that are marked updated.

In pipeline parallelism, the model parameters on different stages are updated at different points in time due to computational dependency. Therefore, surviving workers need to exchange their current iteration number to determine the consensus pre-failure iteration after a failure occurs. Workers with a greater iteration number than the consensus pre-failure iteration need to undo the update.

Logging. We use a dedicated CUDA stream to copy logging data from the GPU to the CPU for asynchronous logging. We insert a CUDA event after the copy operation to check if the copy operation is completed. After the main thread launches asynchronous copying operations at bubble time, it sends the CUDA event with the corresponding tensor and metadata to a queue. A background thread keeps reading items from the queue and checks if the asynchronous copy is complete by checking the CUDA event status. If completed, the thread saves the data to a file. For the global store described in §5.1, we support HDFS [42] and Amazon S3.

Determinism in Logging. Nondeterminism in DNN training may come from the random number seeds and algorithms themselves. For example, in DNN training, some convolution algorithms in cuDNN are nondeterministic¹. We set `torch.backends.cudnn.deterministic=True` to resolve this issue. In addition, for convolutional operations, PyTorch also benchmarks multiple algorithms in the first run, selects the fastest one, and caches this choice so that the same algorithm can be directly selected later. However, there are still some slight differences in the computational results of different deterministic algorithms for the same input. In order

1. cuDNN reproducibility: <https://docs.nvidia.com/deeplearning/cudnn/developer-guide/index.html#reproducibility>

TABLE 2: Benchmark Models. DP = data parallelism. PP = pipeline parallelism.

Model	Dataset	Batch size	#params (billion)	Parallelism
Wide-ResNet-50	ImageNet [45]	256	1.23	DP
ViT-128/32	ImageNet	4096	1.64	PP
BERT-128	Wikipedia [1]	512	1.11	PP

to ensure that the worker selects the same convolutional algorithm after failure recovery as before failure, we save the previous benchmark results for failure recovery.

Usage. We provide an easy-to-use interface for users. A user only needs to provide a user-defined function (UDF) to train for one iteration and specify fault tolerance and training configurations. Then fault tolerance is in place during the user’s model training, and recovery upon a failure can be automatically run without requiring user involvement.

7 EVALUATION

Testbed. We experiment on 16 DGX-2 machines, each equipped with eight 32 GB Tesla V100 GPUs (NVLink interconnect), 160 CPUs, 1.5 TB memory, and 3.6 TB NVMe SSD disks. The machines are connected via 40Gbps Ethernet. We build an HDFS cluster on these machines as the global storage.

Benchmark Models. We evaluate SWIFT on training large image classification and language models with billions of parameters, as given in Table 2. We scale up the original models in their respective papers: for Wide-ResNet-50 [18], we increase the base channel size from 64 to 320; for BERT-Large [1] and ViT-Large/32 [19], we increase the number of transformer layers from 24 to 128, keep the hidden size unchanged and refer to the enlarged models as BERT-128 and ViT-128/32, respectively. We use data parallelism to train Wide-ResNet-50 on two machines and four GPUs on each. To train ViT-128/32 or BERT-128 (with a maximum sequence length of 128), we use a 128-stage pipeline on all 16 machines, with each transformer layer occupying one GPU. We use SGD with momentum for Wide-ResNet-50 and ViT-128/32, and Adam for BERT-128 [1], [19], [19]. We select the micro-batch number to maximize the performance, using 16 and 4 for ViT-128/32 and BERT-128, respectively. SWIFT applies replication-based recovery to Wide-ResNet-50, and logging-based recovery to ViT-128/32 and BERT-128 (by default 16 machine groups and 8 machine groups in selective logging). We run each experiment for 200 iterations, perform a global checkpoint at the beginning of iteration 100, and kill a machine (rank 1) at the beginning of iteration 150.

Baselines. We compare SWIFT with global checkpointing (default in PyTorch), CheckFreq [10] and Elastic Horovod [11]. We use CheckFreq’s open-sourced code [33] and replace Elastic Horovod’s snapshot implementation with CheckFreq’s since it does not implement snapshotting to the CPU. We calculate the optimal snapshot frequency (once per 30 iterations) based on the algorithm suggested by CheckFreq and using the same permissible checkpoint overhead (3.5%) as in CheckFreq’s experiments. For logging-based recovery, we

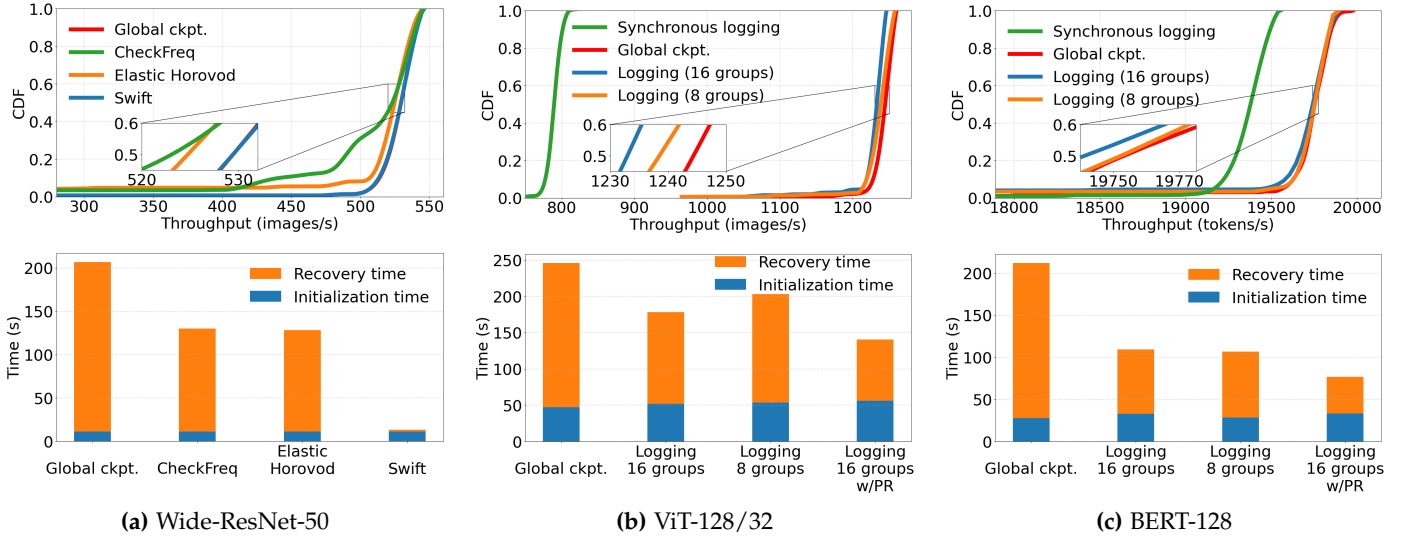


Fig. 9: Failure-free training throughput (top) and recovery time (bottom). PR = parallel recovery.

only compare with global checkpointing, as its checkpointing overhead is already very low (checkpointing is pipelined in pipeline-parallel training), and the performance of CheckFreq would be similar. Elastic Horovod is not applicable since it only supports data parallelism. We also introduce a synchronous logging method (calling `torch.save()` before sending a tensor) as a baseline to evaluate the effect of our asynchronous logging and logging during bubble time (§5.1).

Metrics. We evaluate the training throughput and iteration time during failure-free execution, throughput during recovery, and recovery time. Training throughput is calculated by the number of images (or tokens) processed by all workers per training iteration. Initialization time counts from when workers detect the failure to when the replacements of failed workers join the training job. Recovery time is the duration from when the replacements of workers join the training job to the time they recover to the pre-failure iteration.

7.1 Macro-benchmarks

Replication-based Recovery. Figure 9a shows that SWIFT’s replication-based recovery incurs less runtime overhead than state-of-the-art methods during failure-free training. Training throughput of CheckFreq and Elastic Horovod degrades compared to the normal training (without any checkpoint or snapshot). Figure 9a also presents the recovery time upon a machine failure at iteration 150. Global checkpointing takes a long time to recover, as all workers must load the checkpoint and re-compute the lost iterations (50 iterations in this experiment). CheckFreq and Elastic Horovod do frequent snapshotting but still need to re-compute 30 iterations (the last snapshot was captured at iteration 120). With SWIFT’s replication-based recovery, surviving workers resolve inconsistencies by undoing updates and then broadcast replicas to the replacement workers. It reduces the recovery time by 98.9%, 98.1%, and 98.1% as compared to global checkpointing, CheckFreq, and Elastic Horovod, respectively.

Logging-based Recovery. Figure 9b and Figure 9c show that SWIFT’s logging is slightly slower than global checkpointing




Fig. 10: Training throughput of ViT-128/32 during failure recovery. Blue dashed line indicates completion of recovery with global checkpointing.

during failure-free training of ViT-128/32 and achieves similar throughput for BERT-128. The slight delay for ViT-128/32 is because we use a large batch (4096) when training and the logging data size is relatively large. Synchronous logging significantly degrades training throughput, especially when training ViT-128/32, due to logging more data than BERT-128. SWIFT’s asynchronous logging and logging during bubble time (§5.1) take logging off the critical path, leading to a similar throughput compared to global checkpointing.

Figure 9b and Figure 9c show that the recovery time with logging is substantially shorter than global checkpointing. With 16 machine groups, the recovery time is reduced by 36.0% and 58.5% for ViT-128/32 and BERT-128, respectively. This is because only the 8-stage sub-pipeline on the failed machine needs to be recovered, compared to re-running the whole 128-stage pipeline when using global checkpointing. Note that logging needs slightly more initialization time because it requires additional initialization operations, such as creating a CUDA stream and logging threads.

Machine Group Size. Figures 9b and 9c also show the impact of different machine group sizes on training throughput and recovery time for logging-based recovery. In Figure 9b, with 8 machine groups, the throughput is similar to global checkpointing, due to less logging data than with 16 machine




Fig. 11: Trade-off between recovery time and storage space limit. Marker: (recovery time in seconds, storage limit in gigabytes).




Fig. 12: End-to-end training. The red dash line indicates a failure.

TABLE 3: Space overhead caused by logging per iteration.

Model	#Machine group	Total logging size (GB)	Average consumed bandwidth (GB/s)
ViT-128/32	16	24.66	0.23
	8	11.51	0.11
BERT-128	16	8.05	0.075
	8	3.76	0.035

groups (§5.3). In Figure 10 and Figure 9b, we observe that logging with 8 machine groups requires a longer recovery time due to recovering a 16-stage sub-pipeline on two machines instead of the 8-stage sub-pipeline in the case of 16 machine groups. Table 3 shows the total logging size per iteration and average bandwidth taken by logging in bubble time with different models and different numbers of machine groups. This shows the trade-off between recovery time and space overhead with selective logging (§5.3).

Parallel Recovery. For logging with parallel recovery (§5.2) cases in Figures 10, 9b, and 9c, we use 16 workers (GPUs) to concurrently do the recovery computation for one failed worker. We see that parallel recovery significantly improves training throughput for ViT-128/32 from 12.5x to 15x as compared to global checkpointing (similar results for BERT-128) due to reduced recovery time (by 57.3% and 76.3% for ViT-128/32 and BERT-128, respectively). The throughput fluctuation with parallel recovery is because parallel recovery is so fast that file transfer becomes a bottleneck, i.e., the new logging files are not yet downloaded from HDFS while the replay is already done with the earlier files.

Space-time Trade-off. Figure 11 further evaluates the trade-off between recovery time and space overhead with selective logging. Given a maximal storage capacity, we use the algorithm in §5.3 to decide how to group machines. For both DNN models, the recovery time becomes longer when we lower the space threshold. We can identify good trade-offs in practical usage using the plotted curves. The grouping configurations can be found in Appendix C.

7.2 End-to-end Training

We next run end-to-end training to verify that SWIFT does not affect the trained model accuracy. In Figure 12a, we train ResNet-50 on the ImageNet dataset from scratch using data

TABLE 4: Training workload in the simulation study.

Model	total # of iterations	checkpoint interval	End-to-end training time w/o failure
Wide-ResNet-50	450,360	5,004	479.4hr
ViT-128/32	93,600	312	85.6hr
BERT-128	500,000	5,000	461.1hr

TABLE 5: Simulated end-to-end training time with failures.

Model	#failure	Global ckpt.	SWIFT	Speedup
Wide-ResNet-50	28	557.4hr	480.7hr	1.16x
ViT-128/32	5	86.4hr	86.0hr	1.01x
BERT-128	27	524.2hr	476.1hr	1.10x

parallelism, adopting SGD with momentum on 16 GPUs across four machines. We kill one machine at the beginning of epoch 20, intentionally make an additional update at epoch 20, and then undo this update. We observe that update-undo does not affect the final validation accuracy.

In Figure 12b and Figure 12c, we train ViT-Base/32 [19] and BERT-Base from scratch on the ImageNet and Wikipedia datasets, respectively. For both models, we employ pipeline parallelism using 16 GPUs across four machines and use AdamW optimizer [38]. At the start of epoch 20, we terminate the machine hosting the second stage of the pipeline (machine 1 with workers from rank 4 to rank 7). Similarly, for BERT-Base, we terminate the machine at the beginning of iteration 5000. We do not group the machines for logging nor enable parallel recovery. We observe that our logging-based failure recovery has no loss of accuracy compared to the failure-free counterpart.

7.3 Simulation Study

We further investigate the effects of SWIFT on the end-to-end training time through simulations. Simulation settings are given in Table 4 (others are the same as experimental settings in §7.1). We calculate the expected end-to-end training time without failures based on the iteration time measured in the experiments and the total number of training iterations. For Wide-ResNet-50 and ViT-128/32, we assume storing a checkpoint at the end of each epoch, following common ML practice [18], [19], [46]. For BERT-128, we assume performing

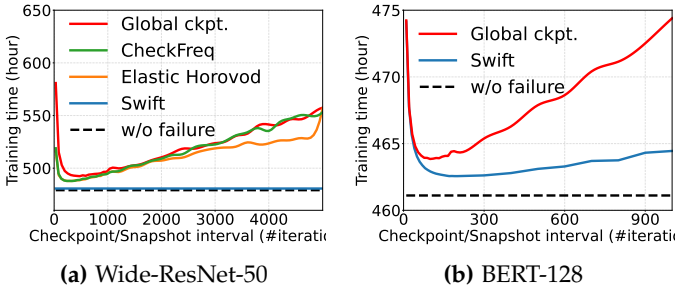


Fig. 13: Impact of checkpoint frequency.

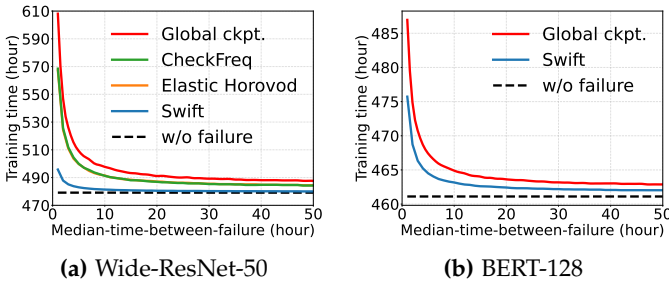


Fig. 14: Impact of failure frequency.

checkpointing once every 5000 iterations, which is 1% of its total number of training iterations. We then inject failures uniformly randomly during training, assuming a 17-hour median-time-between-failure (following [6]). We repeat each simulation ten times and present the average results.

End-to-end Training Time. As shown in Table 5, SWIFT can reduce the end-to-end training time significantly for long-running jobs, as compared to global checkpointing. Specifically, SWIFT can speed up end-to-end training for training Wide-ResNet-50 on ImageNet and pretraining BERT-128 on the Wikipedia dataset by 1.16x and 1.10x, respectively. This translates into saving 77 hours and 48 hours of training time. Short-running jobs like training ViT-128/32 on ImageNet encounter fewer failures and thus benefit less from fast failure recovery. We also compare the end-to-end training time of Wide-ResNet-50 with Elastic Horovod and CheckFreq. We consider the overhead of snapshots in our simulations using data collected in §7.1 and use the same snapshot frequency as in §7.1. End-to-end training with CheckFreq takes 518.9 hours, and with Elastic Horovod, it takes 515.9 hours. SWIFT is 1.08 and 1.07 times faster than CheckFreq and Elastic Horovod, respectively.

Effects of Checkpoint Frequency. We vary the checkpoint/snapshot frequency to investigate its impact on end-to-end training time. We keep the checkpoint frequency unchanged for replication-based recovery in SWIFT since it does not require frequent checkpointing. As shown in Figure 13, SWIFT achieves a shorter training time than other methods in all cases. An optimal checkpoint frequency can be obtained for each method from the curves, which leads to the shortest training time. Comparing the optimal cases of each method, for Wide-ResNet-50, SWIFT saves 11.8 hours, 7.1 hours, and 7.2 hours compared to global checkpointing, CheckFreq, and Elastic Horovod, respectively; for BERT-128, SWIFT saves 1.3 hours compared to global checkpointing

- limited improvement due to the minimal checkpointing overhead of BERT-128 (0.93 seconds).

Effects of Failure Frequency. We further adjust the median-time-between-failure to investigate its effect on end-to-end training time while fixing the checkpoint/snapshot frequency to the optimal frequencies given by Figure 13. Figure 14 shows that SWIFT achieves better speedup when failures are more frequent and also the shortest training time among all methods when failures are infrequent.

8 RELATED WORK

Elastic Training. Most DL jobs use static job configuration (e.g., the number of workers). In elastic training, workers can join and leave. The job can scale out to utilize transient idle resources (e.g., spot instances in cloud computing), or scale in to reserve resources for high-priority jobs. Unfortunately, most elastic training works [47], [48], [49] still rely on checkpoint-restart method to avoid the crash-consistency problem (§2.3). SWIFT can resolve the inconsistency using update-undo (§4) and thus benefit elastic training (e.g., broadcast the worker’s state when new workers come in).

Checkpointing in DL Systems. Check-N-Run [7] proposes incremental checkpointing tailored for training DL recommendation models, exploiting the fact that only a fraction of the recommendation model is updated in each iteration. It is complementary to our work because SWIFT is not limited to recommendation models. The MLP layers in recommendation models are usually trained using data parallelism [6], which can benefit from the replication-based recovery. Orpheus [50] also proposes incremental checkpointing but stores sufficient vectors of gradients, which are much smaller than the gradients themselves. During recovery, the gradients are reconstructed by the stored sufficient vectors and applied to a checkpoint to recover the model state. Our logging method can be seen as an extension of their approach. We consider the sufficient vectors (e.g., intermediate activation/gradient in pipeline parallelism) of the computation graph on a machine rather than a single operator. We log data asynchronously by upstream machines while they require costly synchronous logging for consistency. Recent works also propose partial recovery that loads the checkpoint of the failed machine only and continues with training [6], [35]. Partial recovery avoids global rollback but incurs accuracy loss due to inconsistent model states among workers [6]. In contrast, SWIFT does not degrade final model accuracy while reducing recovery time.

Large-scale DNN Training. In addition to parallelism, other complementary techniques for large-scale DNN training include memory optimization [31] and mixed-precision training [51]. SWIFT can be combined with many of them. For example, we can combine our replication-based recovery with Fully Sharded Data Parallel (FSDP), a popular memory optimization technique that shards the model state across data-parallel workers [31]. We can maintain two copies of each piece of the sharded model state for failure resilience. Moreover, mixed-precision training can reduce the logging size due to using a lower precision for intermediate data [51].

9 CONCLUSION

This paper presents SWIFT, a novel design that expedites failure recovery in distributed DNN training. SWIFT exploits redundancies in data-parallel training for failure recovery and resolves the crash-consistency problem with update-undo. SWIFT advocates logging for pipeline-parallel training, which records inter-machine intermediate data at runtime and limits the computation graph to be re-executed to those on the failed workers. We also design parallel recovery to expedite recovery further and explore the trade-off between recovery time and space overhead with selective logging. Compared to state-of-the-art approaches, extensive evaluations show that SWIFT significantly accelerates failure recovery without affecting training throughput and model accuracy.


ACKNOWLEDGMENTS

This project was supported in part by grants from Hong Kong RGC under the contracts HKU 17204423 and C7004-22G (CRF).

REFERENCES

- [1] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2019.
- [2] M. Shoeybi, M. Patwary, R. Puri, P. LeGresley, J. Casper, and B. Catanzaro, "Megatron-LM: Training Multi-billion Parameter Language Models Using Model Parallelism," *arXiv preprint*, 2019.
- [3] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei, "Language Models are Few-Shot Learners," in *Proceedings of Advances in Neural Information Processing Systems*, 2020.
- [4] M. Jeon, S. Venkataraman, A. Phanishayee, J. Qian, W. Xiao, and F. Yang, "Analysis of Large-scale Multi-tenant GPU Clusters for DNN Training Workloads," in *Proceedings of the 2019 USENIX Annual Technical Conference*, 2019.
- [5] D. Narayanan, M. Shoeybi, J. Casper, P. LeGresley, M. Patwary, V. Korthikanti, D. Vainbrand, P. Kashinkunti, J. Bernauer, B. Catanzaro, A. Phanishayee, and M. Zaharia, "Efficient large-scale language model training on gpu clusters using megatron-lm," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2021.
- [6] K. Maeng, S. Bharuka, I. Gao, M. Jeffrey, V. Saraph, B.-Y. Su, C. Trippel, J. Yang, M. Rabbat, B. Lucia, and C.-J. Wu, "Understanding and Improving Failure Tolerant Training for Deep Learning Recommendation with Partial Recovery," in *Proceedings of the 4th Conference on Machine Learning and Systems*, 2021.
- [7] A. Eisenman, K. K. Matam, S. Ingram, D. Mudigere, R. Krishnamoorthi, K. Nair, M. Smelyanskiy, and M. Annavaram, "Check-N-Run: a Checkpointing System for Training Deep Learning Recommendation Models," in *Proceedings of the 19th USENIX Symposium on Networked Systems Design and Implementation*, 2022.
- [8] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng, "TensorFlow: A System for Large-scale Machine Learning," in *Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation*, 2016.
- [9] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "PyTorch: An Imperative Style, High-Performance Deep Learning Library," in *Proceedings of Advances in Neural Information Processing Systems*, 2019.
- [10] J. Mohan, A. Phanishayee, and V. Chidambaram, "CheckFreq: Frequent, Fine-Grained DNN Checkpointing," in *Proceedings of the 19th USENIX Conference on File and Storage Technologies*, 2021.
- [11] T. H. Authors, "Elastic horovod," 2022, https://horovod.readthedocs.io/en/stable/elastic_include.html.
- [12] S. Fan, Y. Rong, C. Meng, Z. Cao, S. Wang, Z. Zheng, C. Wu, G. Long, J. Yang, L. Xia, L. Diao, X. Liu, and W. Lin, "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*, 2021.
- [13] M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauly, M. J. Franklin, S. Shenker, and I. Stoica, "Resilient Distributed Datasets: A Fault-tolerant Abstraction for In-memory Cluster Computing," in *Proceedings of the 9th USENIX Symposium on Networked Systems Design and Implementation*, 2012.
- [14] S. Wang, J. Liagouris, R. Nishihara, P. Moritz, U. Misra, A. Tumanov, and I. Stoica, "Lineage Stash: Fault Tolerance off the Critical Path," in *Proceedings of the 27th ACM Symposium on Operating Systems Principles*, 2019.
- [15] J.-H. Hwang, M. Balazinska, A. Rasin, U. Cetintemel, M. Stonebraker, and S. Zdonik, "High-availability Algorithms for Distributed Stream Processing," in *Proceedings of International Conference on Data Engineering*, 2005.
- [16] Y. Shen, G. Chen, H. Jagadish, W. Lu, B. C. Ooi, and B. M. Tudor, "Fast Failure Recovery in Distributed Graph Processing Systems," in *Proceedings of the Very Large Data Base Endowment*, 2014.
- [17] D. Narayanan, A. Harlap, A. Phanishayee, V. Seshadri, N. R. Devanur, G. R. Ganger, P. B. Gibbons, and M. Zaharia, "PipeDream: Generalized Pipeline Parallelism for DNN Training," in *Proceedings of the 27th ACM Symposium on Operating Systems Principles*, 2019.
- [18] S. Zagoruyko and N. Komodakis, "Wide Residual Networks," *arXiv preprint*, 2016.
- [19] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale," in *Proceedings of International Conference on Learning Representations*, 2020.
- [20] J. Chen, X. Pan, R. Monga, S. Bengio, and R. Jozefowicz, "Revisiting Distributed Synchronous SGD," *arXiv preprint*, 2016.
- [21] P. Goyal, P. Dollár, R. Girshick, P. Noordhuis, L. Wesolowski, A. Kyrola, A. Tulloch, Y. Jia, and K. He, "Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour," *arXiv preprint*, 2017.
- [22] Y. You, J. Li, S. Reddi, J. Hseu, S. Kumar, S. Bhojanapalli, X. Song, J. Demmel, K. Keutzer, and C.-J. Hsieh, "Large Batch Optimization for Deep Learning: Training BERT in 76 Minutes," in *Proceedings of International Conference on Learning Representations*, 2020.
- [23] Z. Jia, M. Zaharia, and A. Aiken, "Beyond Data and Model Parallelism for Deep Neural Networks," in *Proceedings of the 2nd Conference on Machine Learning and Systems*, 2019.
- [24] Y. Huang, Y. Cheng, A. Bapna, O. Firat, D. Chen, M. Chen, H. Lee, J. Ngiam, Q. V. Le, Y. Wu, and z. Chen, "GPipe: Efficient Training of Giant Neural Networks using Pipeline Parallelism," in *Proceedings of Advances in Neural Information Processing Systems*, 2019.
- [25] D. Narayanan, A. Phanishayee, K. Shi, X. Chen, and M. Zaharia, "Memory-efficient Pipeline-parallel DNN Training," in *Proceedings of International Conference on Machine Learning*, 2021.
- [26] J. Yuan, X. Li, C. Cheng, J. Liu, R. Guo, S. Cai, C. Yao, F. Yang, X. Yi, C. Wu, H. Zhang, and J. Zhao, "OneFlow: Redesign the Distributed Deep Learning Framework from Scratch," *arXiv preprint*, 2021.
- [27] L. Zheng, Z. Li, H. Zhang, Y. Zhuang, Z. Chen, Y. Huang, Y. Wang, Y. Xu, D. Zhuo, E. P. Xing, J. E. Gonzalez, and I. Stoica, "Alpa: Automating Inter- and Intra-Operator Parallelism for Distributed Deep Learning," in *Proceedings of the 16th USENIX Symposium on Operating Systems Design and Implementation*, 2022.
- [28] G. Ma, Y. Xiao, T. Willke, N. Ahmed, S. Nazarian, and P. Bogdan, "A Distributed Graph-theoretic Framework for Automatic Parallelization in Multi-core Systems," in *Proceedings of Machine Learning and Systems*, 2021.
- [29] Y. Xiao, S. Nazarian, and P. Bogdan, "Self-optimizing and Self-programming Computing Systems: A Combined Compiler, Complex Networks, and Machine Learning Approach," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 2019.
- [30] —, "Plasticity-on-Chip Design: Exploiting Self-Similarity for Data Communications," *IEEE Transactions on Computers*, 2021.
- [31] S. Rajbhandari, O. Ruwase, J. Rasley, S. Smith, and Y. He, "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*, 2021.
- [32] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in *Proceedings of International Conference on Learning Representations*, 2015.
- [33] "Checkfreq code at github," 2022, <https://github.com/msr-fiddle/CheckFreq>.
- [34] T. Chen, M. Li, Y. Li, M. Lin, N. Wang, M. Wang, T. Xiao, B. Xu, C. Zhang, and Z. Zhang, "MXNet: A Flexible and Efficient Machine Learning Library for Heterogeneous Distributed Systems," *arXiv preprint*, 2015.
- [35] A. Qiao, B. Aragam, B. Zhang, and E. Xing, "Fault Tolerance in Iterative-Convergent Machine Learning," in *International Conference on Machine Learning*, 2019.
- [36] P. Moritz, R. Nishihara, S. Wang, A. Tumanov, R. Liaw, E. Liang, M. Elibol, Z. Yang, W. Paul, M. I. Jordan, and I. Stoica, "Ray: A Distributed Framework for Emerging AI Applications," in *Proceedings of the 13th USENIX Symposium on Operating Systems Design and Implementation*, 2018.
- [37] R. D. Schlichting and F. B. Schneider, "Fail-stop Processors: An Approach to Designing Fault-tolerant Computing Systems," *ACM Transactions on Computer Systems*, 1983.
- [38] I. Loshchilov and F. Hutter, "Decoupled Weight Decay Regularization," in *Proceedings of International Conference on Learning Representations*, 2019.
- [39] S. J. Reddi, S. Kale, and S. Kumar, "On the Convergence of Adam and Beyond," in *Proceedings of International Conference on Learning Representations*, 2018.
- [40] R. Harte, *Invertibility and Singularity for Bounded Linear Operators*. Courier Dover Publications, 2016.
- [41] A. N. Gomez, M. Ren, R. Urtasun, and R. B. Grosse, "The reversible residual network: Backpropagation without storing activations," in *Proceedings of International Conference on Neural Information Processing Systems*, 2017.
- [42] K. Shvachko, H. Kuang, S. Radia, and R. Chansler, "The Hadoop Distributed File System," in *Proceedings of IEEE 26th Symposium on Mass Storage Systems and Technologies*, 2010.
- [43] M. Rhu, N. Gimelshein, J. Clemons, A. Zulfiqar, and S. W. Keckler, "vDNN: Virtualized Deep Neural Networks for Scalable, Memory-efficient Neural Network Design," in *Proceedings of the 49th Annual IEEE/ACM International Symposium on Microarchitecture*, 2016.
- [44] J. Sylvain, "Optimized inter-gpu collective operations with nccl 2," 2017, <https://developer.nvidia.com/nccl>.
- [45] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "Imagenet Large Scale Visual Recognition Challenge," in *Proceedings of International Journal of Computer Vision*, 2015.
- [46] K. He, X. Zhang, S. Ren, and J. Sun, "Identity Mappings in Deep Residual Networks," in *Proceedings of European Conference on Computer Vision*, 2016.
- [47] L. Mai, G. Li, M. Wagenländer, K. Fertakis, A.-O. Brabete, and P. Pietzuch, "Kungfu: Making Training in Distributed Machine Learning Adaptive," in *Proceedings of the 14th USENIX Symposium on Operating Systems Design and Implementation*, 2020.
- [48] Y. Peng, Y. Bao, Y. Chen, C. Wu, and C. Guo, "Optimus: an Efficient Dynamic Resource Scheduler for Deep Learning Clusters," in *Proceedings of the 13th European Conference on Computer Systems*, 2018.
- [49] A. Qiao, S. K. Choe, S. J. Subramanya, W. Neiswanger, Q. Ho, H. Zhang, G. R. Ganger, and E. P. Xing, "Pollux: Co-adaptive Cluster Scheduling for Goodput-optimized Deep Learning," in *Proceedings of the 15th USENIX Symposium on Operating Systems Design and Implementation*, 2021.
- [50] P. Xie, J. K. Kim, Q. Ho, Y. Yu, and E. Xing, "Orpheus: Efficient Distributed Machine Learning via System and Algorithm Co-design," in *Proceedings of ACM Symposium on Cloud Computing*, 2018.
- [51] P. Micikevicius, S. Narang, J. Alben, G. Diamos, E. Elsen, D. Garcia, B. Ginsburg, M. Houston, O. Kuchaiev, G. Venkatesh, and H. Wu, "Mixed Precision Training," in *Proceedings of International Conference on Learning Representations*, 2018.



Yuchen Zhong received the BEng degree in computer science from Tongji University, China, in 2020. He is working toward a PhD with the Department of Computer Science at the University of Hong Kong. His research interests center on distributed machine learning systems.



Guangming Sheng received the BEng degree from the Hongyi Honor College (Computer Science), Wuhan University, China, in 2022. He is working toward a PhD with the Department of Computer Science at the University of Hong Kong. His research interests include distributed machine learning systems and graph neural network systems.



Juncheng Liu is a seasoned software and deep learning system developer, serving as a core developer for the deep learning framework OneFlow. With over a decade of experience, Juncheng specializes in optimizing the performance of deep learning systems.



Jinhui Yuan is the founder and CEO of OneFlow Inc and the chief architect of the open-sourced Dubhe AI platform proposed by Zhejiang Lab. He received his PhD in computer science from Tsinghua University in 2008 (with honor) and has published numerous papers on computer vision and multimedia retrieval in top conferences and prestigious journals. Jinhui joined MSRA in 2013 and researched distributed systems for large-scale machine learning. In 2014, he and his collaborators invented LightLDA, which was once the fastest system for training large-scale topic models. Since 2016, Jinhui focused on building a new generation deep learning framework OneFlow.



Chuan Wu (Senior Member, IEEE) received the PhD from the Department of Electrical and Computer Engineering, University of Toronto, Canada, in 2008. Since September 2008, she has been with the Department of Computer Science, University of Hong Kong, where she is currently a professor. Her current research interests include the areas of cloud computing, distributed machine learning systems, network function virtualization, and intelligent elderly care technologies.

APPENDIX A

UPDATE-UNDO ALGORITHMS

Algorithm 3 SGD

- 1: **Input:** learning rate sequence $\{\eta_t\}_{t=1}^T$; weight decay $\lambda > 0$.
 - 2: **Initialize:** $x_1 \in \mathbb{R}^d$.
 - 3: **for** $i = 1$ **to** T **do**
 - 4: $g_t = \nabla f(x_t)$
 - 5: $x_{t+1} = x_t - \eta_t(g_t + \lambda x_t)$
 - 6: **end for**
-

Algorithm 4 Undo SGD

- 1: **Input:** learning rate sequence $\{\eta_t\}_{t=1}^T$; weight decay $\lambda > 0$; $x_{t+1} \in \mathbb{R}^d$; $g_t \in \mathbb{R}^d$.
 - 2: $x_t = (x_{t+1} + \eta_t g_t) / (1 - \eta_t \lambda)$
-

Algorithm 5 Adam [32]

- 1: **Input:** learning rate sequence $\{\eta_t\}_{t=1}^T$; weight decay $\lambda > 0$; Exponential decay rates for moment estimates $0 \leq \beta_1, \beta_2 < 1$; $\epsilon = 10^{-8}$.
 - 2: **Initialize:** $x_1 \in \mathbb{R}^d$; $m_0 = 0 \in \mathbb{R}^d$; $v_0 = 0 \in \mathbb{R}^d$.
 - 3: **for** $i = 1$ **to** T **do**
 - 4: $g_t = \nabla f(x_t)$
 - 5: $g'_t = g_t + \lambda \cdot x_t$
 - 6: $m_t = \beta_1 \cdot m_{t-1} + (1 - \beta_1) \cdot g'_t$
 - 7: $v_t = \beta_2 \cdot v_{t-1} + (1 - \beta_2) \cdot g_t'^2$
 - 8: $\hat{m}_t = m_t / (1 - \beta_1^t)$
 - 9: $\hat{v}_t = v_t / (1 - \beta_2^t)$
 - 10: $x_t = x_{t-1} - \eta_t \cdot \hat{m}_t / (\sqrt{\hat{v}_t} + \epsilon)$
 - 11: **end for**
-

Algorithm 6 Undo Adam

- 1: **Input:** learning rate sequence $\{\eta_t\}_{t=1}^T$; weight decay $\lambda > 0$; Exponential decay rates for moment estimates $0 \leq \beta_1, \beta_2 < 1$; $\epsilon = 10^{-8}$; $x_{t+1} \in \mathbb{R}^d$; $g_t \in \mathbb{R}^d$; $m_t \in \mathbb{R}^d$; $v_t \in \mathbb{R}^d$.
 - 2: $\hat{m}_t = m_t / (1 - \beta_1^t)$
 - 3: $\hat{v}_t = v_t / (1 - \beta_2^t)$
 - 4: $x_t = x_{t+1} + \eta_t \cdot \hat{m}_t / (\sqrt{\hat{v}_t} + \epsilon)$
 - 5: $g'_t = g_t + \lambda \cdot x_t$
 - 6: $m_{t-1} = (m_t - (1 - \beta_1) \cdot g'_t) / \beta_1$
 - 7: $v_{t-1} = (v_t - (1 - \beta_2) \cdot g_t^2) / \beta_2$
-

Algorithm 7 AdamW [38]

- 1: **Input:** learning rate sequence $\{\eta_t\}_{t=1}^T$; weight decay $\lambda > 0$; Exponential decay rates for moment estimates $0 \leq \beta_1, \beta_2 < 1$; $\epsilon = 10^{-8}$.
 - 2: **Initialize:** $x_1 \in \mathbb{R}^d$; $m_0 = 0 \in \mathbb{R}^d$; $v_0 = 0 \in \mathbb{R}^d$.
 - 3: **for** $i = 1$ **to** T **do**
 - 4: $g_t = \nabla f(x_t)$
 - 5: $m_t = \beta_1 \cdot m_{t-1} + (1 - \beta_1) \cdot g_t$
 - 6: $v_t = \beta_2 \cdot v_{t-1} + (1 - \beta_2) \cdot g_t^2$
 - 7: $\hat{m}_t = m_t / (1 - \beta_1^t)$
 - 8: $\hat{v}_t = v_t / (1 - \beta_2^t)$
 - 9: $x_t = x_{t-1} - \eta_t \cdot (\alpha \cdot \hat{m}_t / (\sqrt{\hat{v}_t} + \epsilon) + \lambda \cdot x_{t-1})$
 - 10: **end for**
-

Algorithm 8 Undo AdamW

- 1: **Input:** learning rate sequence $\{\eta_t\}_{t=1}^T$; weight decay $\lambda > 0$; Exponential decay rates for moment estimates $0 \leq \beta_1, \beta_2 < 1$; $\epsilon = 10^{-8}$; $x_1 \in \mathbb{R}^d$; $m_0 = 0 \in \mathbb{R}^d$; $v_0 = 0 \in \mathbb{R}^d$.
 - 2: $\hat{m}_t = m_t / (1 - \beta_1^t)$
 - 3: $\hat{v}_t = v_t / (1 - \beta_2^t)$
 - 4: $x_{t-1} = (x_t + \eta_t \cdot (\alpha \cdot \hat{m}_t / \sqrt{\hat{v}_t} + \epsilon)) / (1 - \eta_t \cdot \lambda)$
 - 5: $m_{t-1} = (m_t - (1 - \beta_1) \cdot g_t) / \beta_1$
 - 6: $v_{t-1} = (v_t - (1 - \beta_2) \cdot g_t^2) / \beta_2$
-

APPENDIX B

SUPPORT MULTIPLE FAILURES AND CASCADING FAILURES

SWIFT can tolerate multiple simultaneous failures. Even if there are multiple failures, replication-based recovery can be used if there is still a copy of the lost model state. For logging-based recovery, if the failed workers (on failure machines) constitute a consecutive portion of the pipeline, they can be recovered jointly; if they span different portions of the pipeline, we can recover those portions independently.

SWIFT can also tolerate cascading failures. If another failure occurs while recovering from one failure, you can use replication-based recovery as long as that lost state is still backed up. For logging-based recovery, during the recovery of a failed machine, if another machine crashes which hosts workers which are connected in the pipeline to workers on the being-recovered machine, the ongoing recovery process is aborted, and the two machines are recovered together; if the newly crashed machine does not host workers connected to those on the being-recovered machine, it does not affect the ongoing recovery process, and the new failure can be recovered independently.

APPENDIX C

GROUPING DETAILS

We use the model and settings as described in §7.1. We profile the computation time for 5 iterations of each stage and take the average time. The communication size between each stage can be calculated directly based on the model configuration and batch size.

TABLE 6: BERT-128's grouping results. The number in the outcome represents the machine rank.

Storage limit (Bytes)	Outcome
5.00E+11	[[0], [1], [2], [3], [4], [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15]]
4.00E+11	[[0], [1], [2], [3], [4], [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15]]
3.50E+11	[[0], [1], [2], [3], [4], [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15]]
3.00E+11	[[0], [1], [2], [3], [4], [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15]]
2.50E+11	[[0], [1], [2], [3, 4], [5, 6], [7, 8], [9, 10], [11, 12], [13, 14, 15]]
2.20E+11	[[0], [1, 2], [3, 4], [5, 6], [7, 8], [9, 10, 11, 12], [13, 14, 15]]
1.50E+11	[[0, 1, 2], [3, 4], [5, 6, 7, 8], [9, 10, 11, 12], [13, 14, 15]]
1.00E+11	[[0, 1, 2], [3, 4, 5, 6, 7, 8], [9, 10, 11, 12, 13, 14, 15]]
8.00E+10	[[0, 1, 2, 3, 4, 5, 6, 7, 8], [9, 10, 11, 12, 13, 14, 15]]
5.00E+10	[[0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15]]

TABLE 7: ViT-128/32's grouping results. The number in the outcome represents the machine rank.

Storage limit (Bytes)	Outcome
1.40E+12	[[0], [1], [2], [3], [4], [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15]]
1.20E+12	[[0], [1], [2], [3], [4], [5], [6], [7], [8], [9], [10], [11, 12], [13, 14], [15]]
1.10E+12	[[0], [1], [2], [3], [4], [5], [6], [7], [8], [9], [10], [11, 12], [13, 14, 15]]
1.00E+12	[[0], [1], [2], [3], [4], [5], [6], [7, 8], [9, 10], [11, 12], [13, 14, 15]]
9.00E+11	[[0], [1], [2], [3], [4], [5], [6, 7], [8, 9], [10, 11, 12], [13, 14, 15]]
8.00E+11	[[0], [1], [2], [3], [4, 5], [6, 7], [8, 9], [10, 11, 12], [13, 14, 15]]
7.00E+11	[[0], [1], [2], [3], [4, 5], [6, 7], [8, 9], [10, 11, 12, 13, 14, 15]]
6.00E+11	[[0], [1], [2, 3], [4, 5], [6, 7], [8, 9], [10, 11, 12, 13, 14, 15]]
5.00E+11	[[0, 1], [2, 3], [4, 5], [6, 7, 8, 9], [10, 11, 12, 13, 14, 15]]
4.00E+11	[[0, 1], [2, 3, 4, 5], [6, 7, 8, 9], [10, 11, 12, 13, 14, 15]]
3.00E+11	[[0, 1], [2, 3, 4, 5], [6, 7, 8, 9, 10, 11, 12, 13, 14, 15]]
2.00E+11	[[0, 1, 2, 3, 4, 5], [6, 7, 8, 9, 10, 11, 12, 13, 14, 15]]
1.00E+11	[[0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15]]