

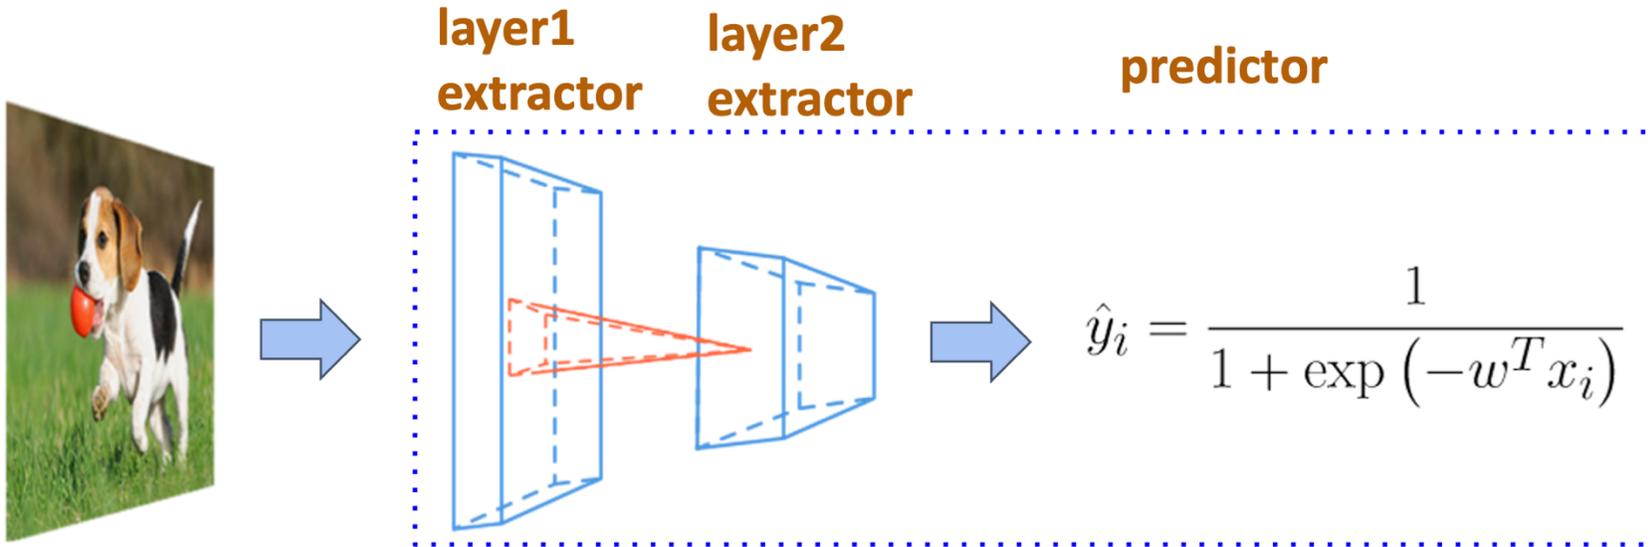
# **15-442/15-642: Machine Learning Systems**

## **Parallelization Part 1 (Data Parallelism and Zero Redundancy)**

**Tianqi Chen and Zhihao Jia**

Carnegie Mellon University

# Recap: DNN Training Overview



**Objective**

$$L(w) = \sum_{i=1}^n l(y_i, \hat{y}_i) + \lambda \|w\|^2$$

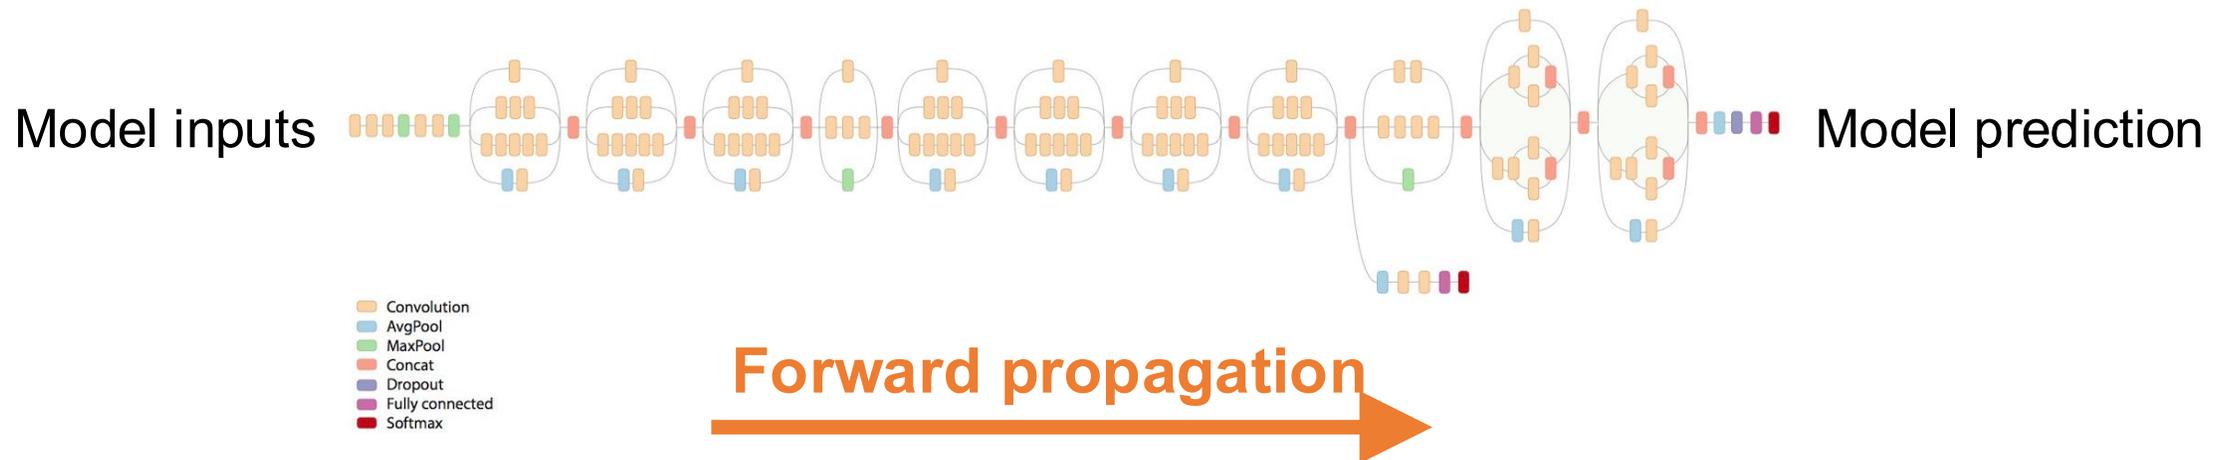
**Training**

$$w \leftarrow w - \eta \nabla_w L(w)$$

# DNN Training Process

Train ML models through many iterations of 3 stages

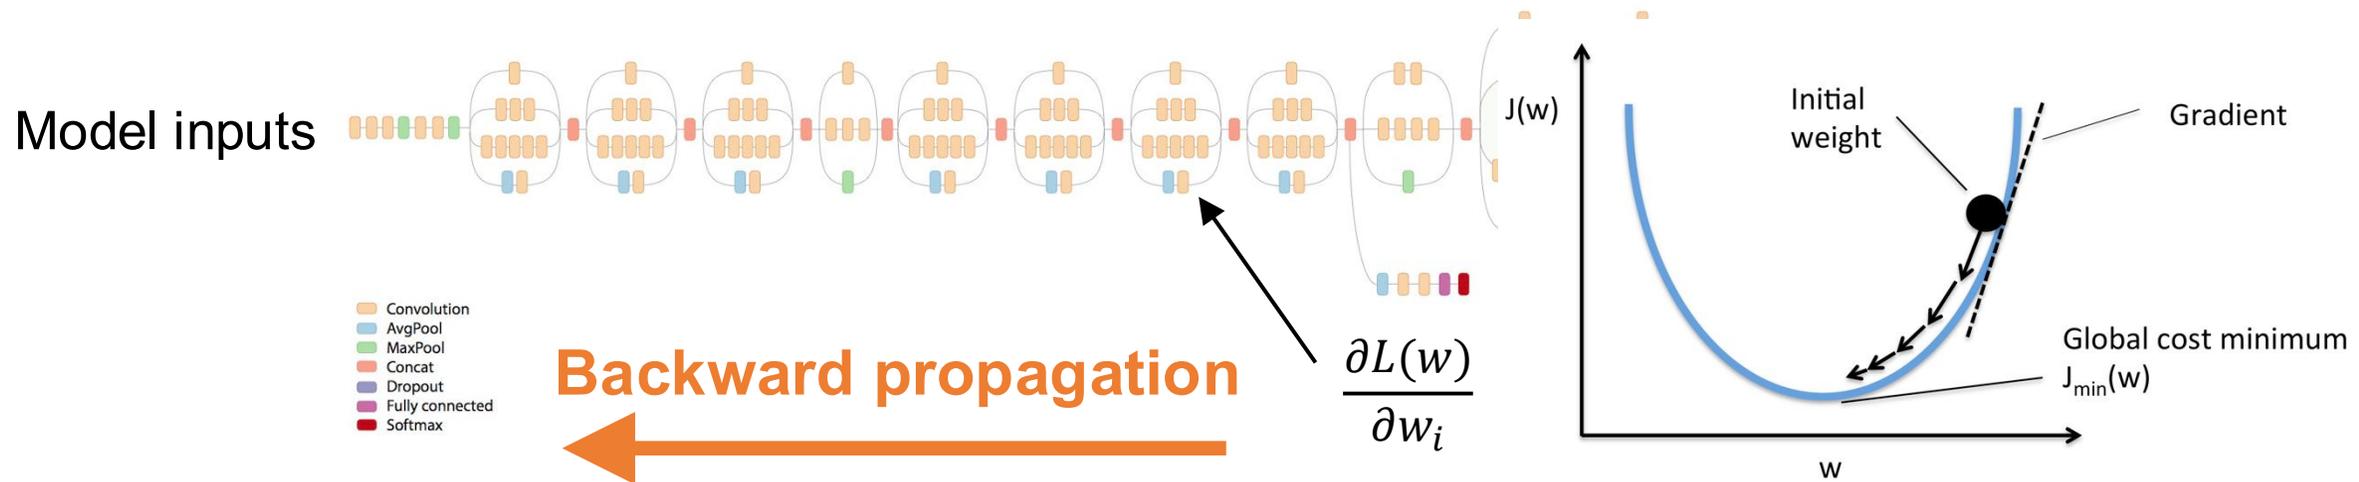
1. **Forward propagation**: apply model to a batch of input samples and run calculation through operators to produce a prediction
2. **Backward propagation**: run the model in reverse to produce error for each trainable weight
3. **Weight update**: use the loss value to update model weights



# DNN Training Process

Train ML models through many iterations of 3 stages

- 1. Forward propagation:** apply model to a batch of input samples and run calculation through operators to produce a prediction
- 2. Backward propagation:** run the model in reverse to produce a gradient for each trainable weight
- 3. Weight update:** use the loss value to update model weights



# DNN Training Process

Train ML models through many iterations of 3 stages

1. **Forward propagation**: apply model to a batch of input samples and run calculation through operators to produce a prediction
2. **Backward propagation**: run the model in reverse to produce a gradient for each trainable weight
3. **Weight update**: use the gradients to update model weights

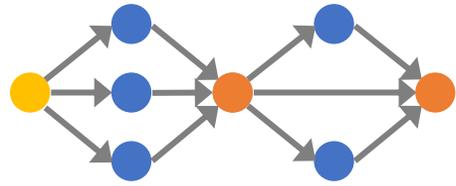
$$w_i := w_i - \gamma \frac{\partial L(w)}{\partial w_i} = w_i - \frac{\gamma}{n} \sum_{j=1}^n \frac{\partial l_i(w)}{\partial w_i}$$

Gradients of individual samples

# How can we parallelize DNN training?

$$w_i := w_i - \gamma \nabla L(w_i) = w_i - \frac{\gamma}{n} \sum_{j=1}^n \nabla L_j(w_i)$$

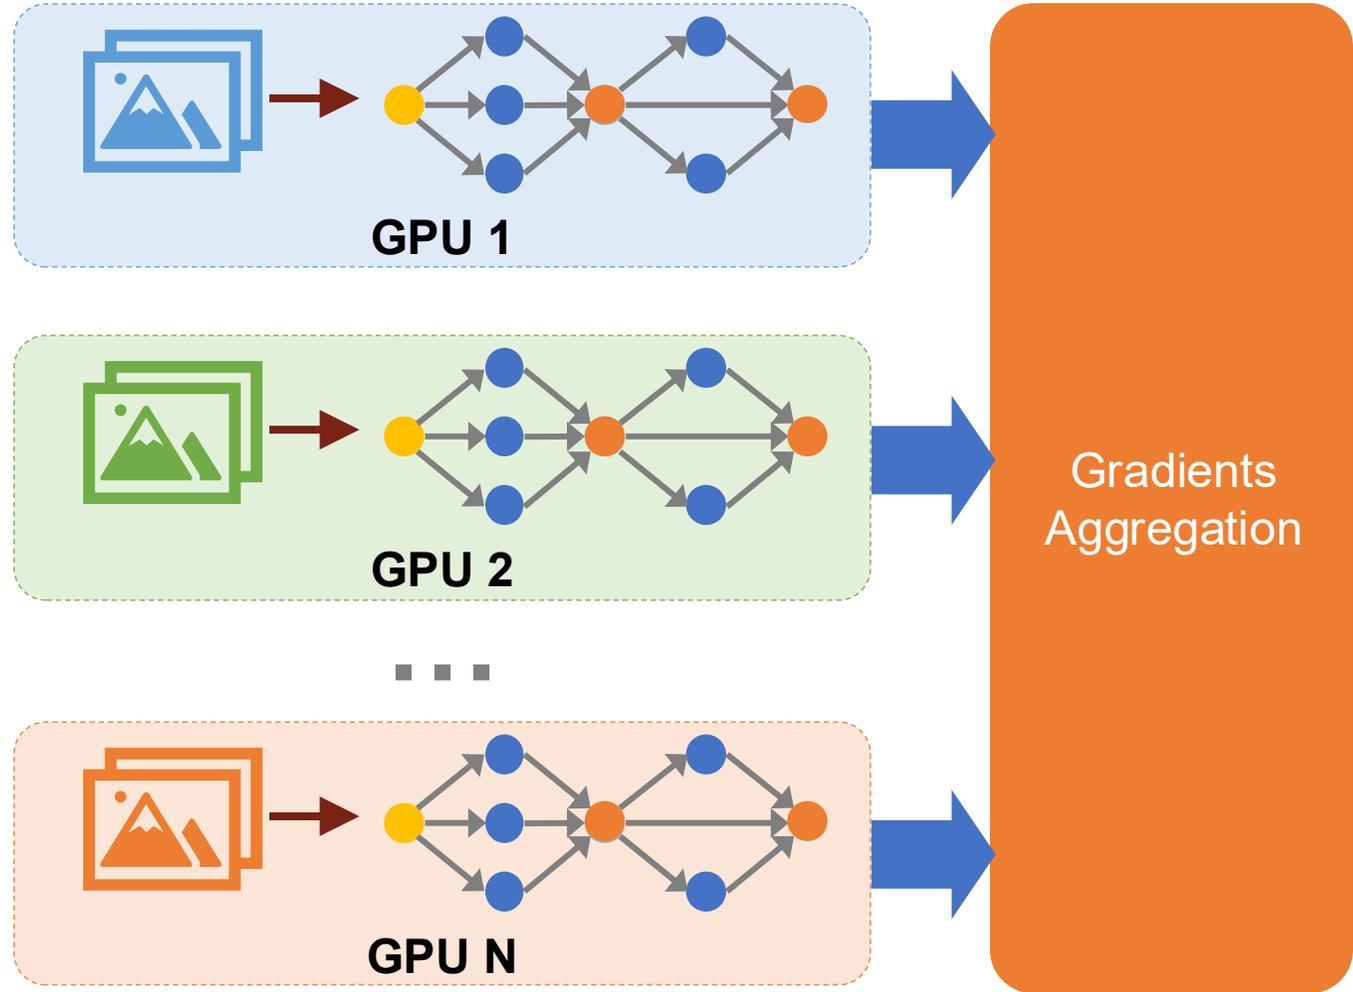
# Data Parallelism



Training Dataset

$$w_i := w_i - \gamma \nabla L(w_i) = w_i - \frac{\gamma}{n} \sum_{j=1}^n \nabla L_j(w_i)$$

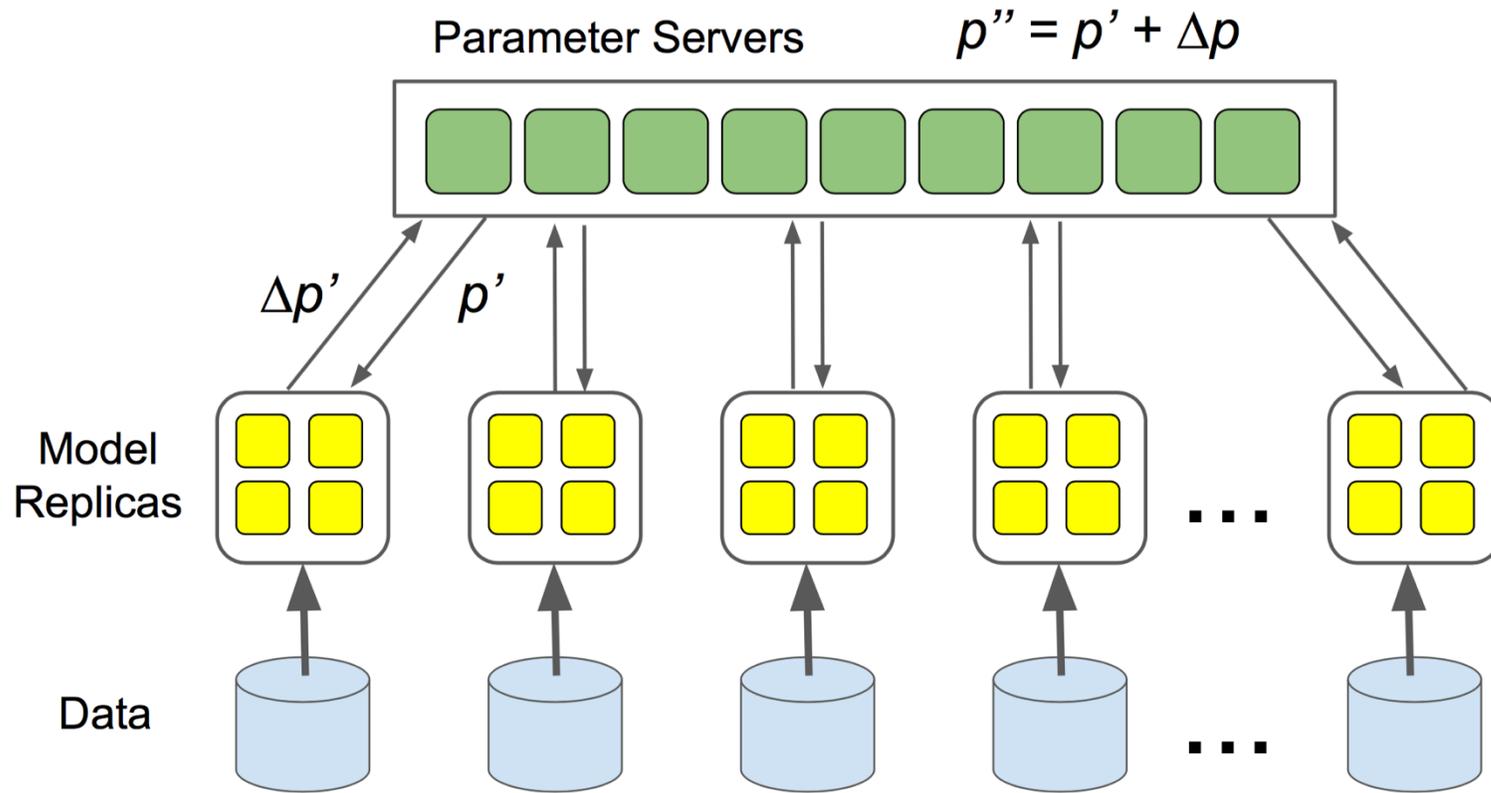
1. Partition training data into batches



2. Compute the gradients of each batch on a GPU

3. Aggregate gradients across GPUs

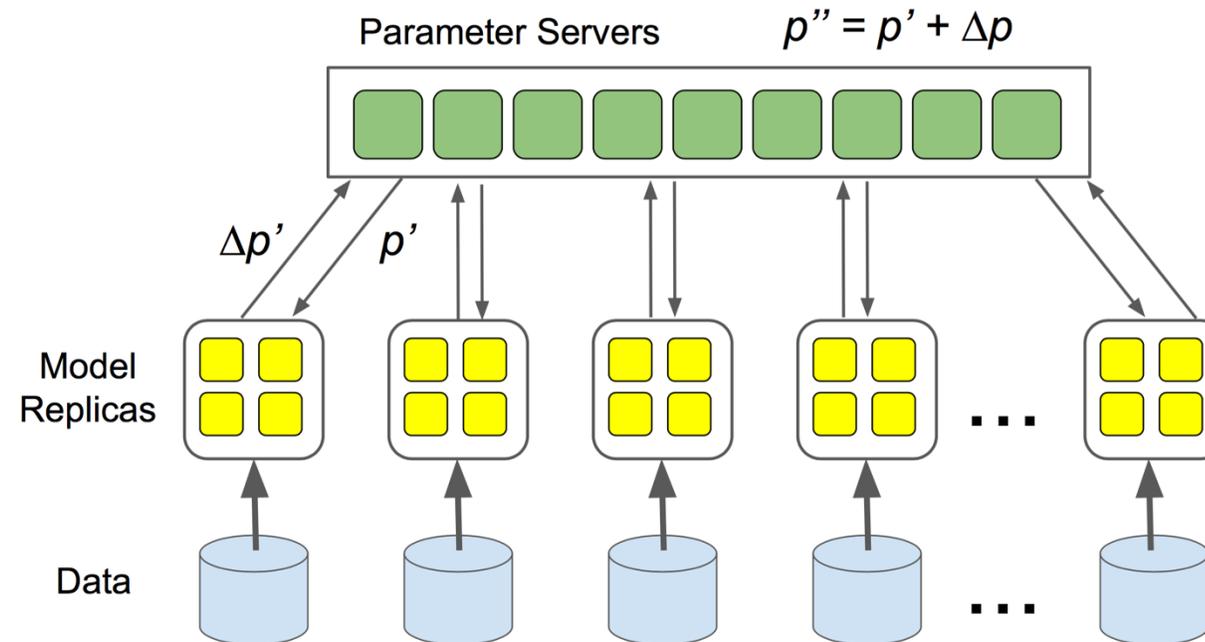
# Data Parallelism: Parameter Server



Workers push gradients to parameter servers and pull updated parameters back

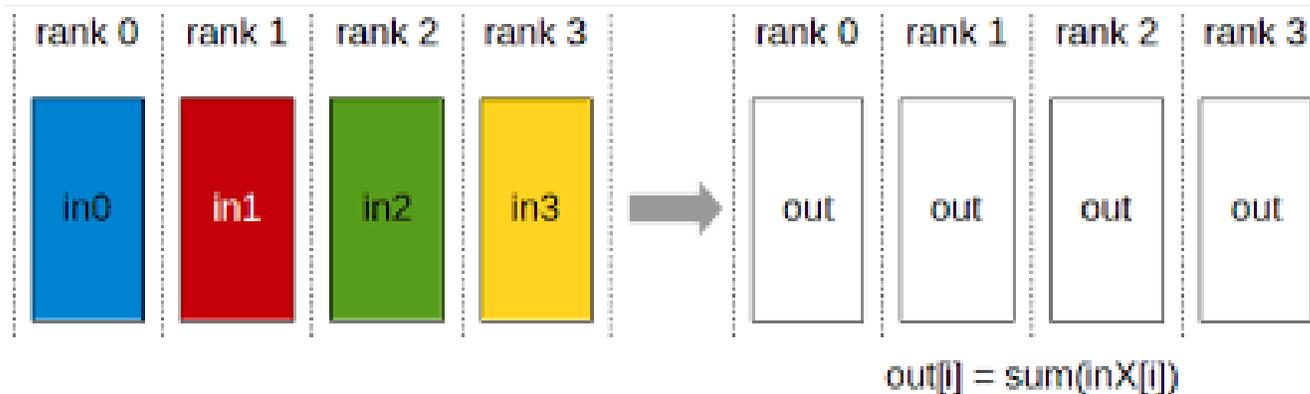
# Inefficiency of Parameter Server

- **Centralized communication:** all workers communicate with parameter servers for weights update; cannot scale to large numbers of workers
- How can we decentralize communication in DNN training?



# Inefficiency of Parameter Server

- **Centralized communication**: all workers communicate with parameter servers for weights update; cannot scale to large numbers of workers
- How can we decentralize communication in DNN training?
- **AllReduce**: perform element-wise reduction across multiple devices

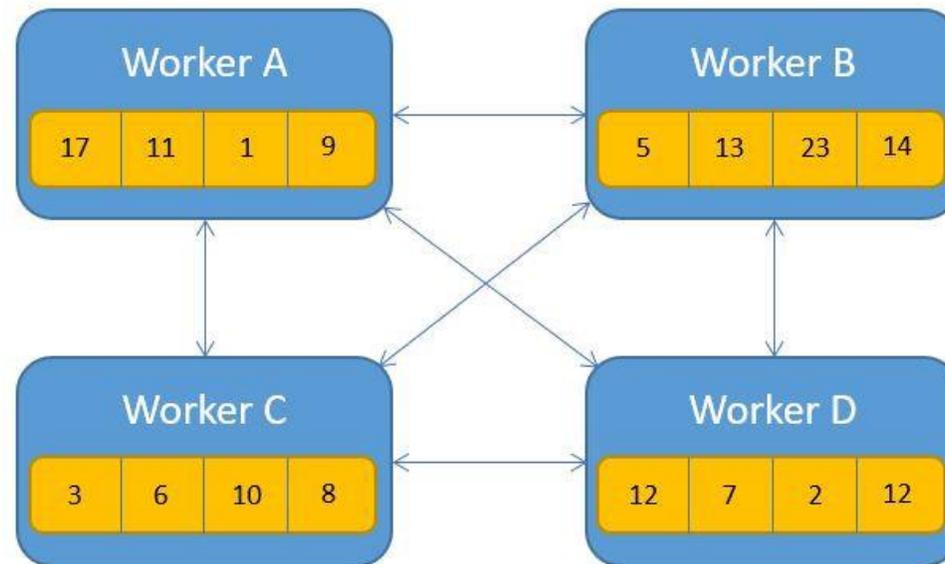


# Different Ways to Perform AllReduce

- Naïve AllReduce
- Ring AllReduce
- Tree AllReduce
- Butterfly AllReduce

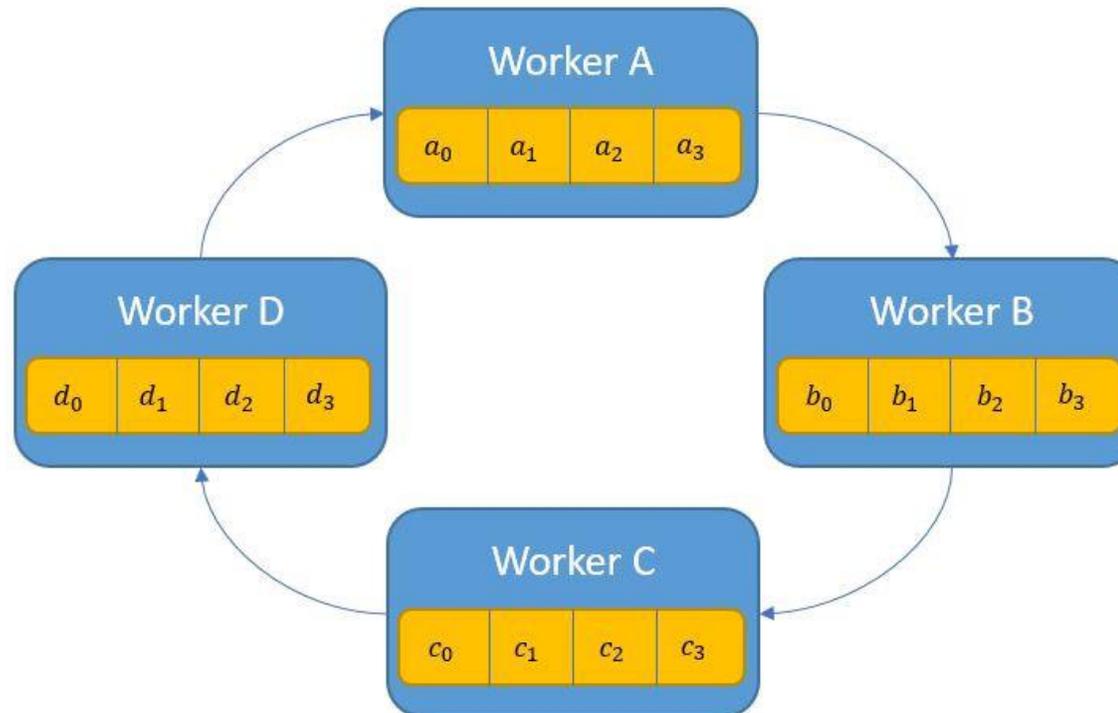
# Naïve AllReduce

- Each worker can send its local gradients to all other workers
- If we have N workers and each worker contains M parameters
- Overall communication:  $N * (N-1) * M$  parameters
- **Issue**: each worker communicates with all other workers; same scalability issue as parameter server



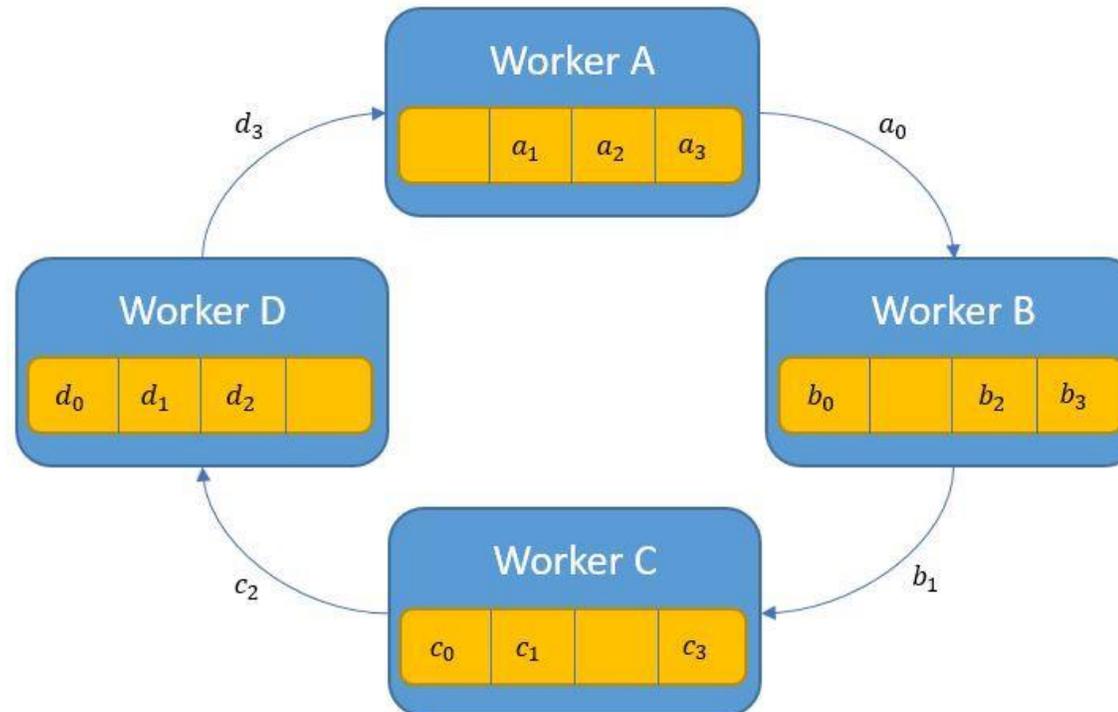
# Ring AllReduce

- Construct a ring of  $N$  workers, divide  $M$  parameters into  $N$  slices
- Step 1 (Aggregation): each worker send one slice ( $M/N$  parameters) to the next worker on the ring; repeat  $N$  times



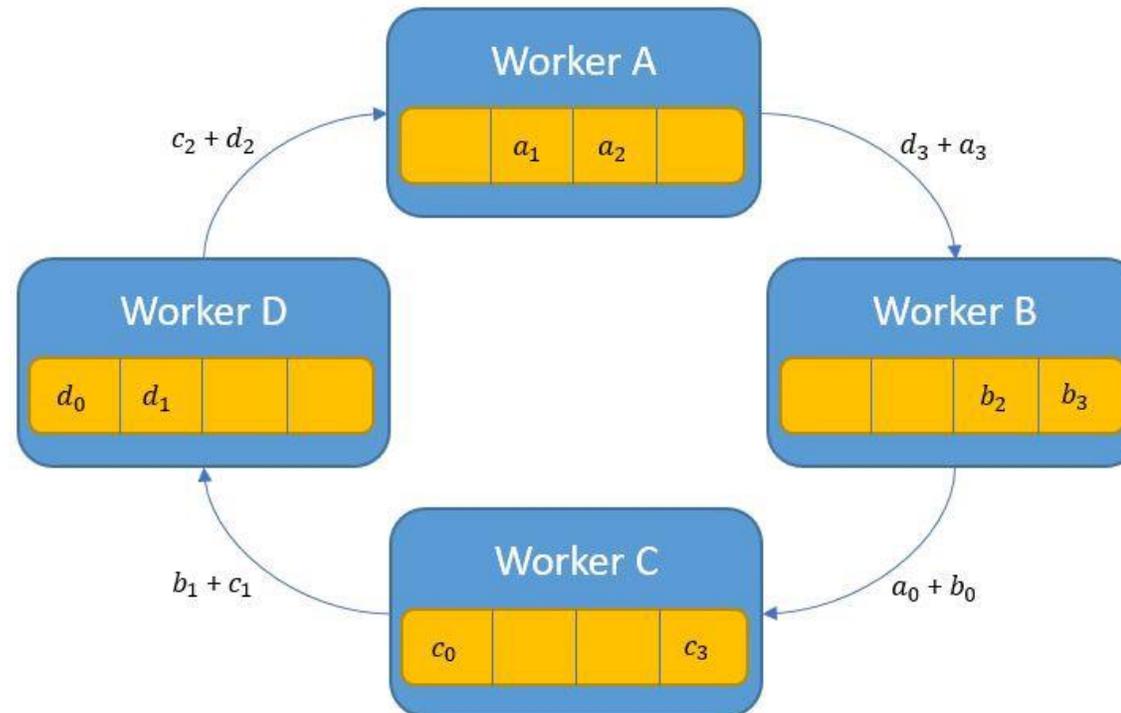
# Ring AllReduce

- Construct a ring of  $N$  workers, divide  $M$  parameters into  $N$  slices
- Step 1 (Aggregation): each worker send one slice ( $M/N$  parameters) to the next worker on the ring; repeat  $N$  times



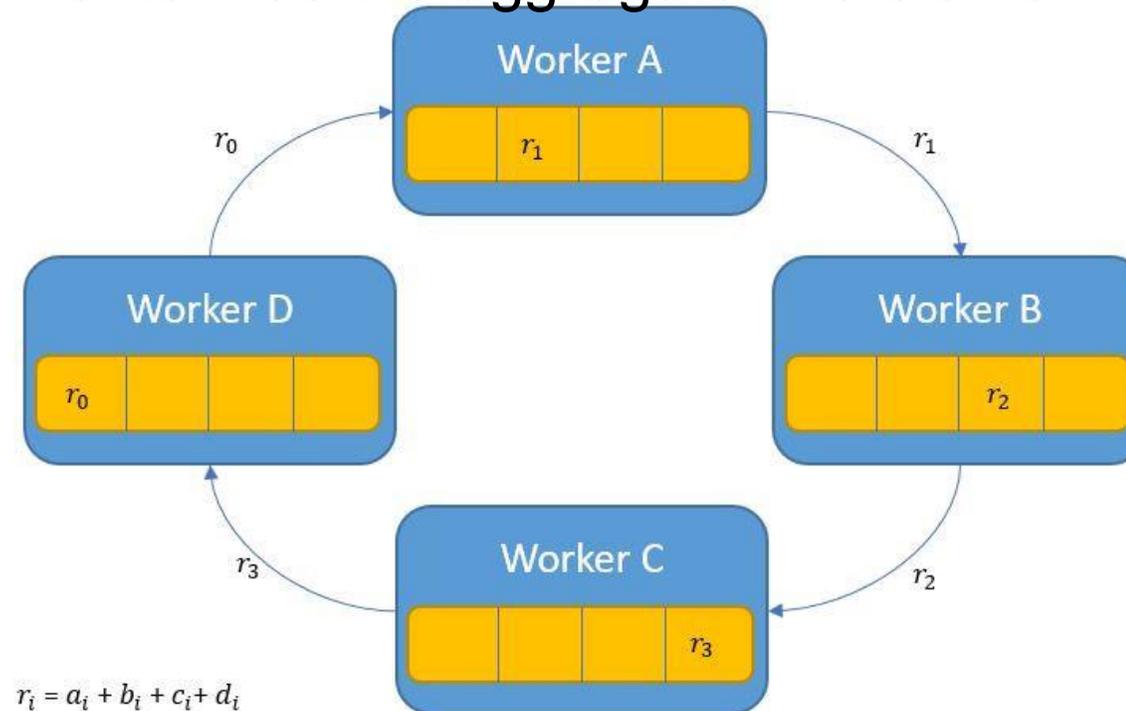
# Ring AllReduce

- Construct a ring of N workers, divide M parameters into N slices
- Step 1 (Aggregation): each worker send one slice (M/N parameters) to the next worker on the ring; repeat N times



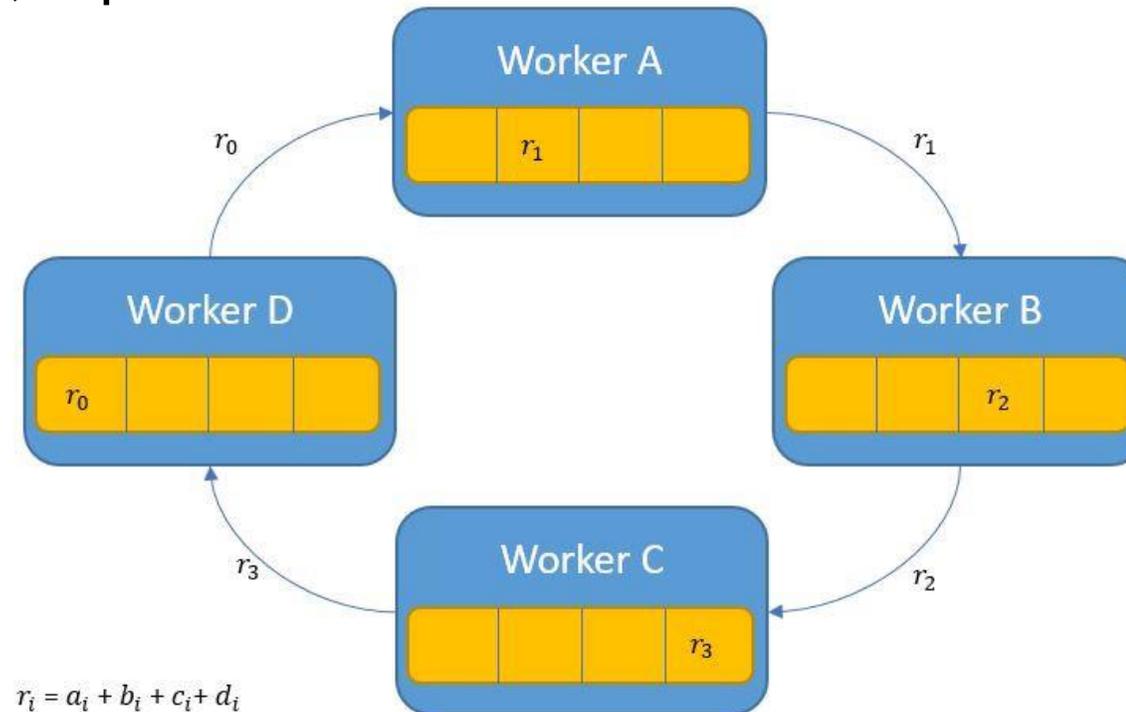
# Ring AllReduce

- Construct a ring of N workers, divide M parameters into N slices
- Step 1 (Aggregation): each worker send one slice (M/N parameters) to the next worker on the ring; repeat N times
- After step 1, each worker has the aggregated version of M/N parameters



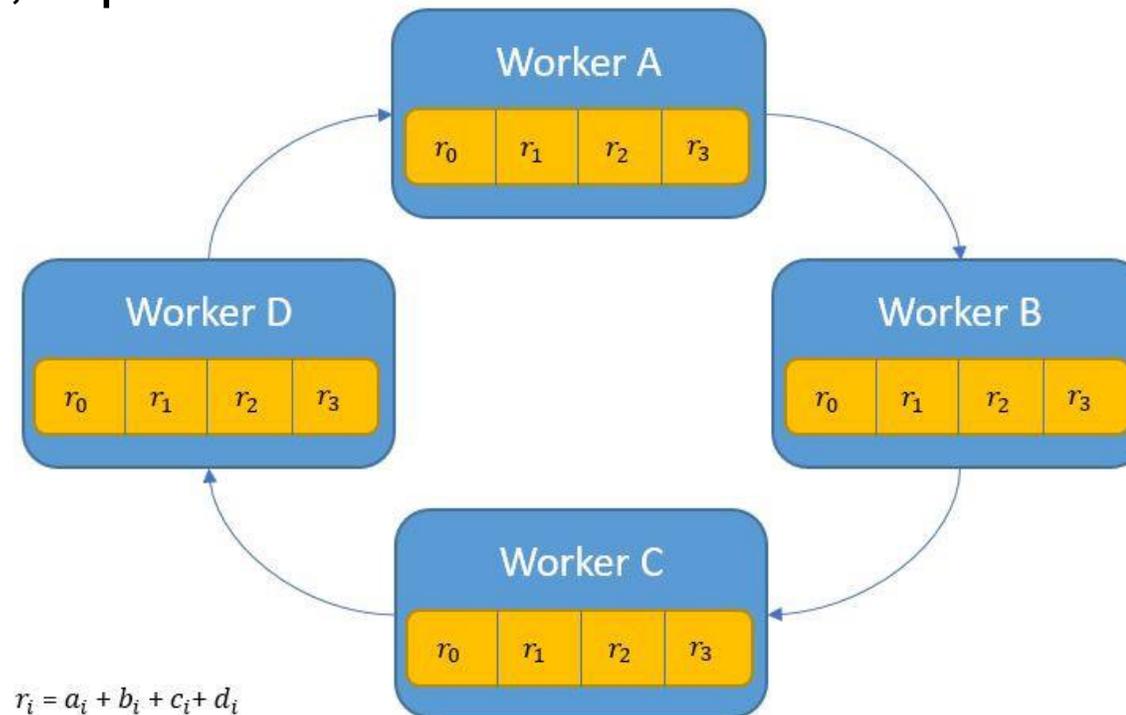
# Ring AllReduce

- Construct a ring of N workers, divide M parameters into N slices
- Step 1 (Aggregation): each worker send one slice (M/N parameters) to the next worker on the ring; repeat N times
- Step 2 (Broadcast): each worker send one slice of aggregated parameters to the next worker; repeat N times



# Ring AllReduce

- Construct a ring of N workers, divide M parameters into N slices
- Step 1 (Aggregation): each worker send one slice (M/N parameters) to the next worker on the ring; repeat N times
- Step 2 (Broadcast): each worker send one slice of aggregated parameters to the next worker; repeat N times

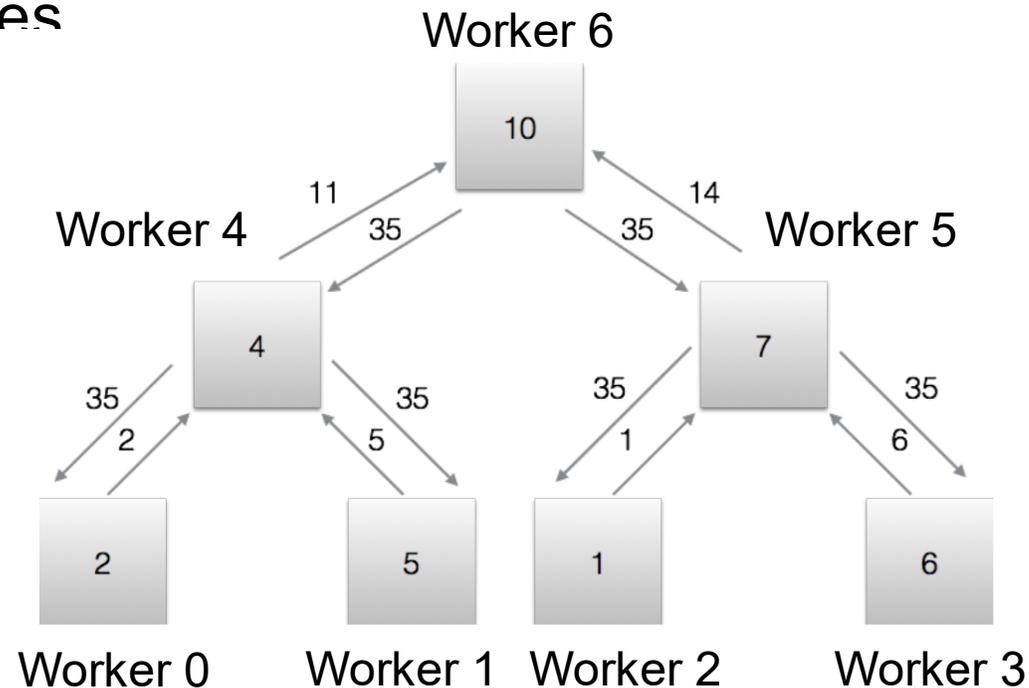


# Ring AllReduce

- Construct a ring of  $N$  workers, divide  $M$  parameters into  $N$  slices
- Step 1 (Aggregation): each worker send one slice ( $M/N$  parameters) to the next worker on the ring; repeat  $N$  times
- Step 2 (Broadcast): each worker send one slice of aggregated parameters to the next worker; repeat  $N$  times
- Overall communication:  $2 * M * N$  parameters
  - Aggregation:  $M * N$  parameters
  - Broadcast:  $M * N$  parameters

# Tree AllReduce

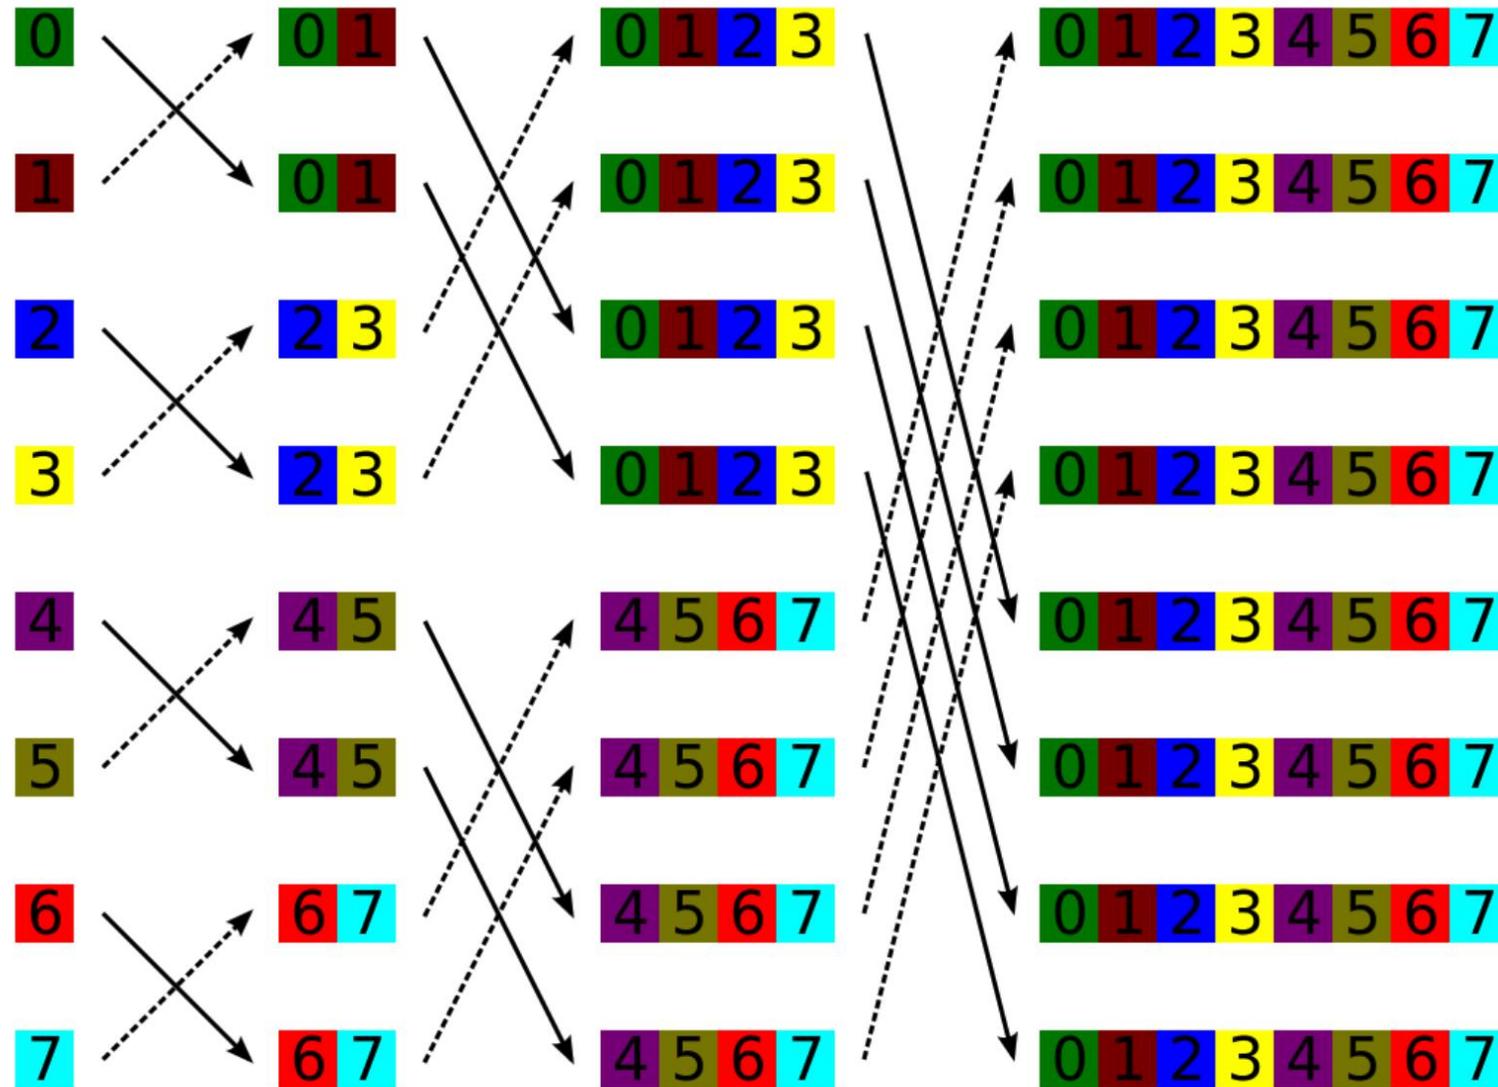
- Construct a tree of N workers;
- Step 1 (Aggregation): each worker sends M parameters to its parent; repeat  $\log(N)$  times
- Step 2 (Broadcast): each worker sends M parameters to its children; repeat  $\log(N)$  times



# Tree AllReduce

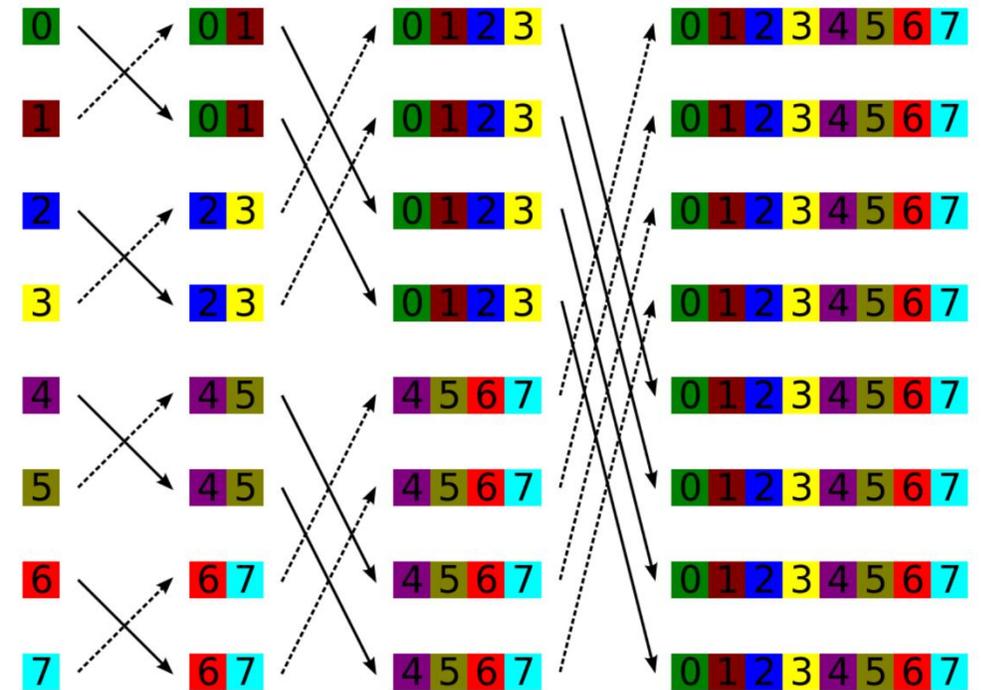
- Construct a tree of  $N$  workers;
- Step 1 (Aggregation): each worker sends  $M$  parameters to its parent; repeat  $\log(N)$  times
- Step 2 (Broadcast): each worker sends  $M$  parameters to its children; repeat  $\log(N)$  times
- Overall communication:  $2 * N * M$  parameters
  - Aggregation:  $M * N$  parameters
  - Broadcast:  $M * N$  parameters

# Butterfly Network



# Butterfly AllReduce

- Repeat  $\log(N)$  times:
  1. Each worker sends  $M$  parameters to its target node in the butterfly network
  2. Each worker aggregates gradients locally
- Overall communication:  $N * M * \log(N)$  parameters



# Comparing different AllReduce Methods

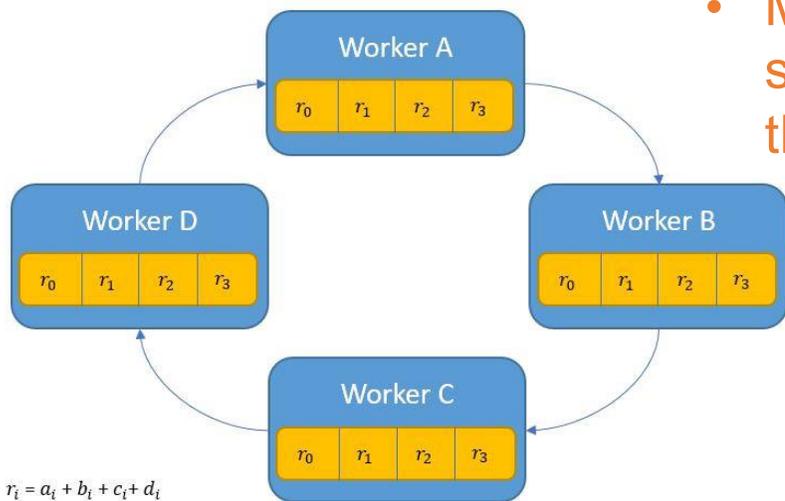
	Parameter Server	Naïve AllReduce	Ring AllReduce	Tree AllReduce	Butterfly AllReduce
<b>Overall communication</b>	$2 \times N \times M$	$N^2 \times M$	$2 \times N \times M$	$2 \times N \times M$	$N \times M \times \log N$

Question: Ring AllReduce is more efficient and scalable than Tree AllReduce and Parameter Server, why?

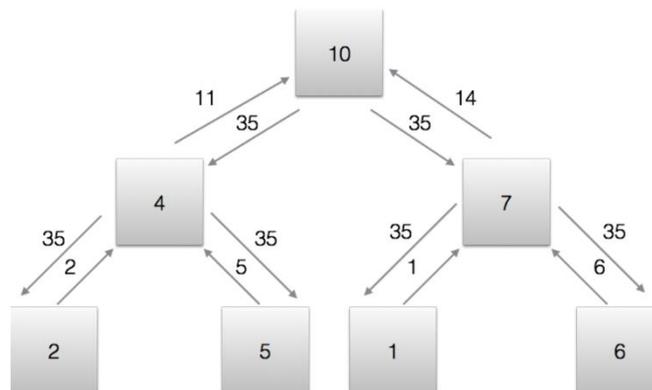
# Ring AllReduce v.s. Tree AllReduce v.s. Parameter Server

Ring AllReduce:

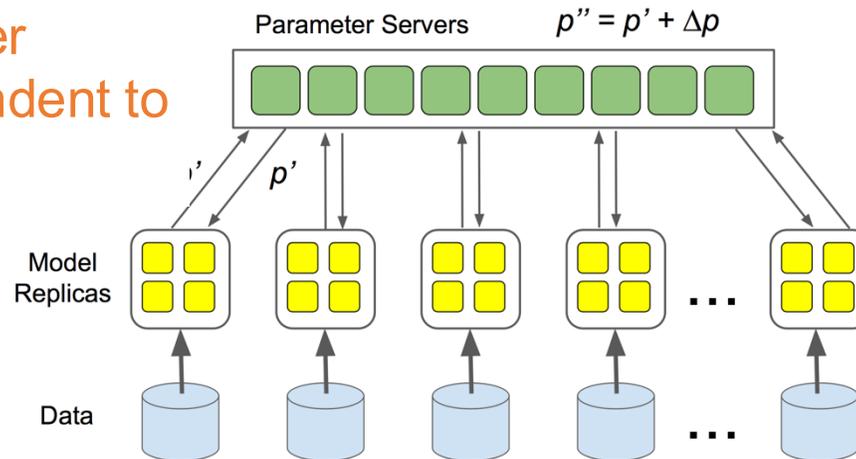
- Best latency
- Balanced workload across workers
- More scalable since each worker sends  $2 \cdot M$  parameters (independent to the number of workers)



Each worker sends  $M/N$  parameters per iteration; repeat for  $2 \cdot N$  iterations  
**Latency:**  $M/N * (2 \cdot N) / \text{bandwidth}$



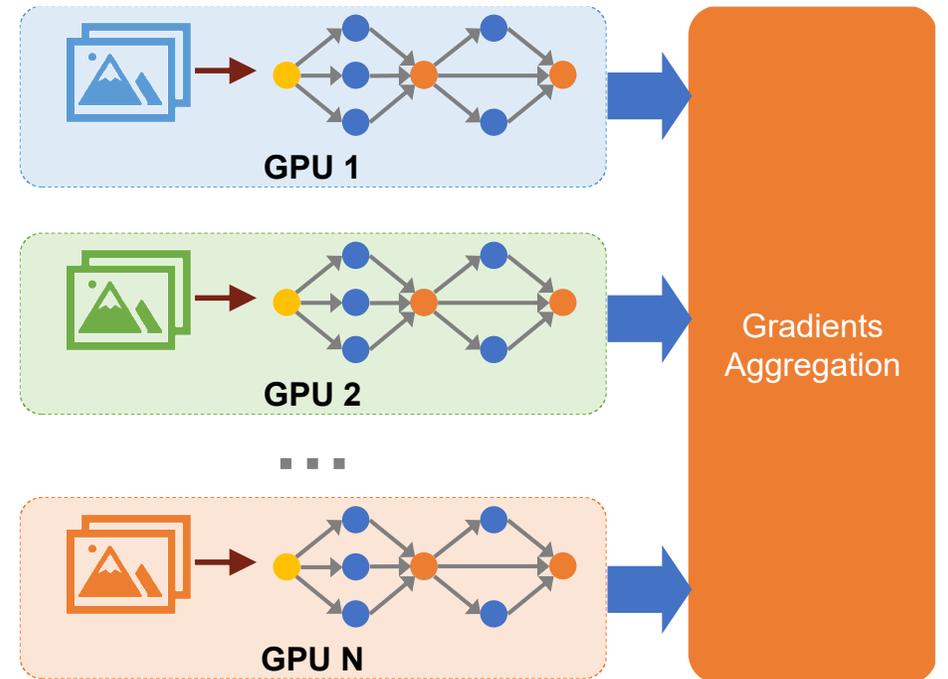
Each worker sends  $M$  parameters per iteration; repeat for  $2 \cdot \log(N)$  iterations  
**Latency:**  $M * 2 * \log(N) / \text{bandwidth}$



All workers send  $M$  parameters to parameter servers and receive  $M$  parameters from servers  
**Latency:**  $M * N / \text{bandwidth}$

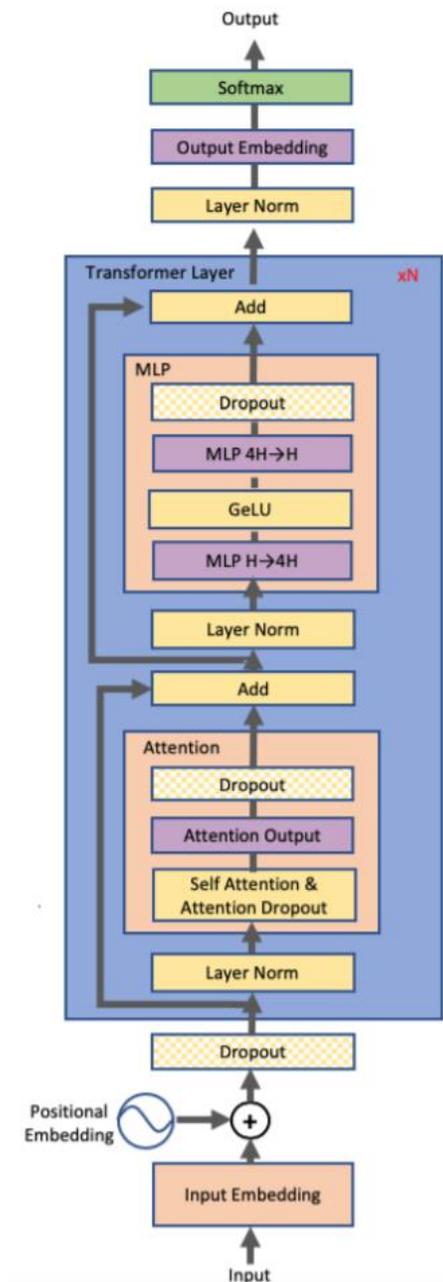
# An Issue with Data Parallelism

- Each GPU saves a replica of the entire model
- Cannot train large models that exceed GPU device memory



# Large Model Training Challenges

	<b>Bert-Large</b>	<b>GPT-2</b>	<b>Turing 17.2 NLG</b>	<b>GPT-3</b>
Parameters	0.32B	1.5B	17.2B	175B
Layers	24	48	78	96
Hidden Dimension	1024	1600	4256	12288
Relative Computation	1x	4.7x	<b>54x</b>	<b>547x</b>
Memory Footprint	5.12GB	24GB	<b>275GB</b>	<b>2800GB</b>

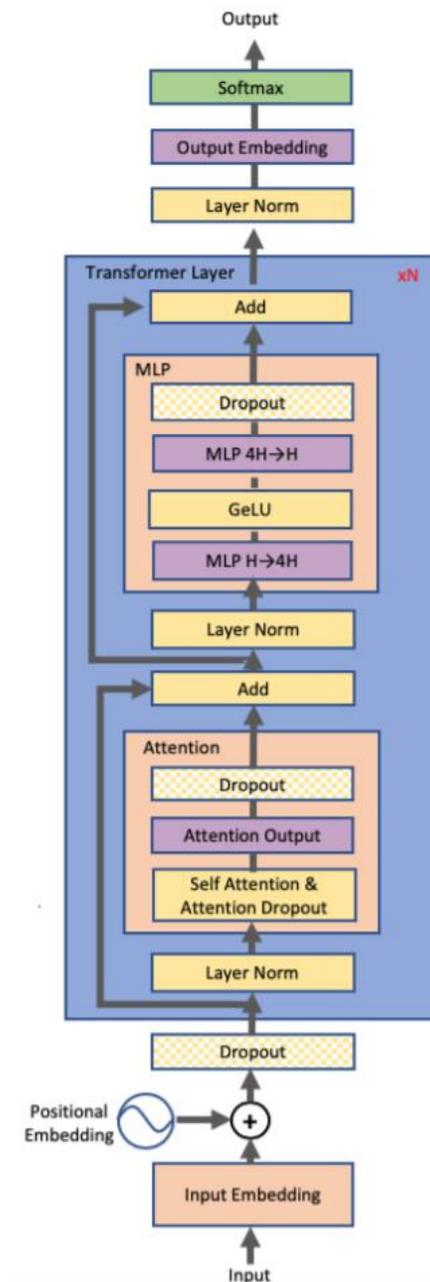


# Large Model Training Challenges

	<b>Bert-Large</b>	<b>GPT-2</b>	<b>Turing 17.2 NLG</b>	<b>GPT-3</b>
Parameters	0.32B	1.5B	17.2B	175B
Layers	24	48	78	96
Hidden Dimension	1024	1600	4256	12288
Relative Computation	1x	4.7x	<b>54x</b>	<b>547x</b>
Memory Footprint	5.12GB	24GB	<b>275GB</b>	<b>2800GB</b>

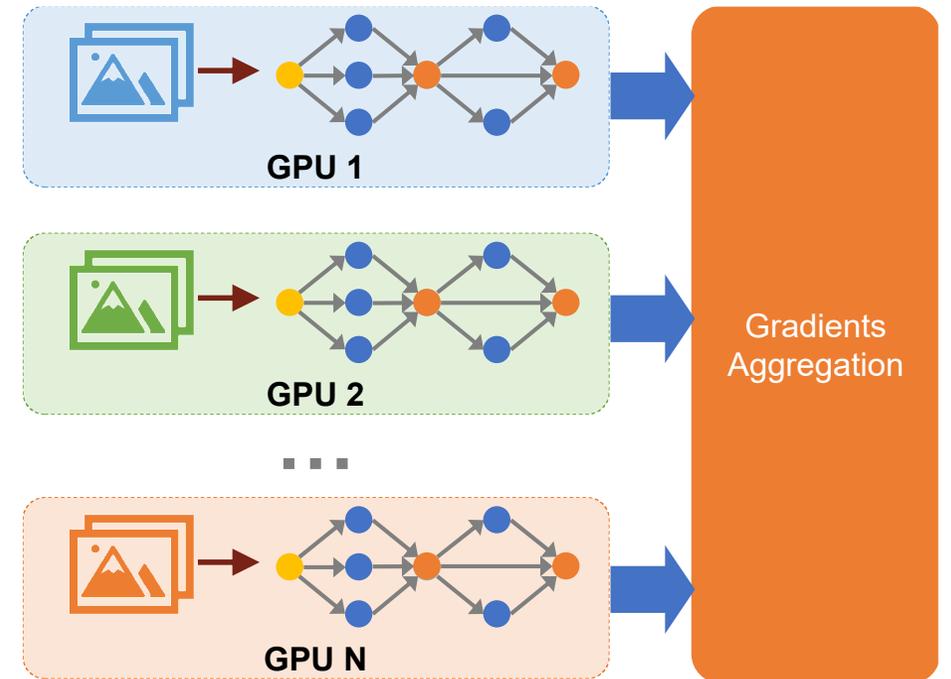
NVIDIA V100 GPU memory capacity: 16G/32G  
 NVIDIA A100 GPU memory capacity: 40G/80G

Out of Memory



# ZeRO: Zero Redundancy Optimizer

- Eliminating data redundancy in data parallel training
- A widely used technique for data parallel training of large models



# Revisit: Stochastic Gradient Descent

For t = 1 to T

Backward pass      Forward pass

$\Delta w = \eta \times \frac{1}{b} \sum_{i=1}^b \nabla \left( \text{loss}(f_w(x_i, y_i)) \right)$  // compute derivative and update

w -= Δw // apply update

End

# Adaptive Learning Rates (Adam)

For  $t = 1$  to  $T$

$$g = \frac{1}{b} \sum_{i=1}^b \nabla \left( \text{loss}(f_w(x_i, y_i)) \right)$$

$$\Delta w = \text{adam}(g)$$

$w \ -= \ \Delta w$  // apply update

End

$$\begin{aligned} \nu_t &= \beta_1 * \nu_{t-1} - (1 - \beta_1) * g_t \\ s_t &= \beta_2 * s_{t-1} - (1 - \beta_2) * g_t^2 \end{aligned}$$

$$\Delta \omega_t = -\eta \frac{\nu_t}{\sqrt{s_t + \epsilon}} * g_t$$

$g_t$  : Gradient at time  $t$  along  $\omega^j$

$\nu_t$  : Exponential Average of gradients along  $\omega_j$

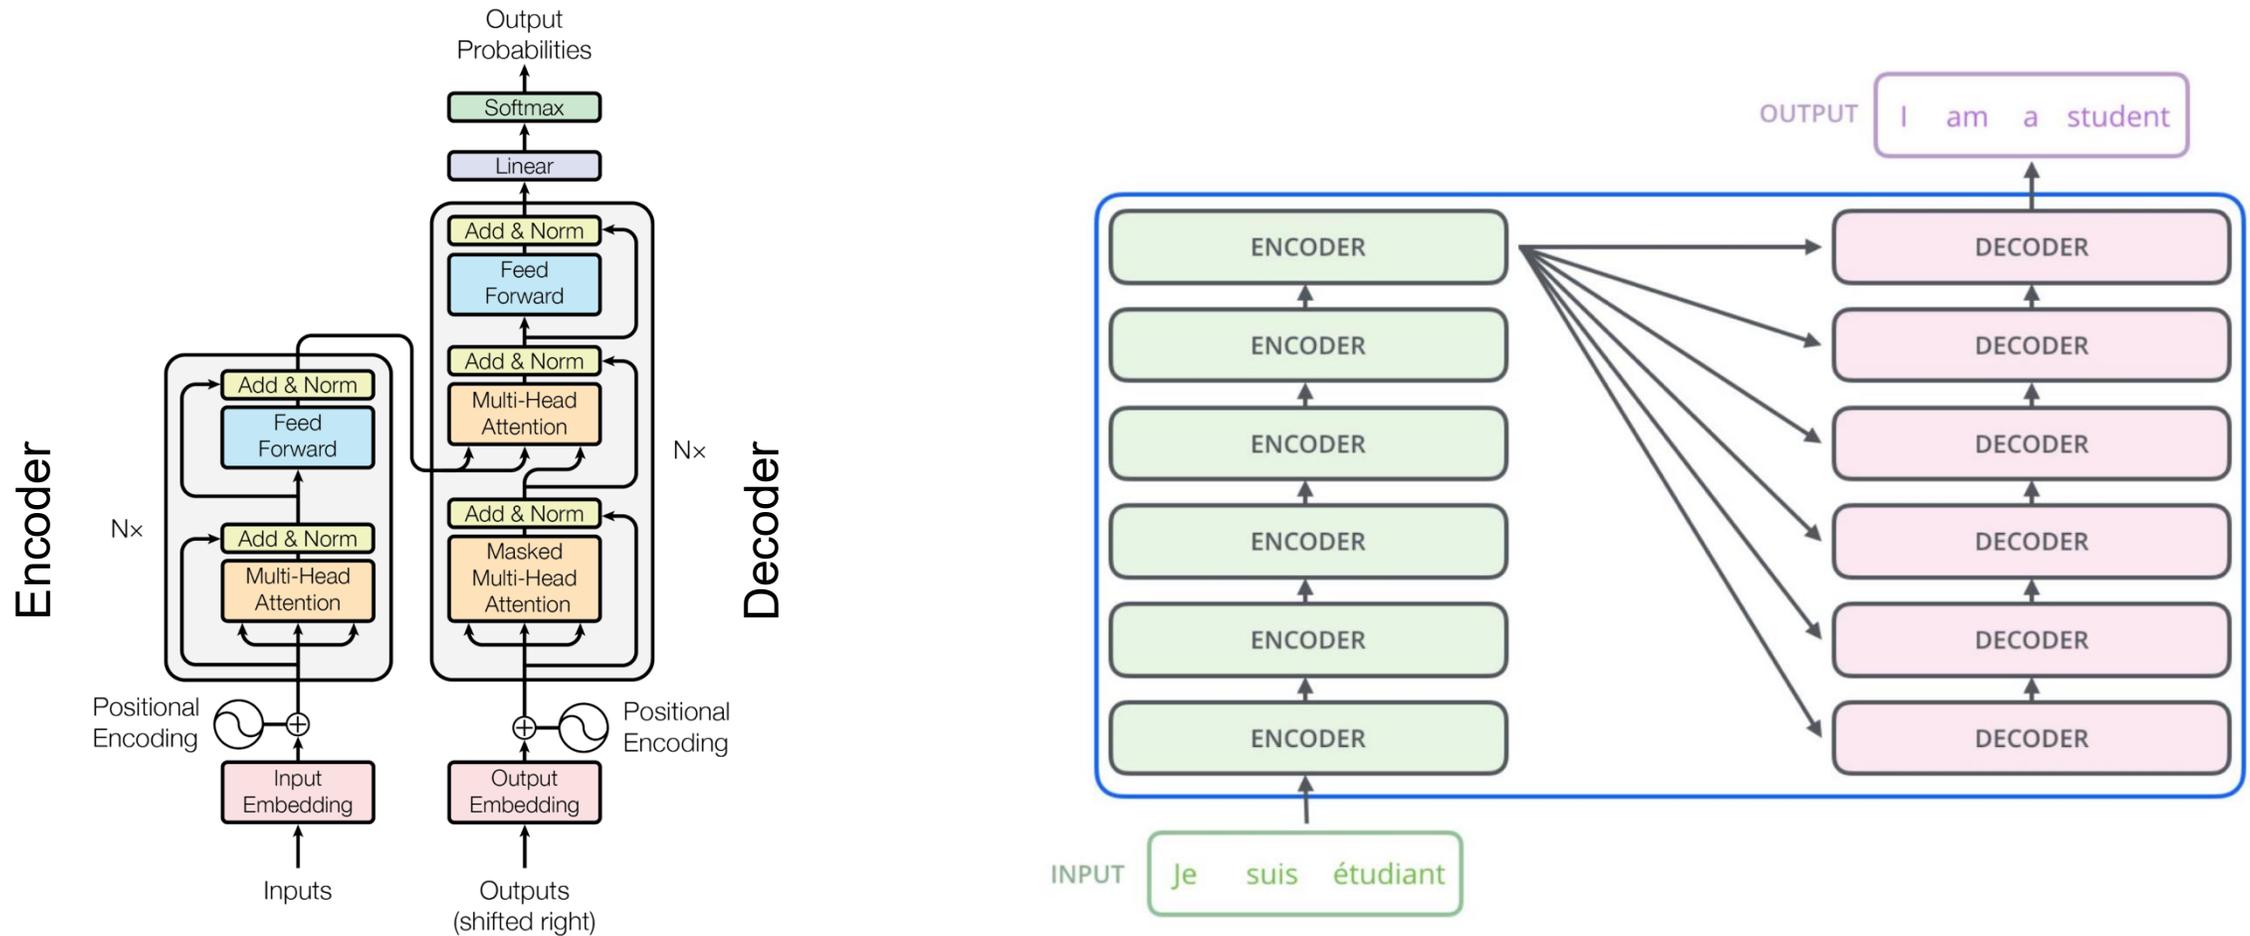
$s_t$  : Exponential Average of squares of gradients along  $\omega_j$

$\beta_1, \beta_2$  : Hyperparameters

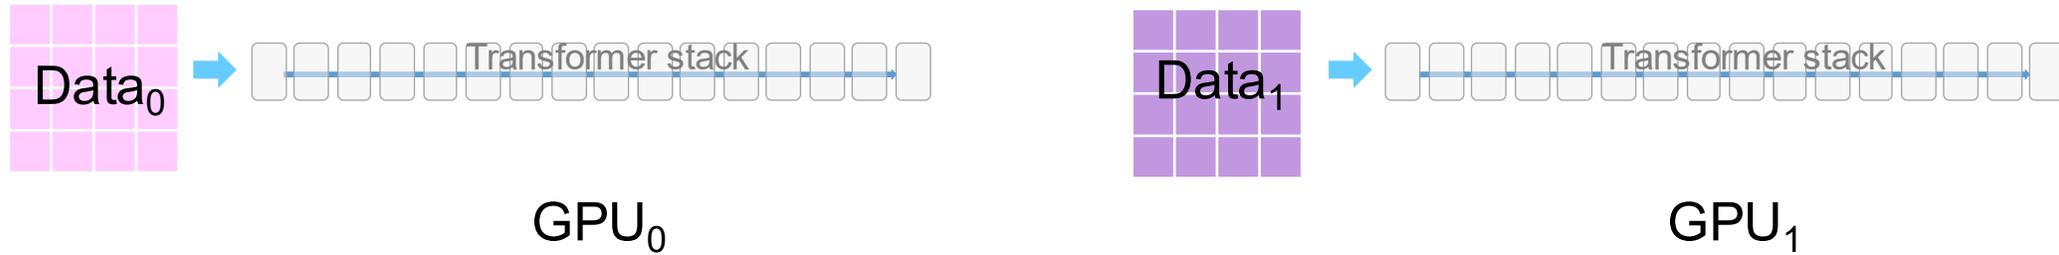
[1] Kingma and Ba, "Adam: A Method for Stochastic Optimization", 2014,

<https://arxiv.org/abs/1412.6980>

# Transformer for Language Models

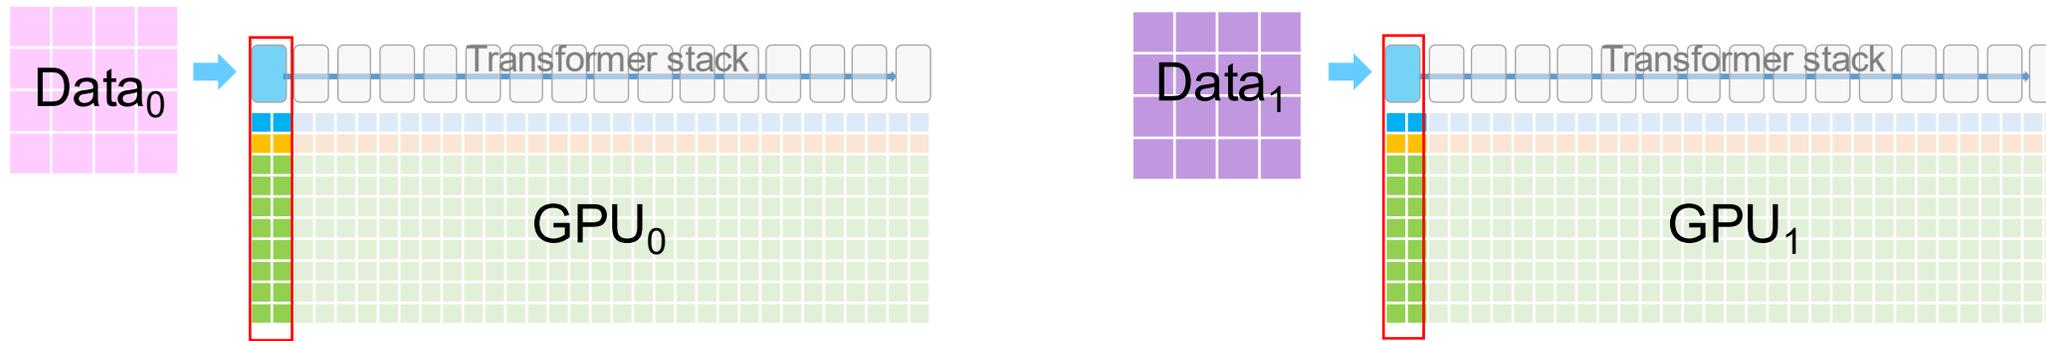


# Understanding Memory Consumption



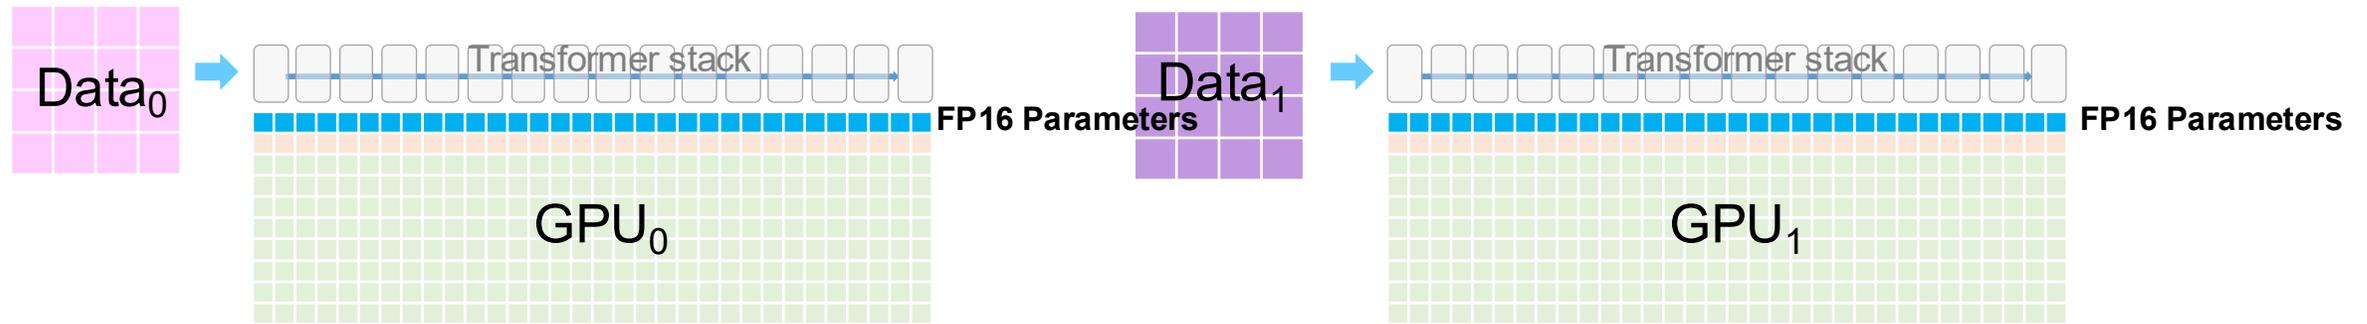
A 16-layer transformer model  = 1 layer

# Understanding Memory Consumption



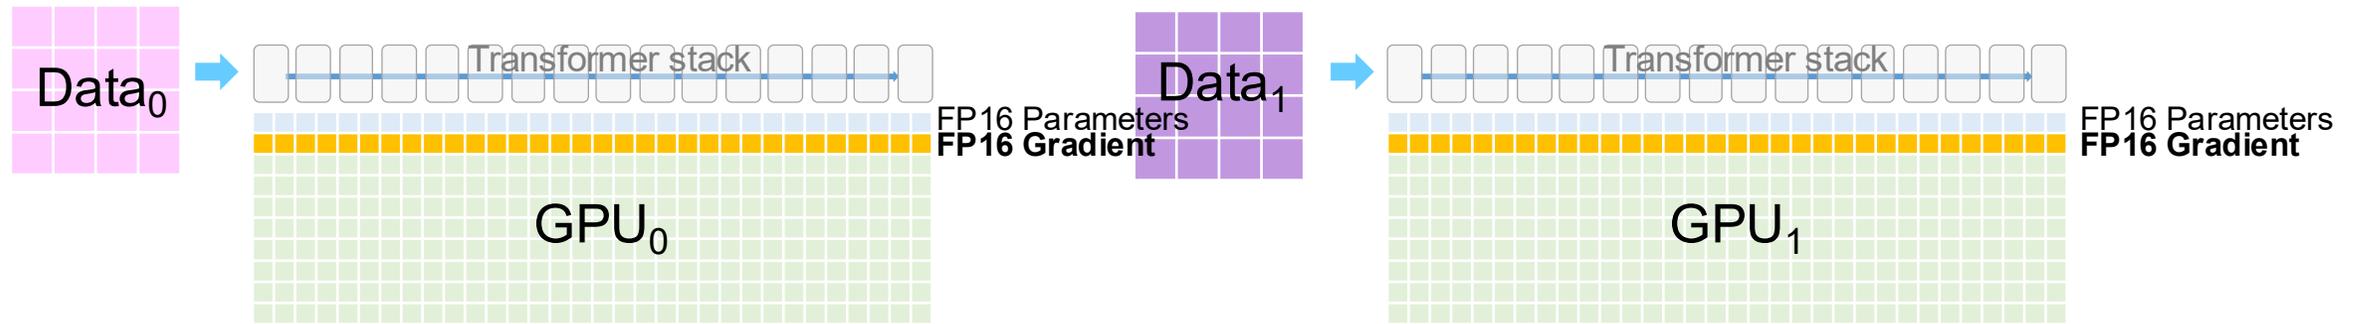
Each cell  represents GPU memory used by its corresponding transformer layer 

# Understanding Memory Consumption



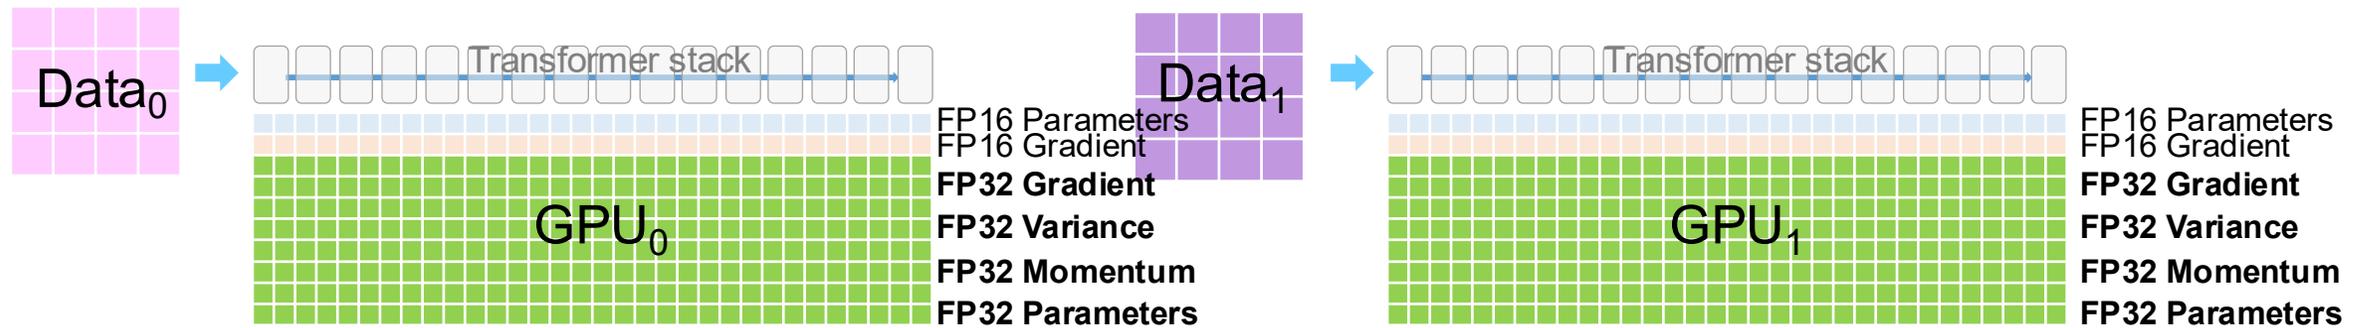
- FP16 parameter

# Understanding Memory Consumption



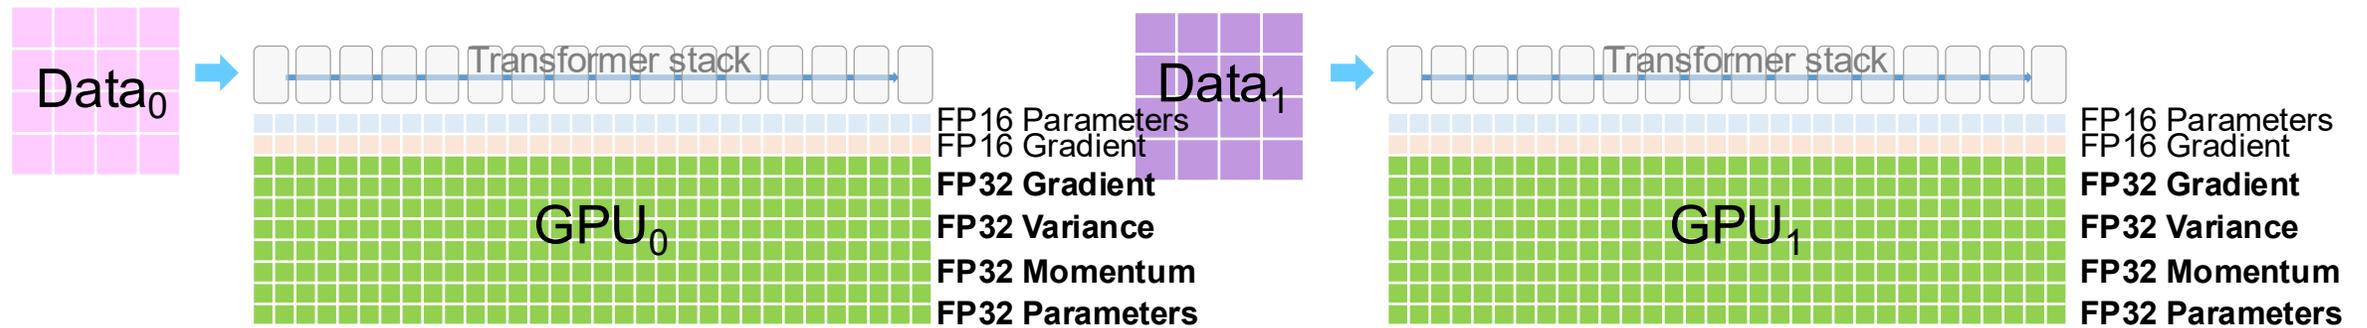
- FP16 parameter
- FP16 Gradients

# Understanding Memory Consumption



- FP16 parameter
- FP16 Gradients
- FP32 Optimizer States
  - Gradients, Variance, Momentum, Parameters

# Understanding Memory Consumption



- FP16 parameter : **2M bytes**
- FP16 Gradients : **2M bytes**
- FP32 Optimizer States : **16M bytes**
  - Gradients, Variance, Momentum, Parameters

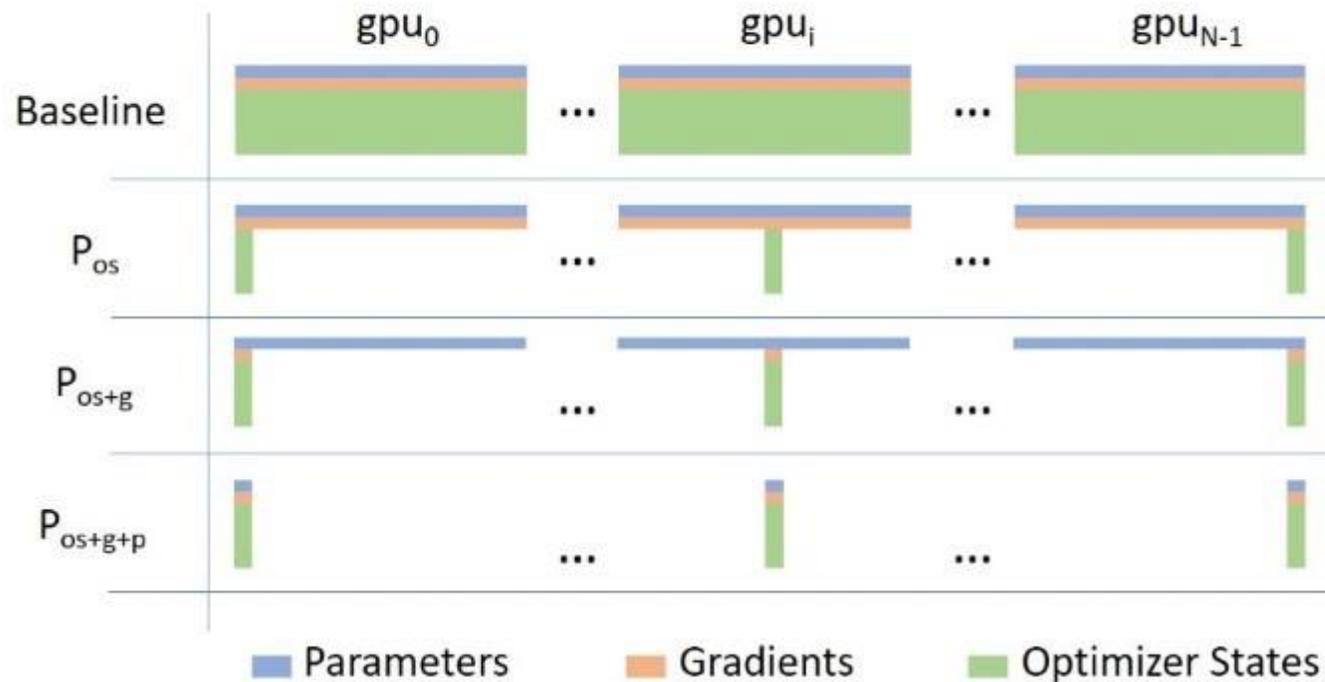
Example 1B parameter model ->  
20GB/GPU

- Memory consumption doesn't include:
- Input batch + activations

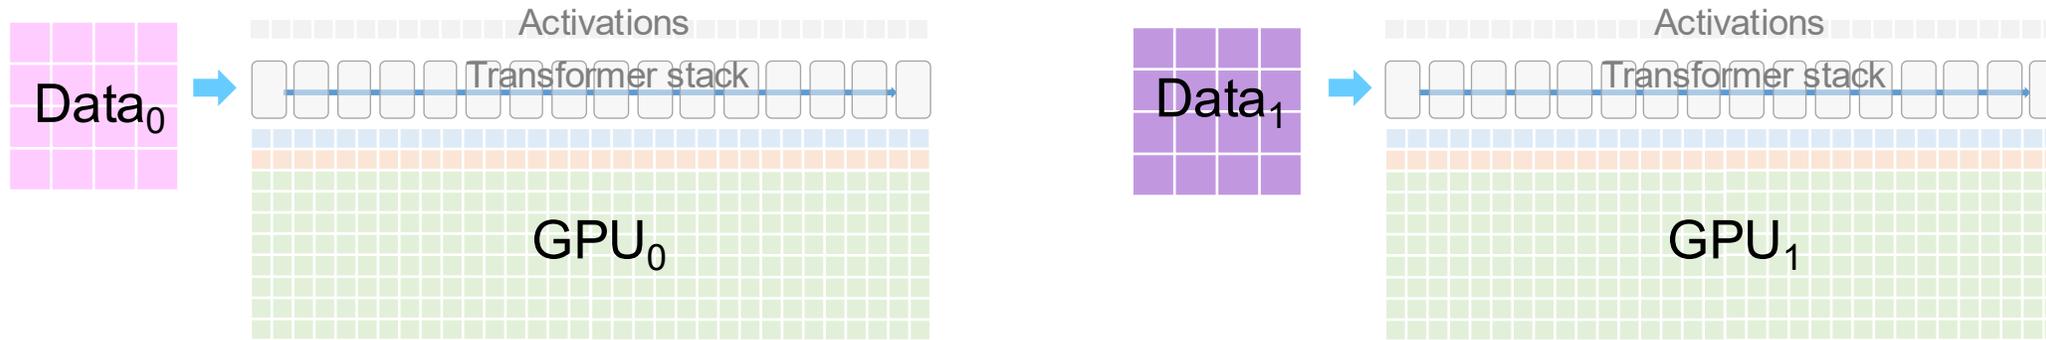
M = number of parameters in the model

# ZeRO-DP: ZeRO powered Data Parallelism

- ZeRO removes the redundancy across data parallel process
- Stage 1: partitioning optimizer states
- Stage 2: partitioning gradients
- Stage 3: partitioning parameters

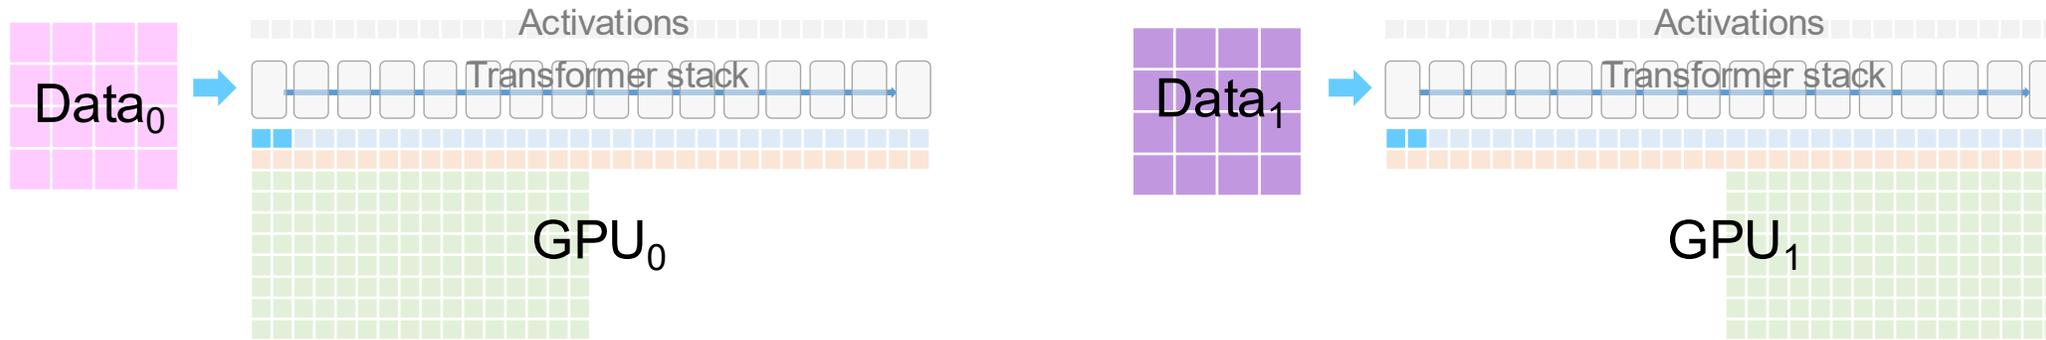


# ZeRO Stage 1: Partitioning Optimizer States



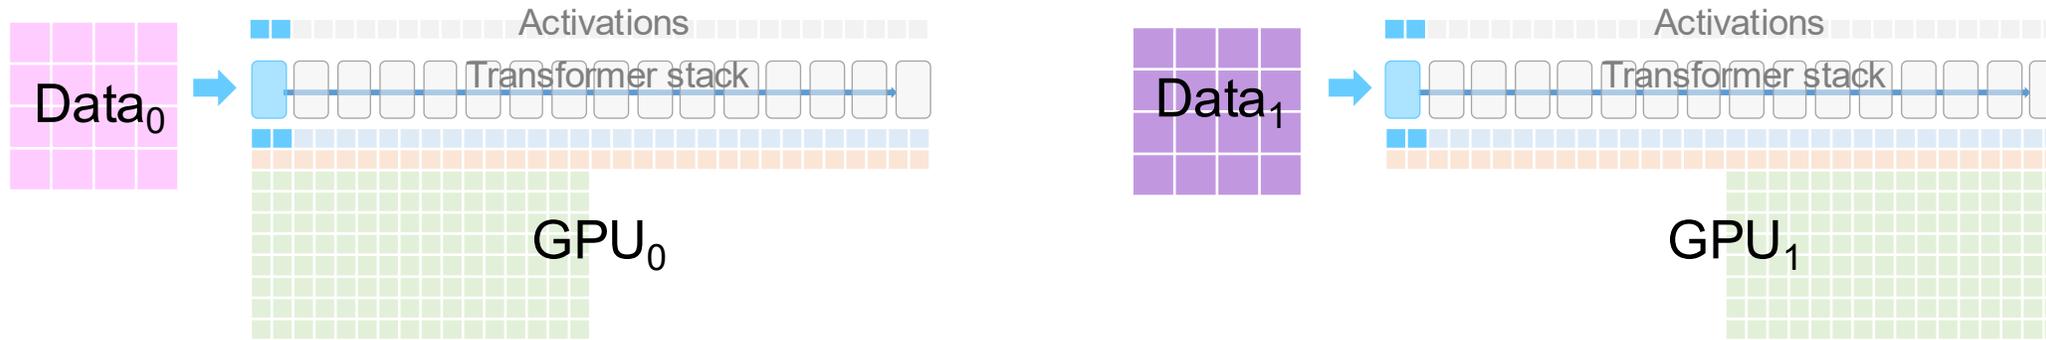
- ZeRO Stage 1

# ZeRO Stage 1: Partitioning Optimizer States



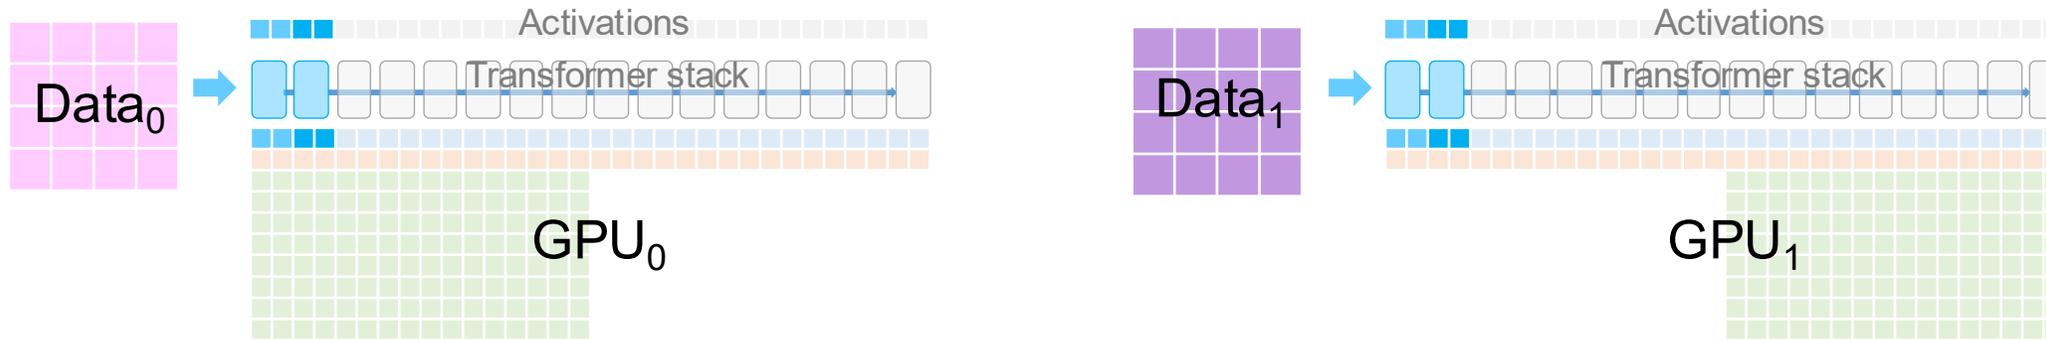
- ZeRO Stage 1
- Partitions optimizer states across GPUs

# ZeRO Stage 1: Partitioning Optimizer States



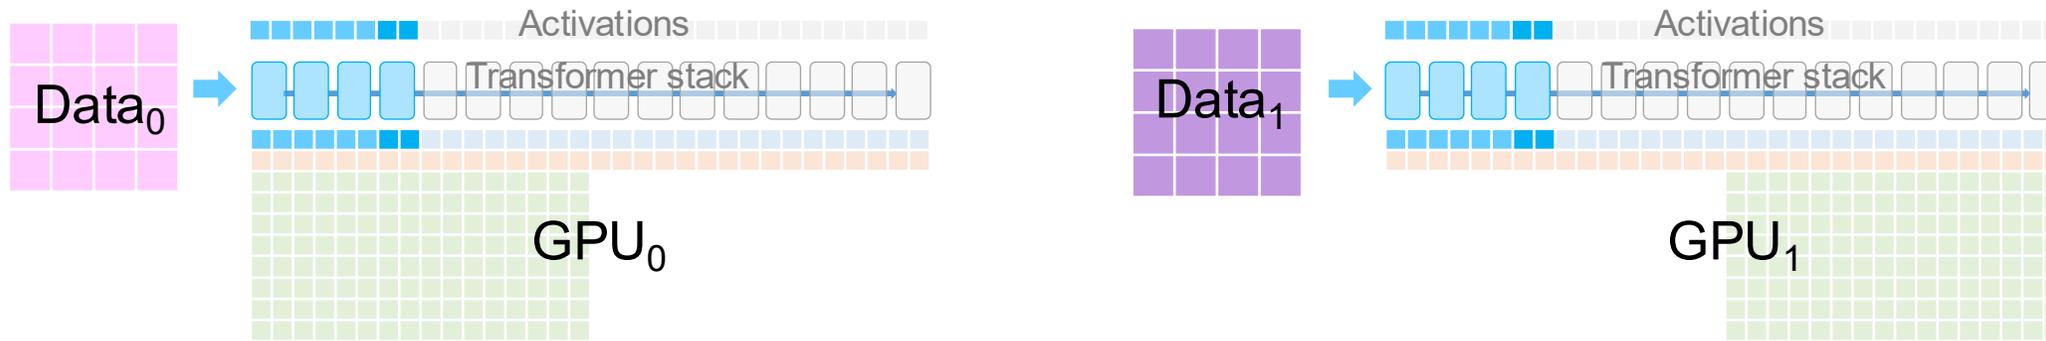
- ZeRO Stage 1
- Partitions optimizer states across GPUs
- Run Forward across the transformer blocks

# ZeRO Stage 1: Partitioning Optimizer States



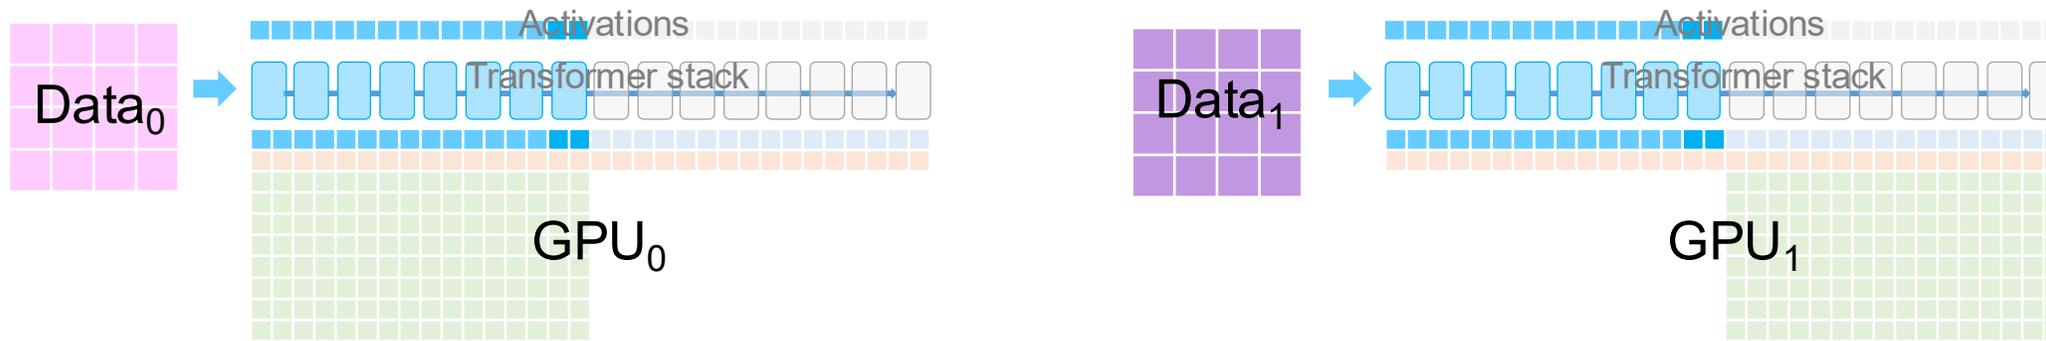
- ZeRO Stage 1
- Partitions optimizer states across GPUs
- Run Forward across the transformer blocks

# ZeRO Stage 1: Partitioning Optimizer States



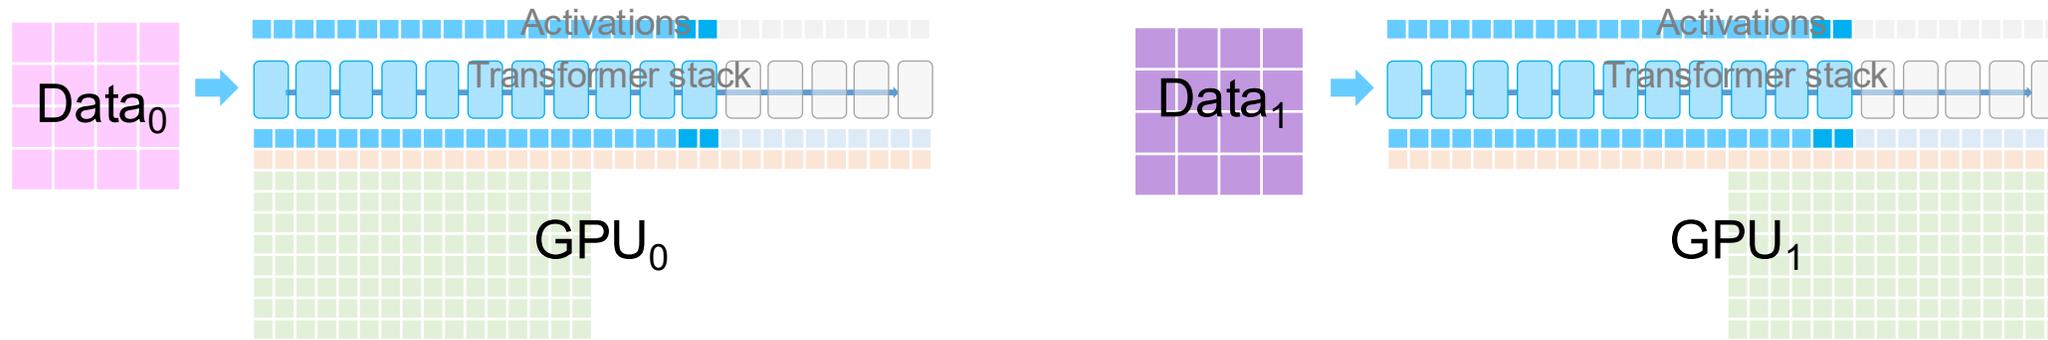
- ZeRO Stage 1
- Partitions optimizer states across GPUs
- Run Forward across the transformer blocks

# ZeRO Stage 1: Partitioning Optimizer States



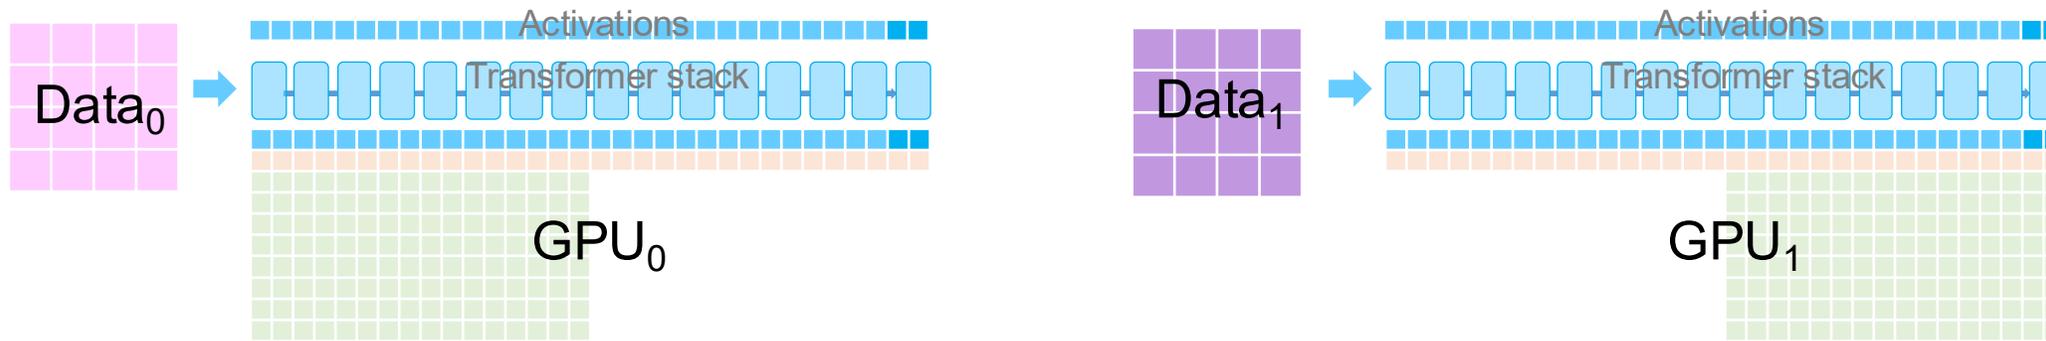
- ZeRO Stage 1
- Partitions optimizer states across GPUs
- Run Forward across the transformer blocks

# ZeRO Stage 1: Partitioning Optimizer States



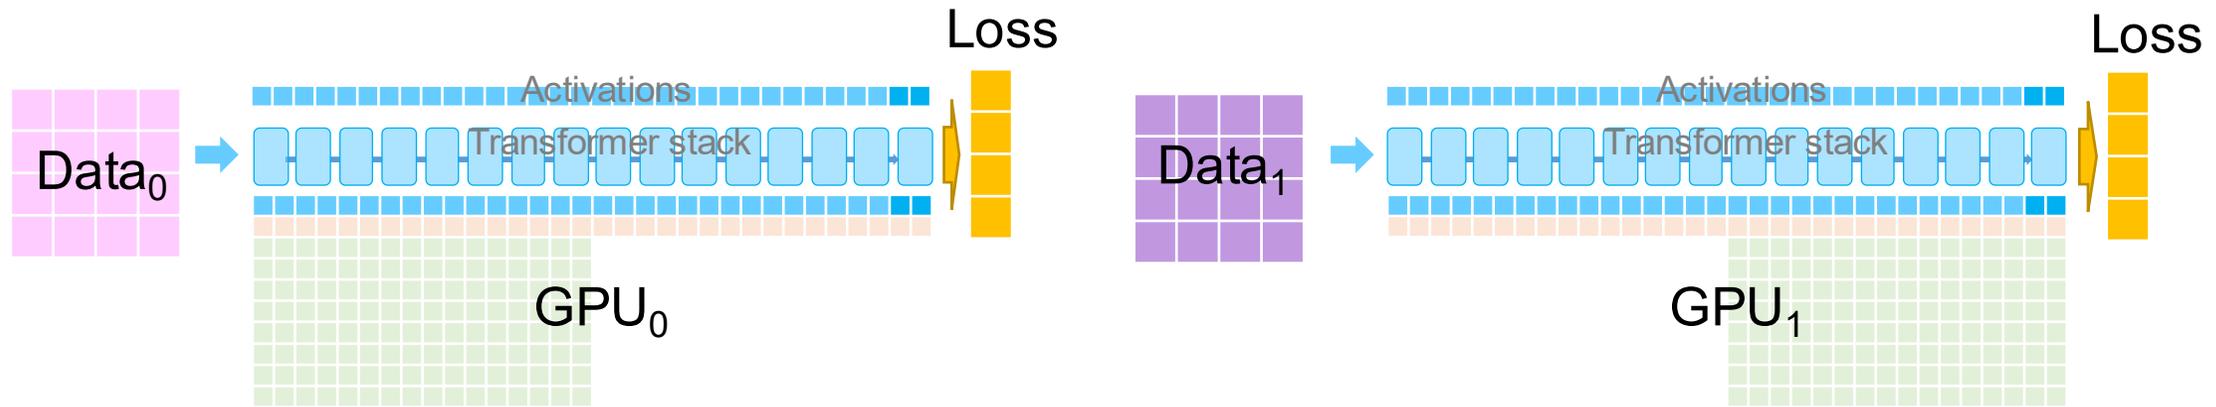
- ZeRO Stage 1
- Partitions optimizer states across GPUs
- Run Forward across the transformer blocks

# ZeRO Stage 1: Partitioning Optimizer States



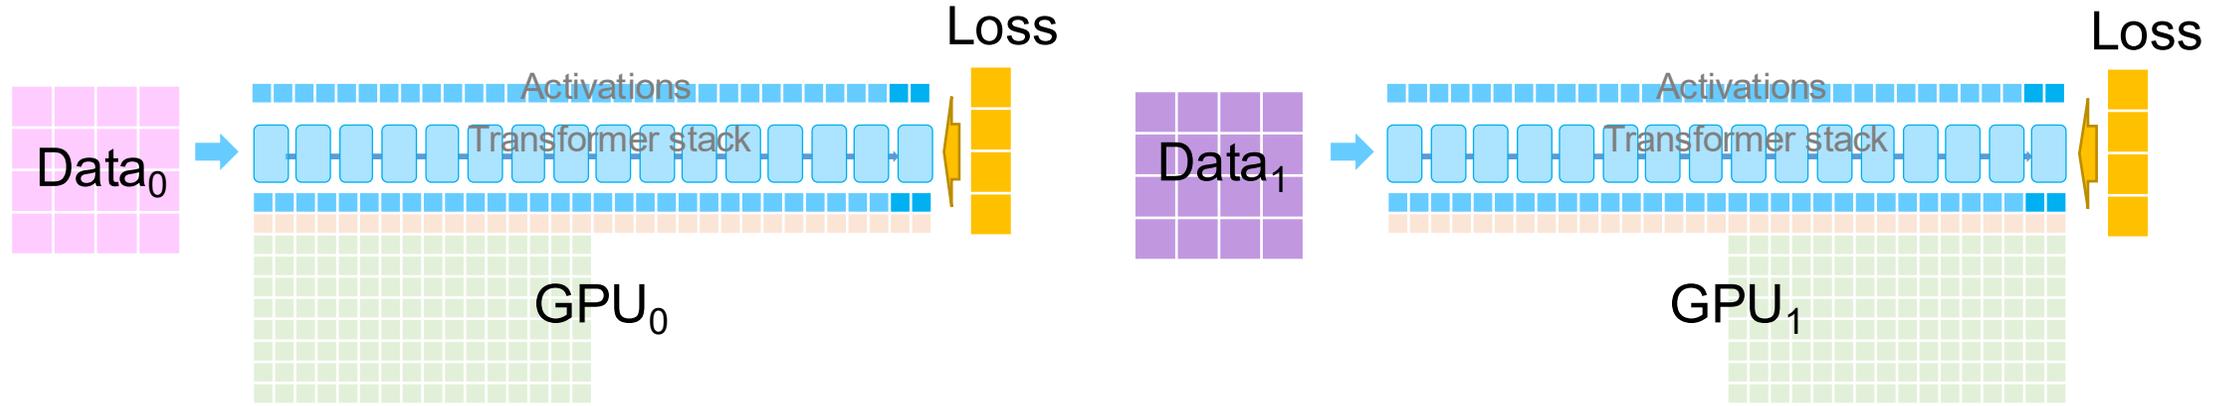
- ZeRO Stage 1
- Partitions optimizer states across GPUs
- Run Forward across the transformer blocks

# ZeRO Stage 1: Partitioning Optimizer States



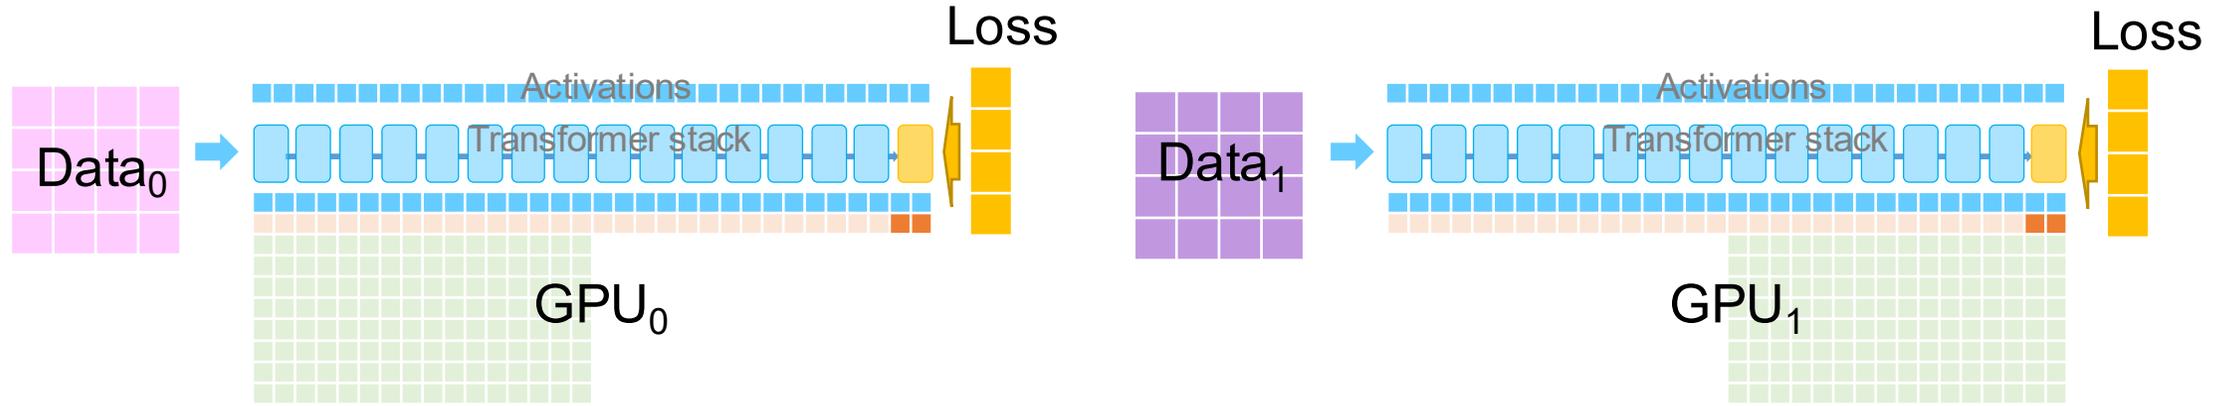
- ZeRO Stage 1
- Partitions optimizer states across GPUs
- Run Forward across the transformer blocks

# ZeRO Stage 1: Partitioning Optimizer States



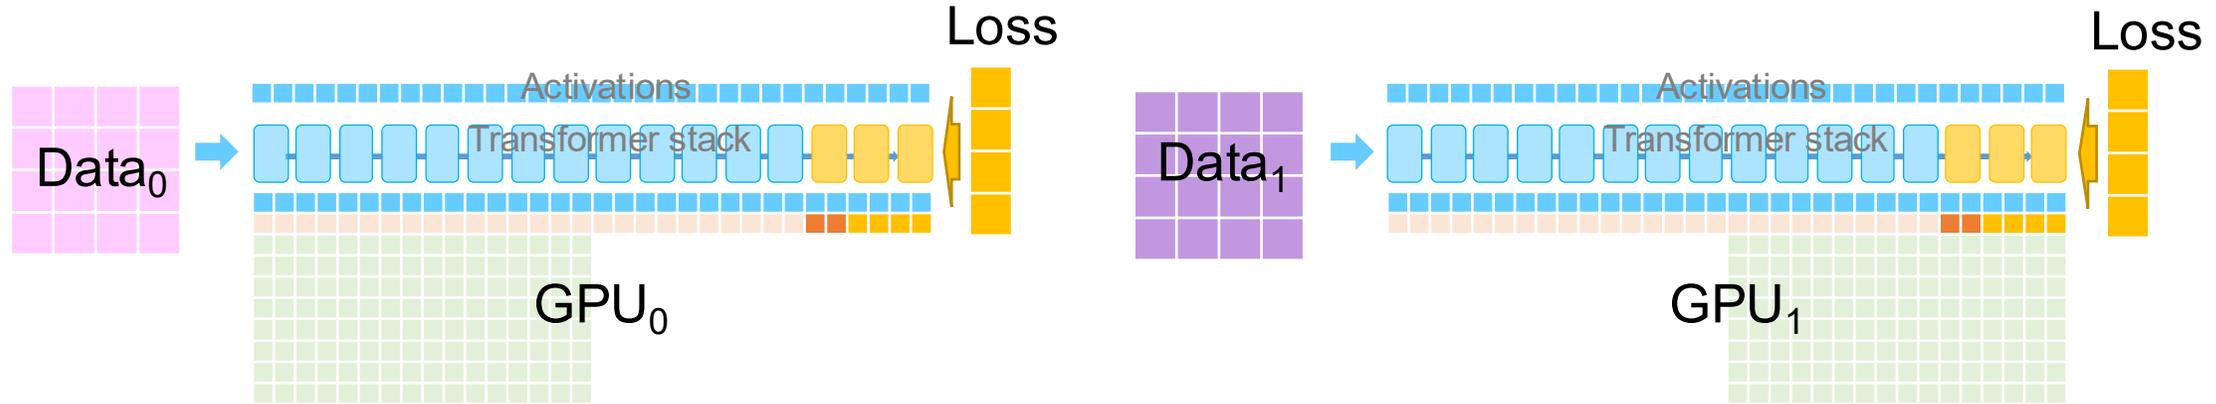
- ZeRO Stage 1
- Partitions optimizer states across GPUs
- Run Forward across the transformer blocks
- Backward propagation to generate FP16 gradients

# ZeRO Stage 1: Partitioning Optimizer States



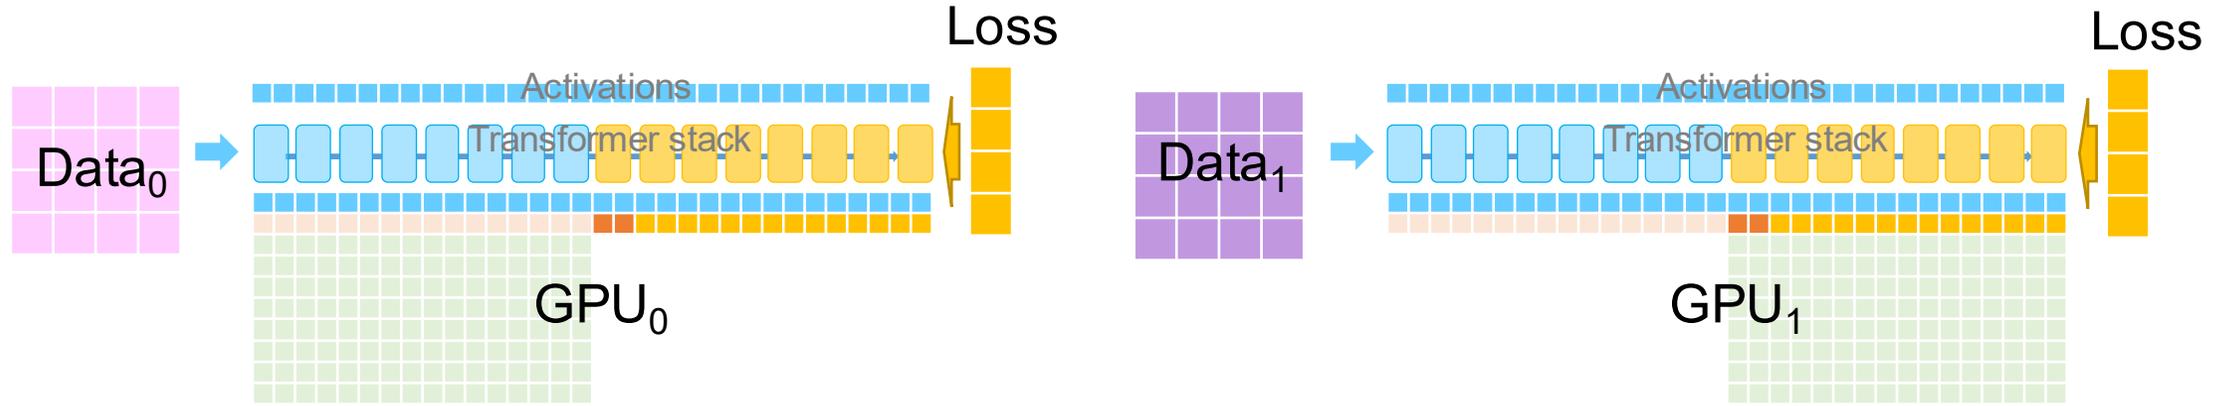
- ZeRO Stage 1
- Partitions optimizer states across GPUs
- Run Forward across the transformer blocks
- Backward propagation to generate FP16 gradients

# ZeRO Stage 1: Partitioning Optimizer States



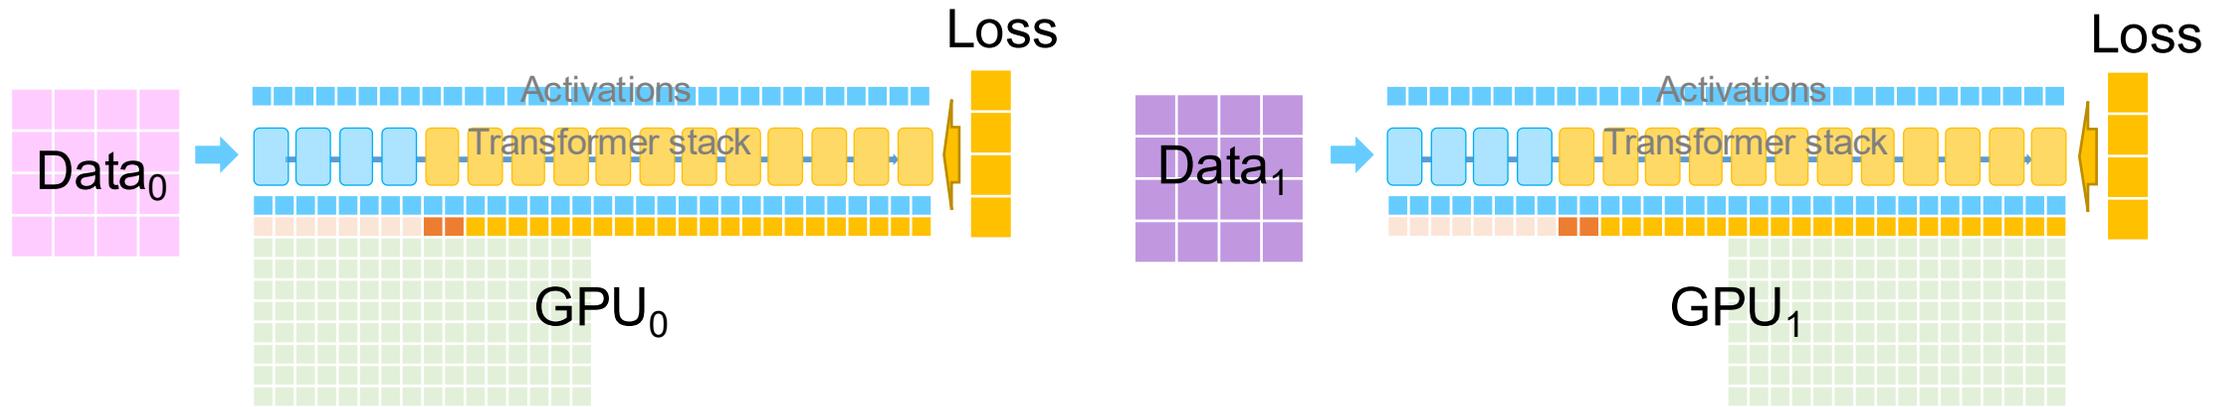
- ZeRO Stage 1
- Partitions optimizer states across GPUs
- Run Forward across the transformer blocks
- Backward propagation to generate FP16 gradients

# ZeRO Stage 1: Partitioning Optimizer States



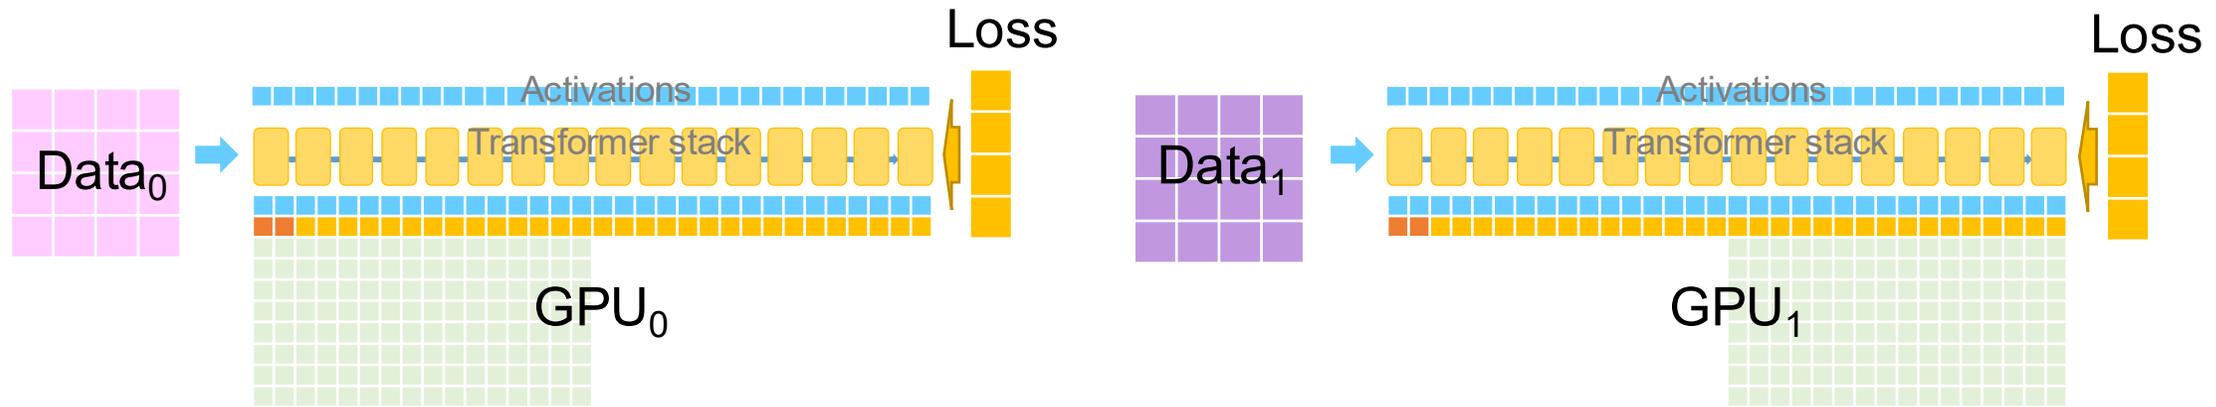
- ZeRO Stage 1
- Partitions optimizer states across GPUs
- Run Forward across the transformer blocks
- Backward propagation to generate FP16 gradients

# ZeRO Stage 1: Partitioning Optimizer States



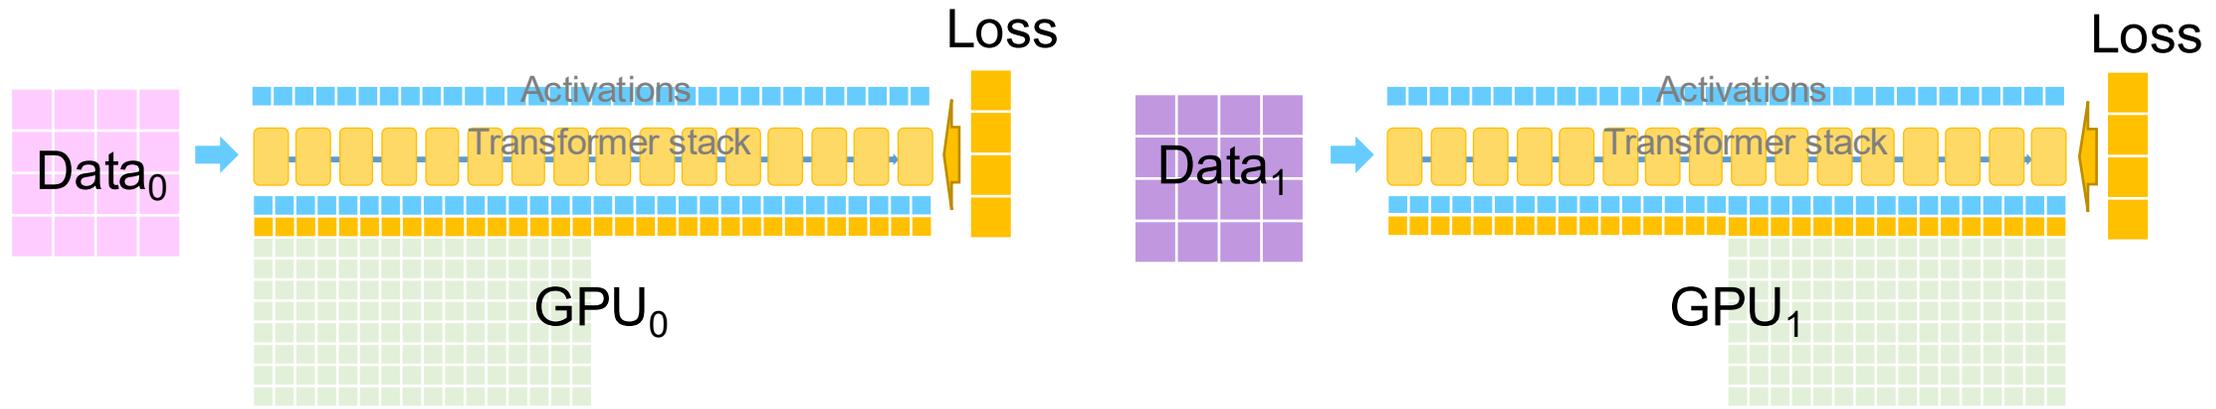
- ZeRO Stage 1
- Partitions optimizer states across GPUs
- Run Forward across the transformer blocks
- Backward propagation to generate FP16 gradients

# ZeRO Stage 1: Partitioning Optimizer States



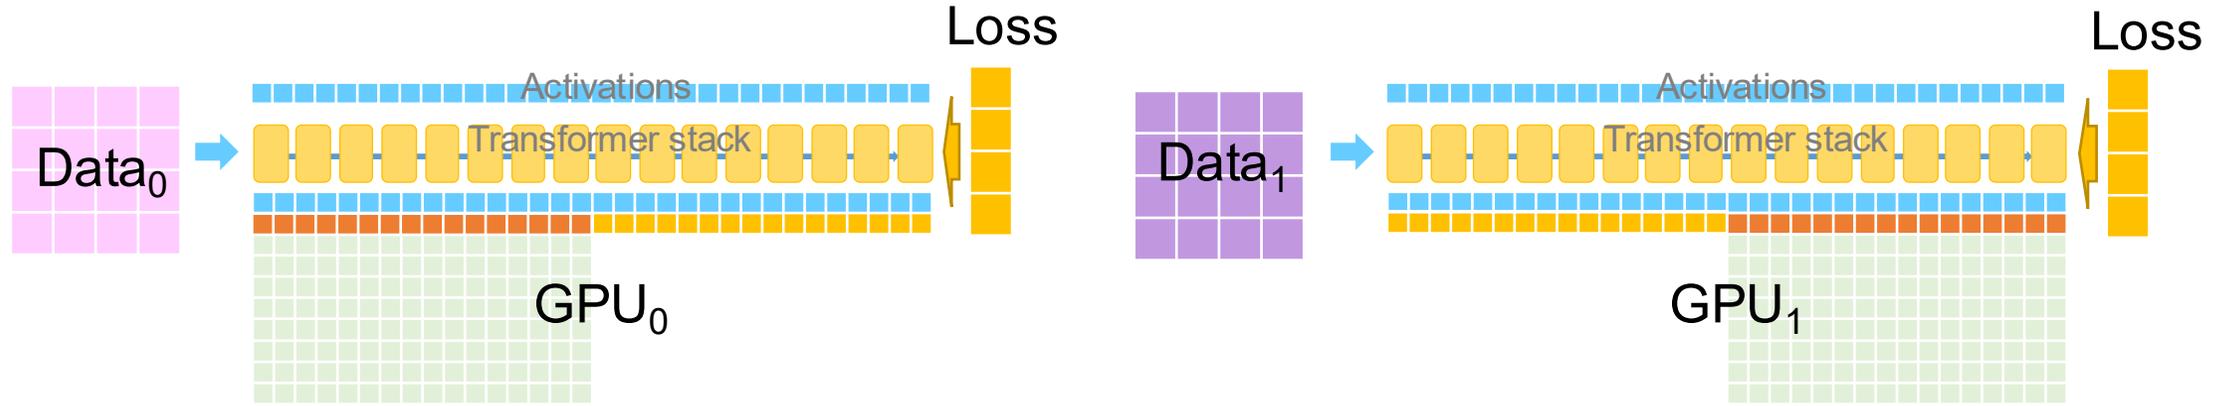
- ZeRO Stage 1
- Partitions optimizer states across GPUs
- Run Forward across the transformer blocks
- Backward propagation to generate FP16 gradients

# ZeRO Stage 1: Partitioning Optimizer States



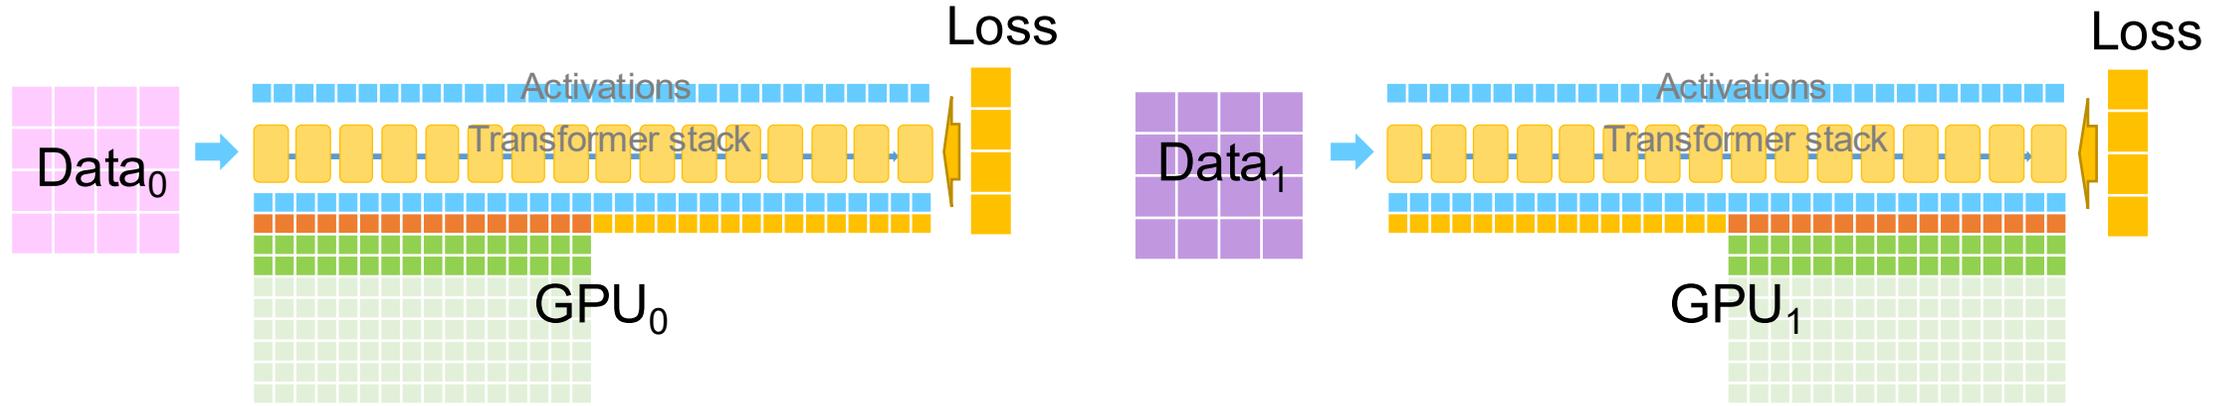
- ZeRO Stage 1
- Partitions optimizer states across GPUs
- Run Forward across the transformer blocks
- Backward propagation to generate FP16 gradients and AllReduce to average

# ZeRO Stage 1: Partitioning Optimizer States



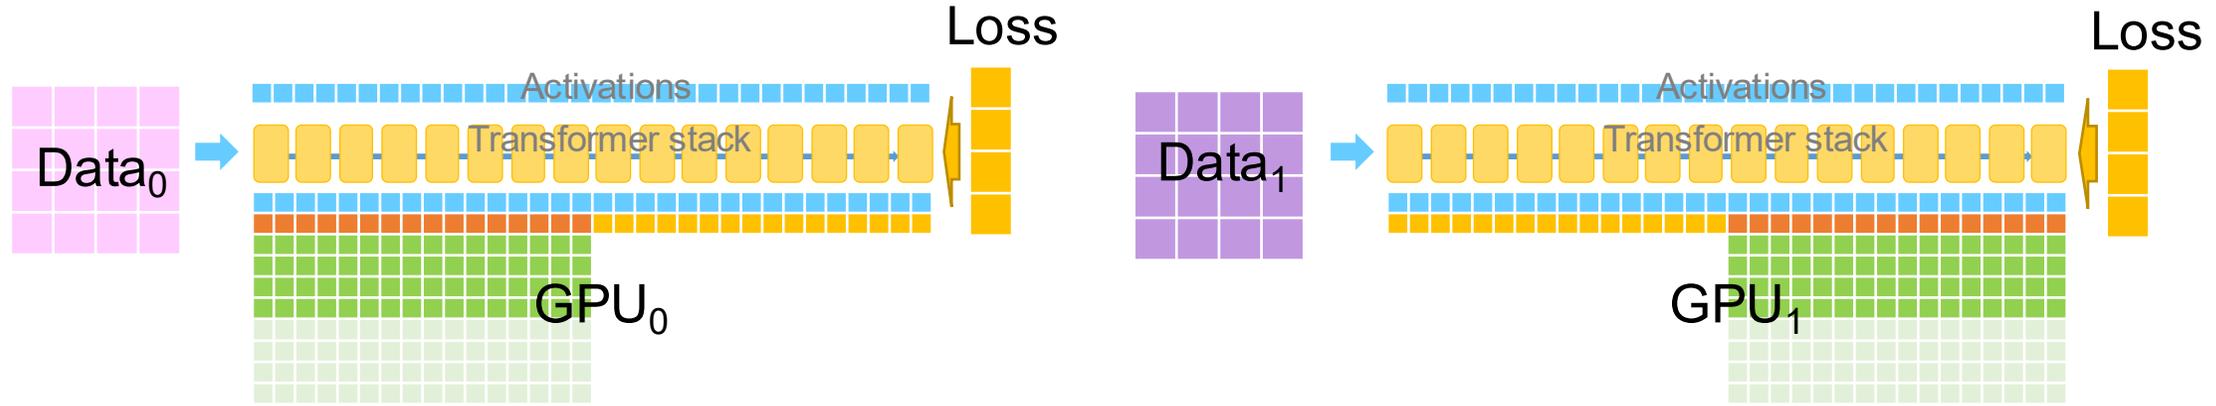
- ZeRO Stage 1
- Partitions optimizer states across GPUs
- Run Forward across the transformer blocks
- Backward propagation to generate FP16 gradients and AllReduce to average

# ZeRO Stage 1: Partitioning Optimizer States



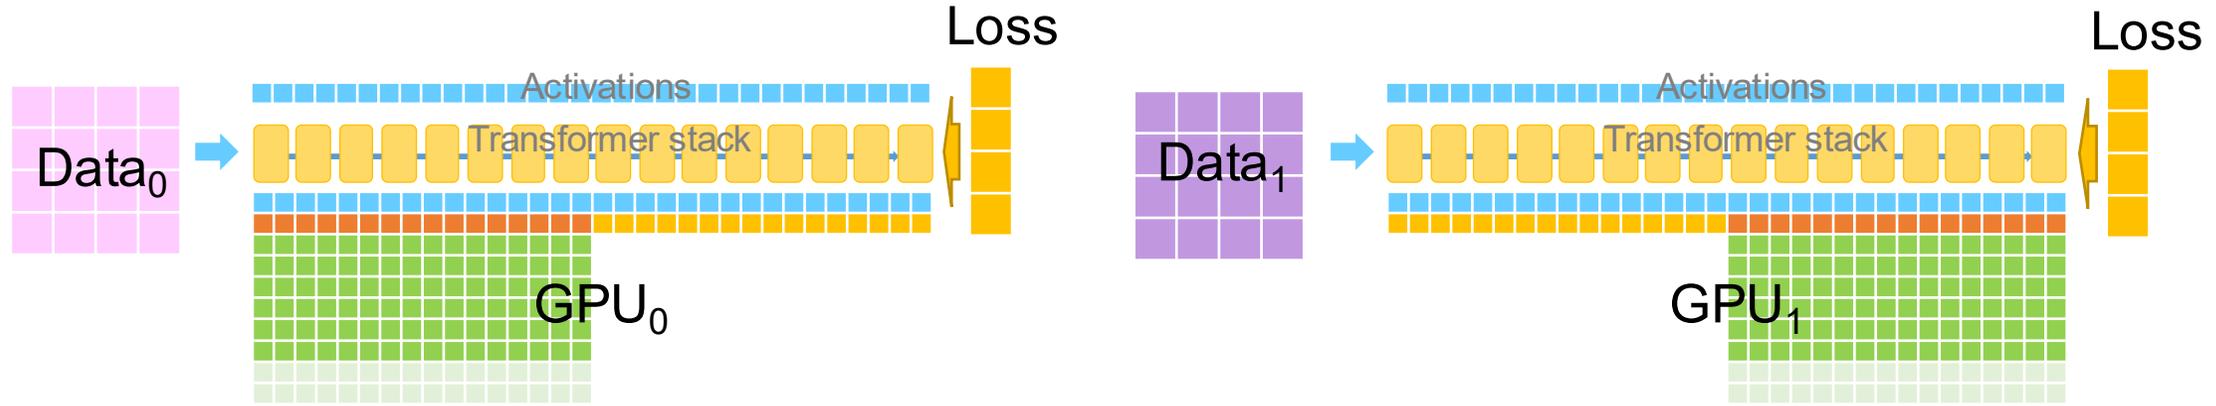
- ZeRO Stage 1
- Partitions optimizer states across GPUs
- Run Forward across the transformer blocks
- Backward propagation to generate FP16 gradients and AllReduce to average
- Update the FP32 weights with ADAM optimizer

# ZeRO Stage 1: Partitioning Optimizer States



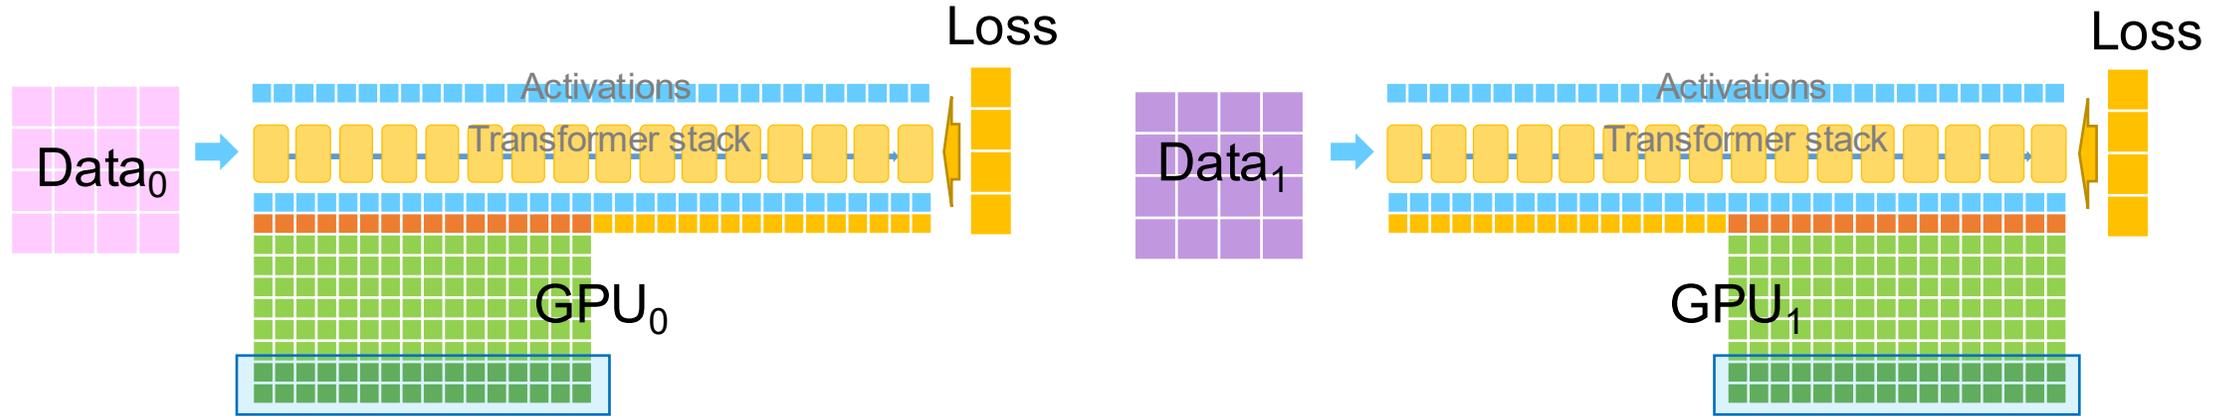
- Run Forward across the transformer blocks
- Backward propagation to generate FP16 gradients and AllReduce to average
- Update the FP32 weights with ADAM optimizer

# ZeRO Stage 1: Partitioning Optimizer States



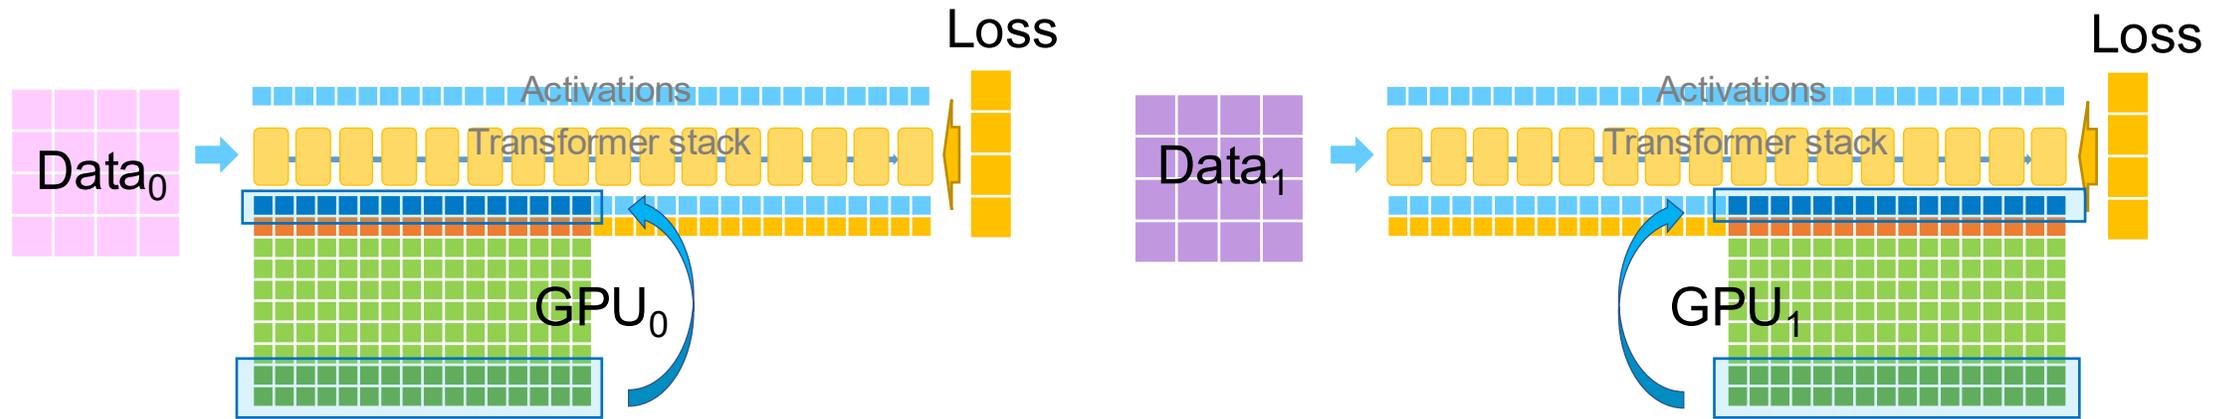
- Run Forward across the transformer blocks
- Backward propagation to generate FP16 gradients and AllReduce to average
- Update the FP32 weights with ADAM optimizer

# ZeRO Stage 1: Partitioning Optimizer States



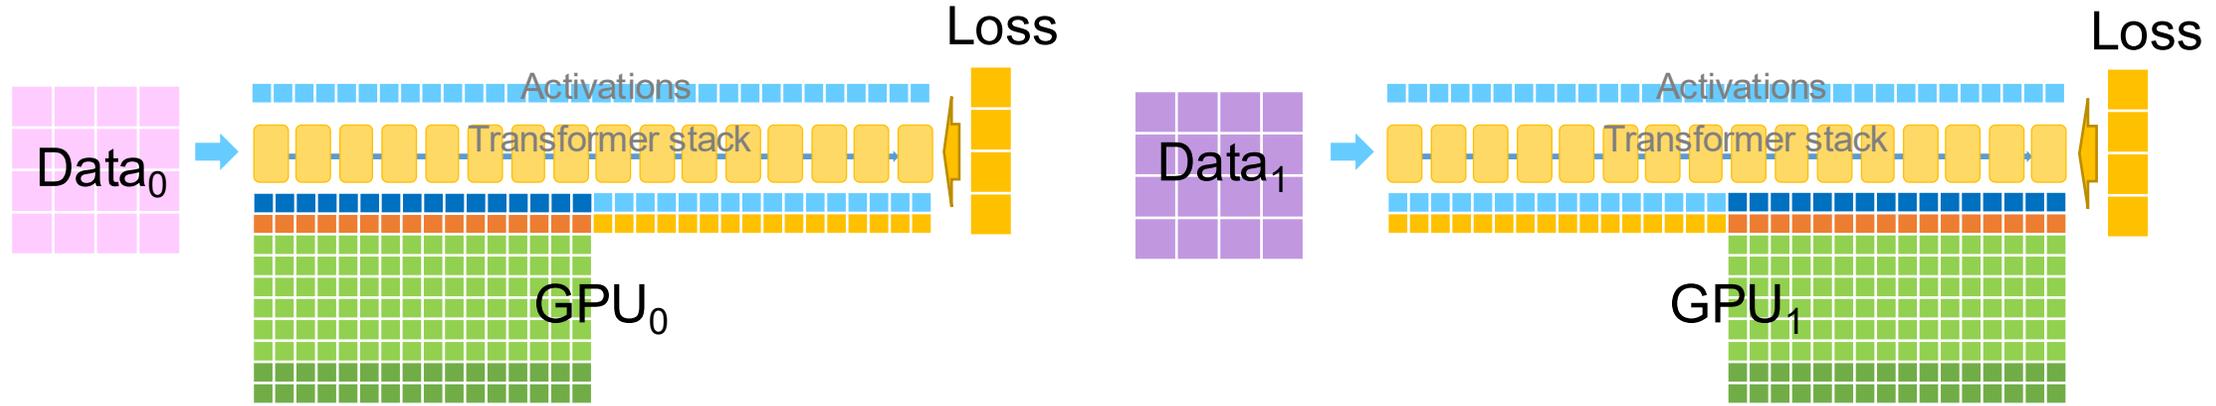
- Run Forward across the transformer blocks
- Backward propagation to generate FP16 gradients and AllReduce to average
- Update the FP32 weights with ADAM optimizer

# ZeRO Stage 1: Partitioning Optimizer States



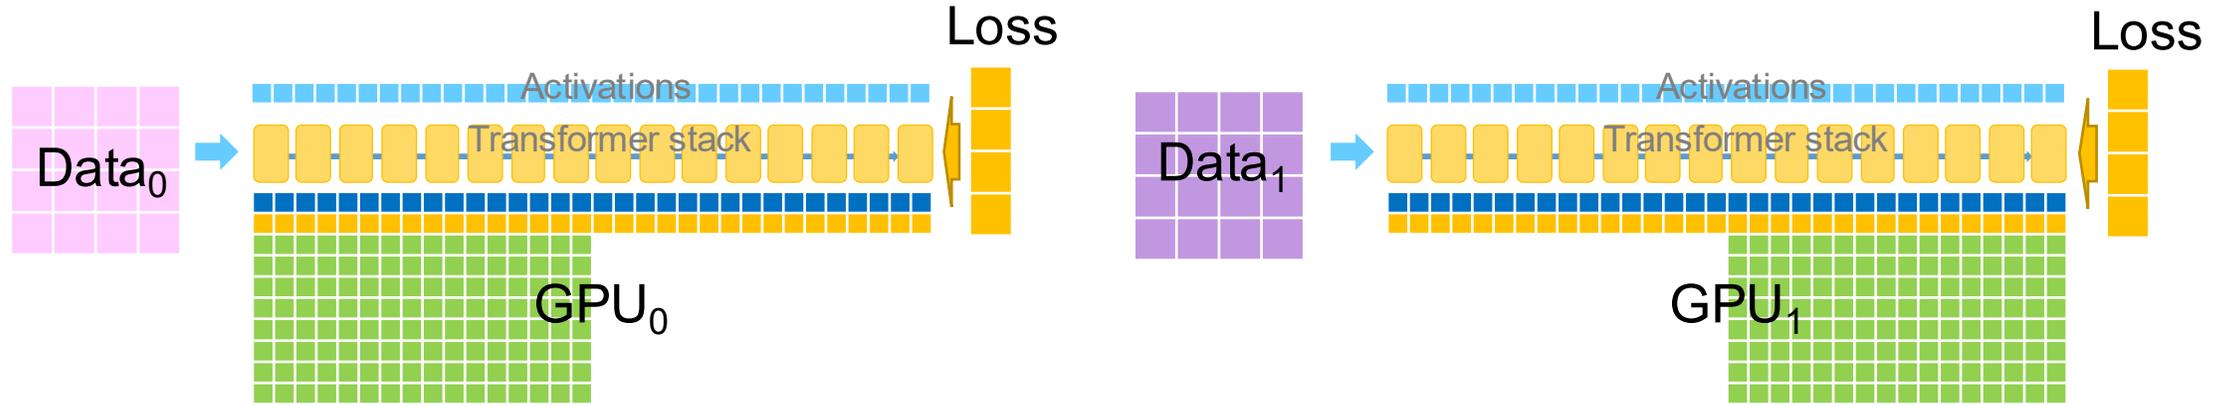
- Run Forward across the transformer blocks
- Backward propagation to generate FP16 gradients and AllReduce to average
- Update the FP32 weights with ADAM optimizer
- Update the FP16 weights

# ZeRO Stage 1: Partitioning Optimizer States



- Run Forward across the transformer blocks
- Backward propagation to generate FP16 gradients and AllReduce to average
- Update the FP32 weights with ADAM optimizer
- Update the FP16 weights
- All Gather the FP16 weights to complete the iteration

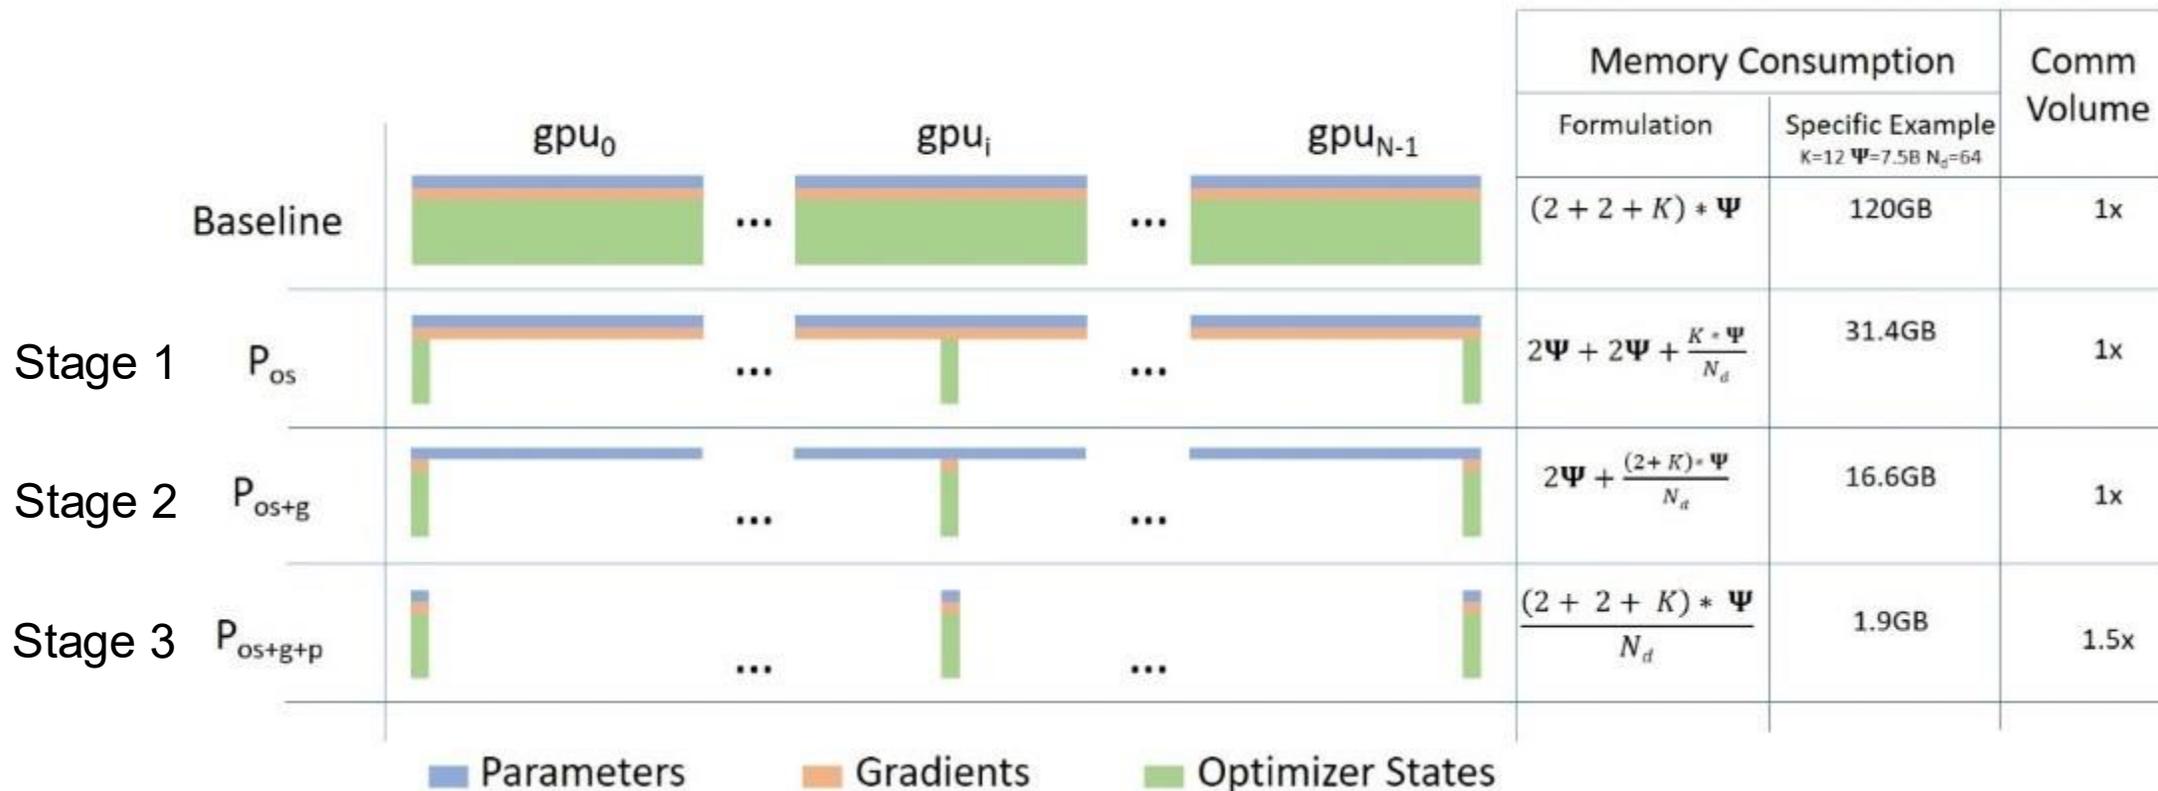
# ZeRO Stage 1: Partitioning Optimizer States



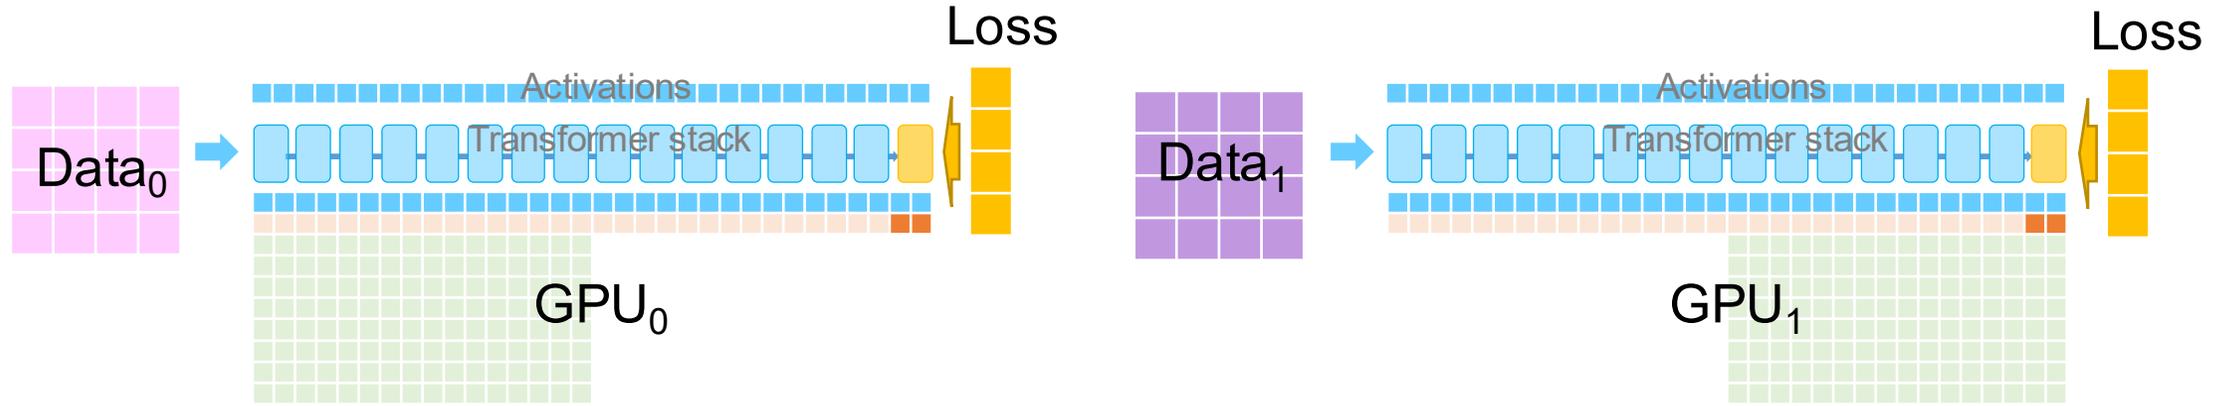
- Run Forward across the transformer blocks
- Backward propagation to generate FP16 gradients and AllReduce to average
- Update the FP32 weights with ADAM optimizer
- Update the FP16 weights
- All Gather the FP16 weights to complete the iteration

# ZeRO: Zero Redundancy Optimizer

- Progressive memory savings and communication volume
- Turning NLR 17.2B is powered by Stage 1 and Megatron

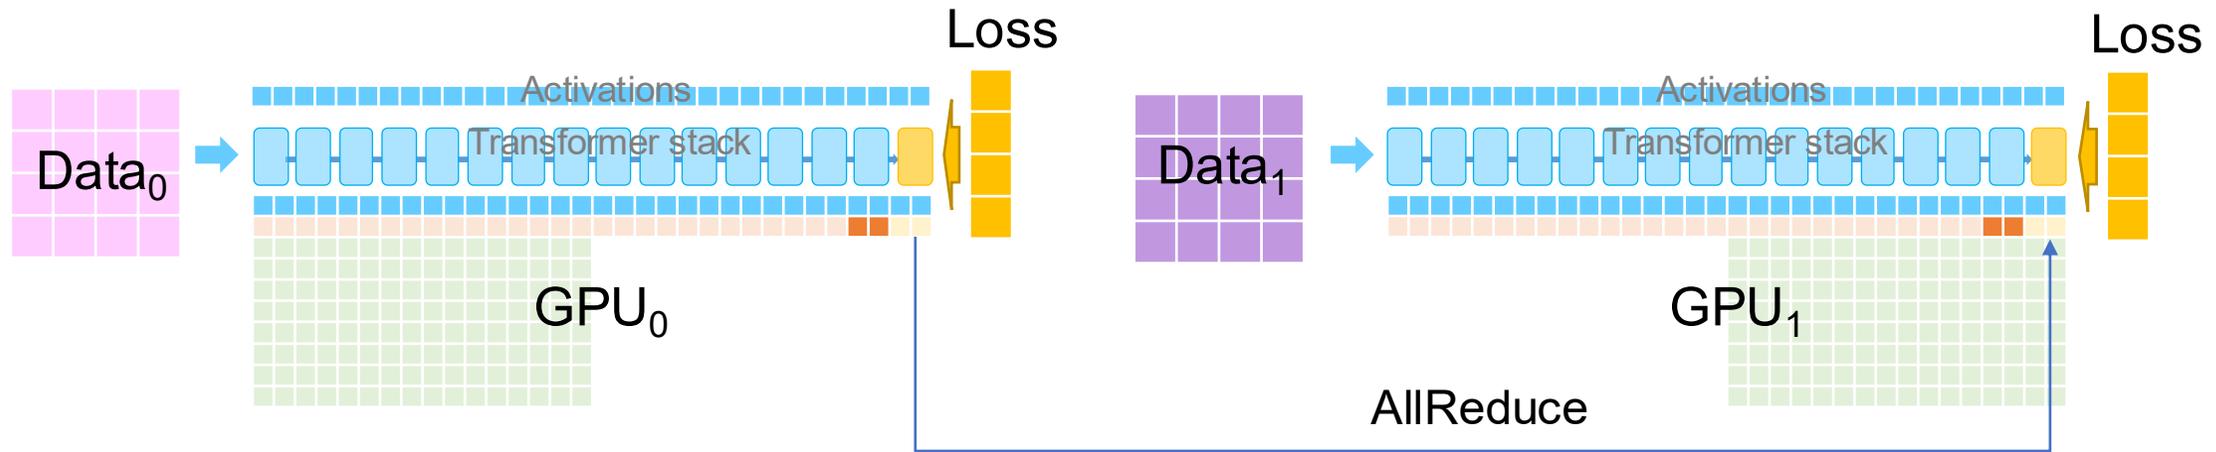


# ZeRO Stage 2: Partitioning Gradients



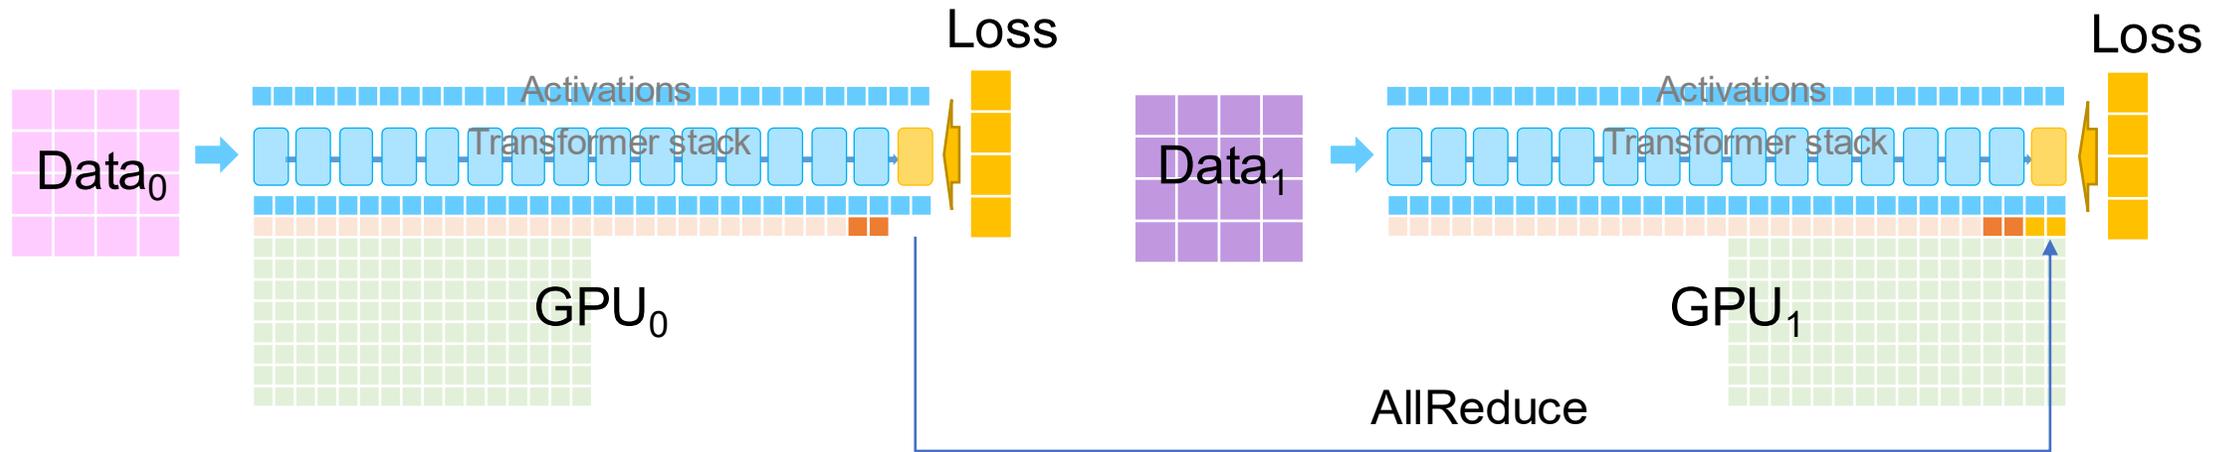
- Partitioning gradients across GPUs
- The forward process remains the same as stage 1

# ZeRO Stage 2: Partitioning Gradients



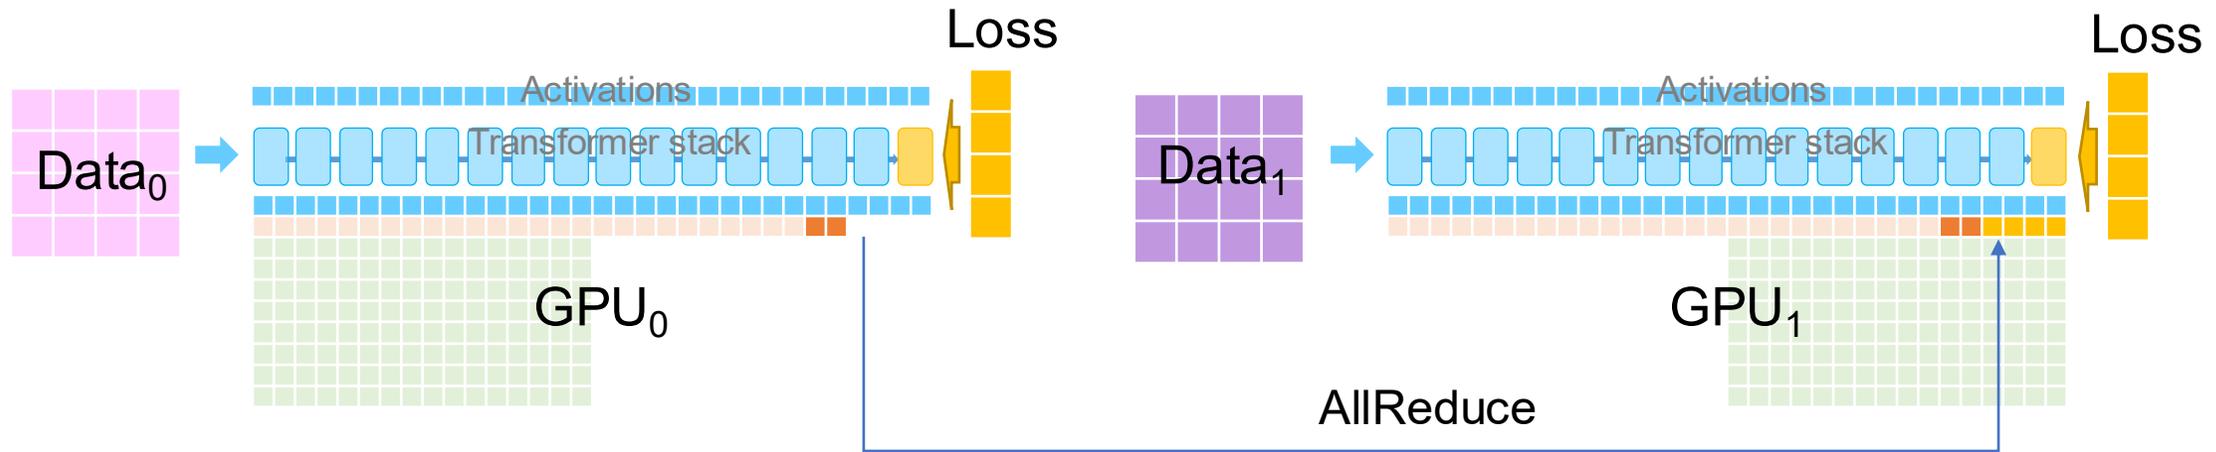
- Partitioning gradients across GPUs
- Perform AllReduce right after back propagation of each layer

# ZeRO Stage 2: Partitioning Gradients



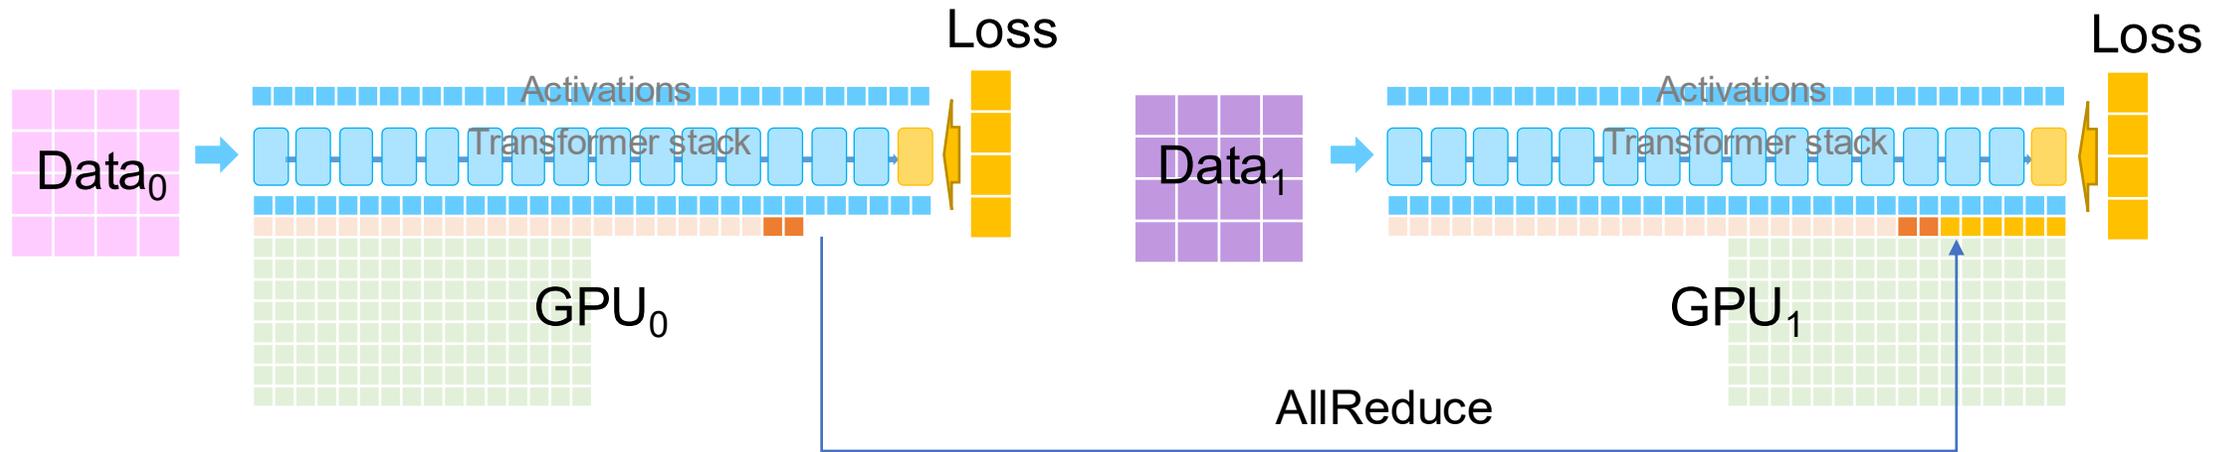
- Partitioning gradients across GPUs
- Only one GPU keeps the gradients after AllReduce

# ZeRO Stage 2: Partitioning Gradients



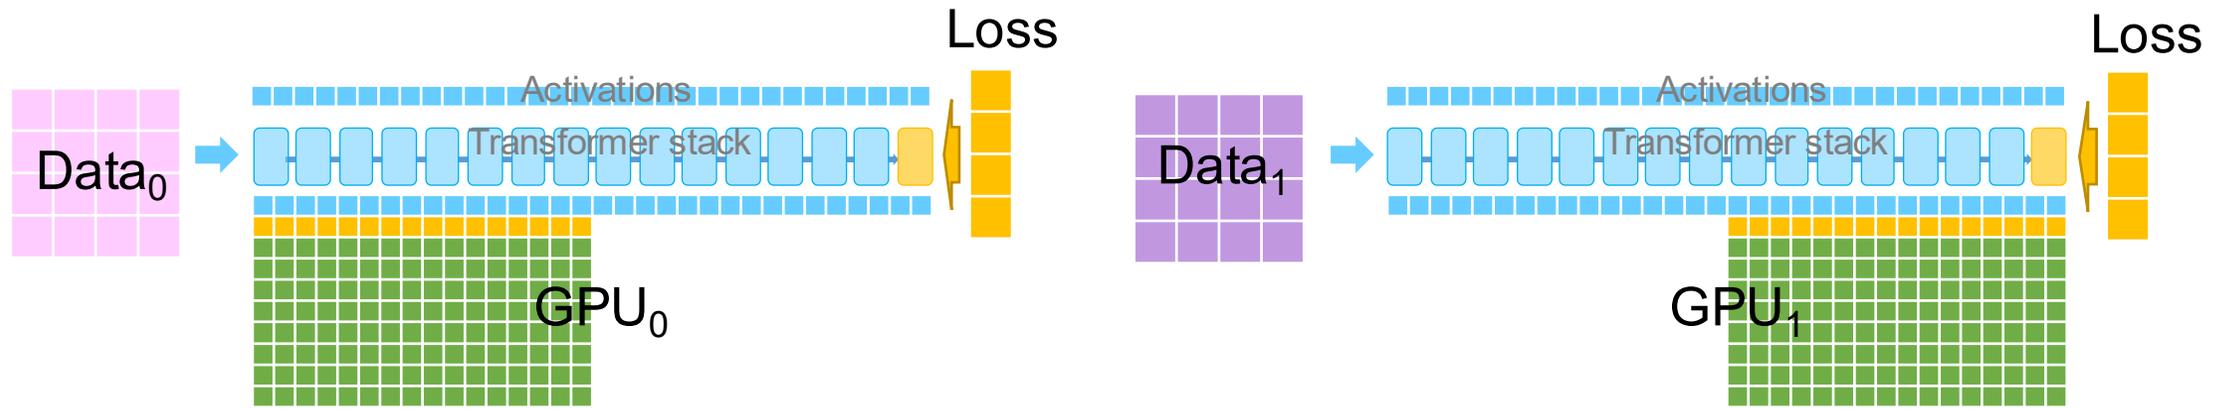
- Partitioning gradients across GPUs
- Reduce gradients on GPUs responsible for updating parameters

# ZeRO Stage 2: Partitioning Gradients



- Partitioning gradients across GPUs
- Reduce gradients on GPUs responsible for updating parameters

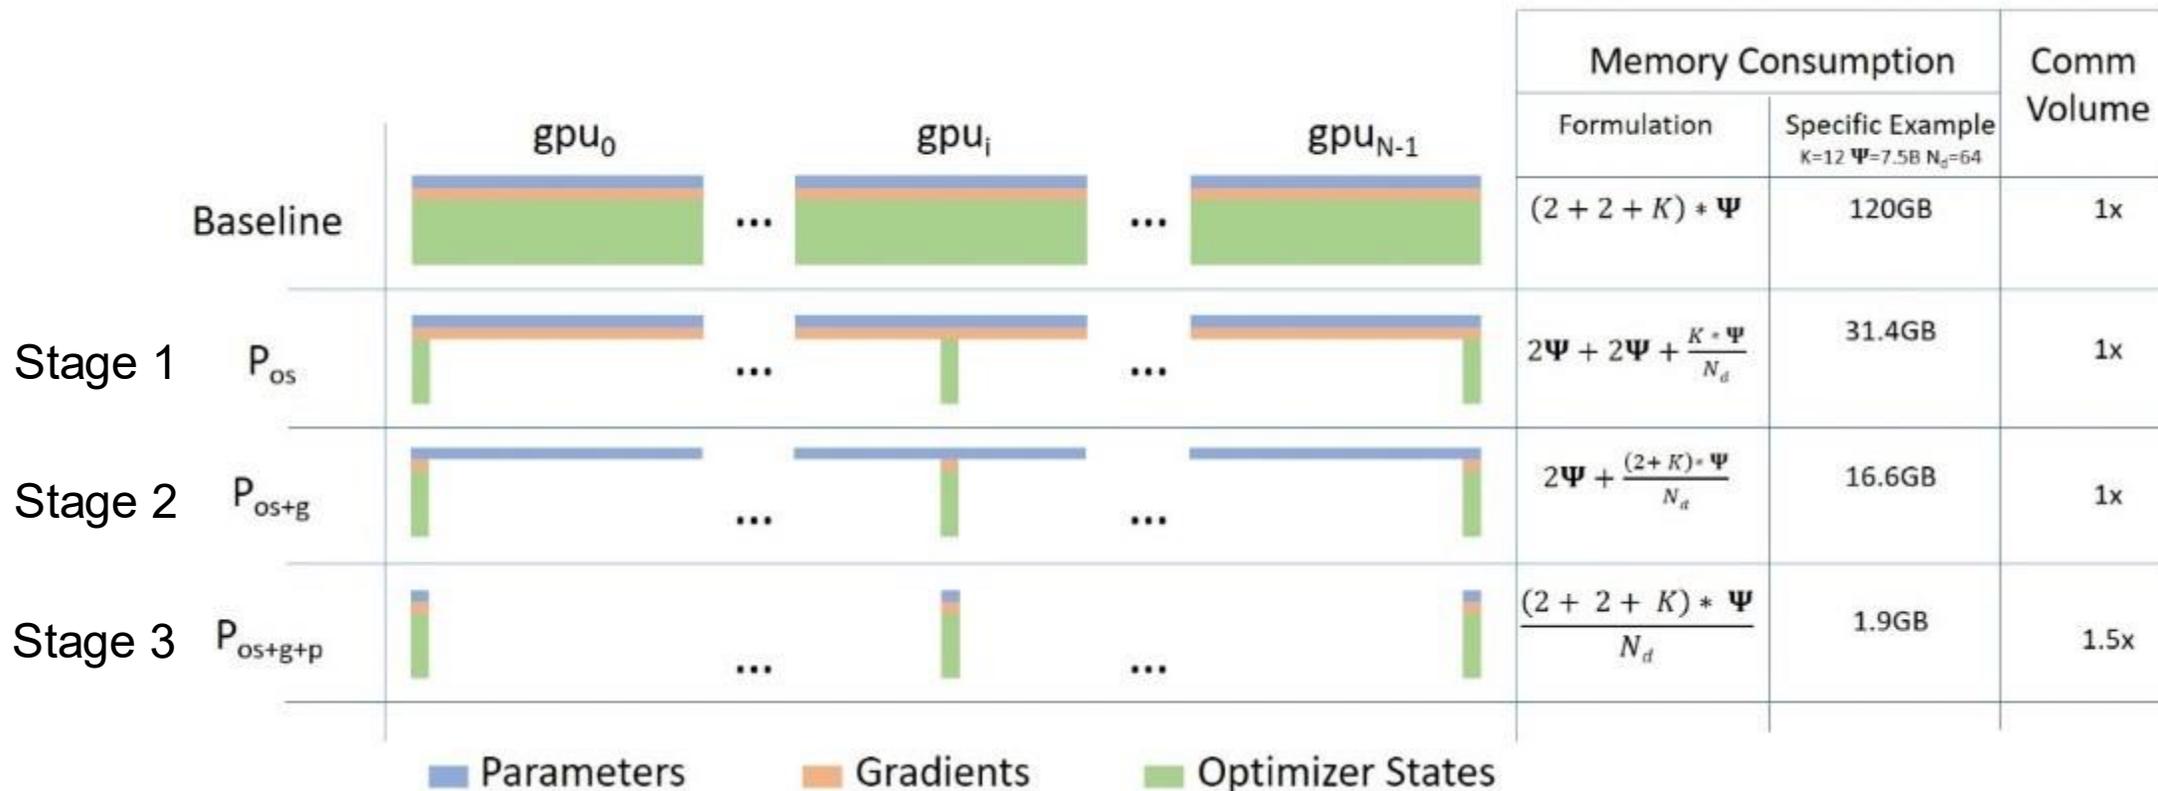
# ZeRO Stage 2: Partitioning Gradients



- Partitioning gradients across GPUs
- Reduce gradients on GPUs responsible for updating parameters

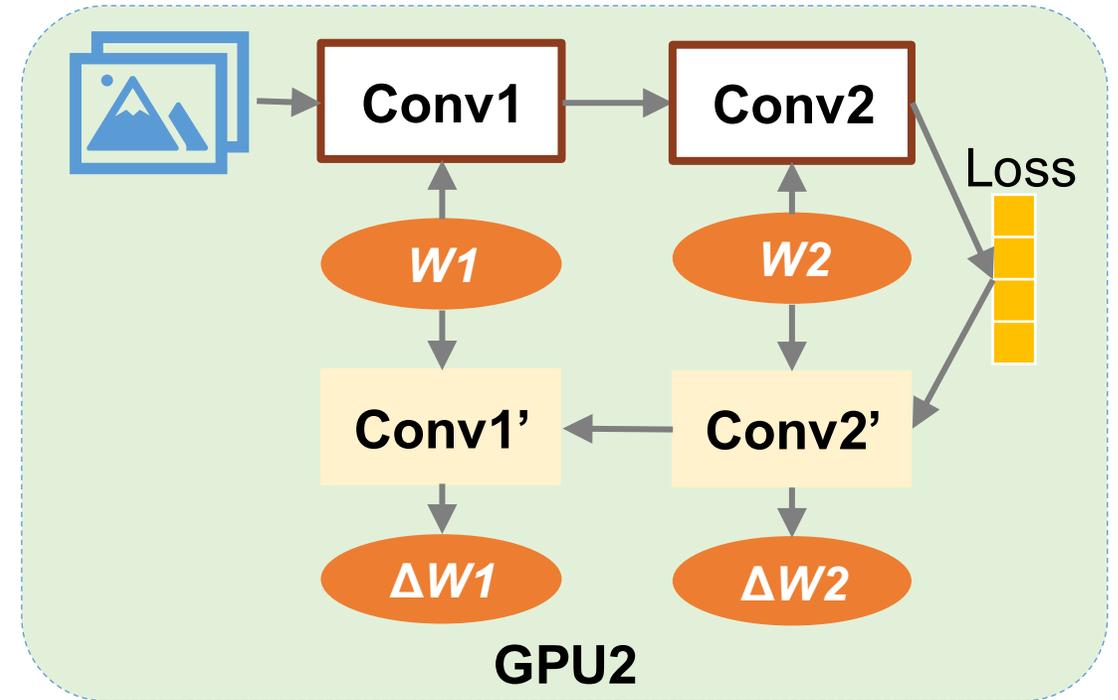
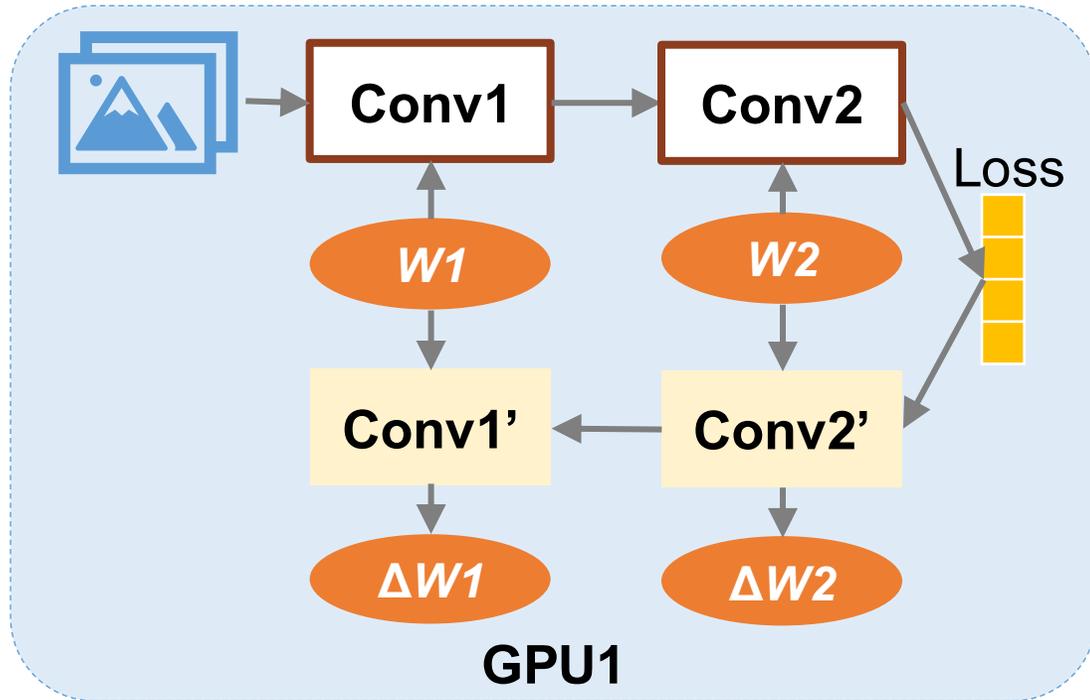
# ZeRO: Zero Redundancy Optimizer

- Progressive memory savings and communication volume
- Turning NLR 17.2B is powered by Stage 1 and Megatron



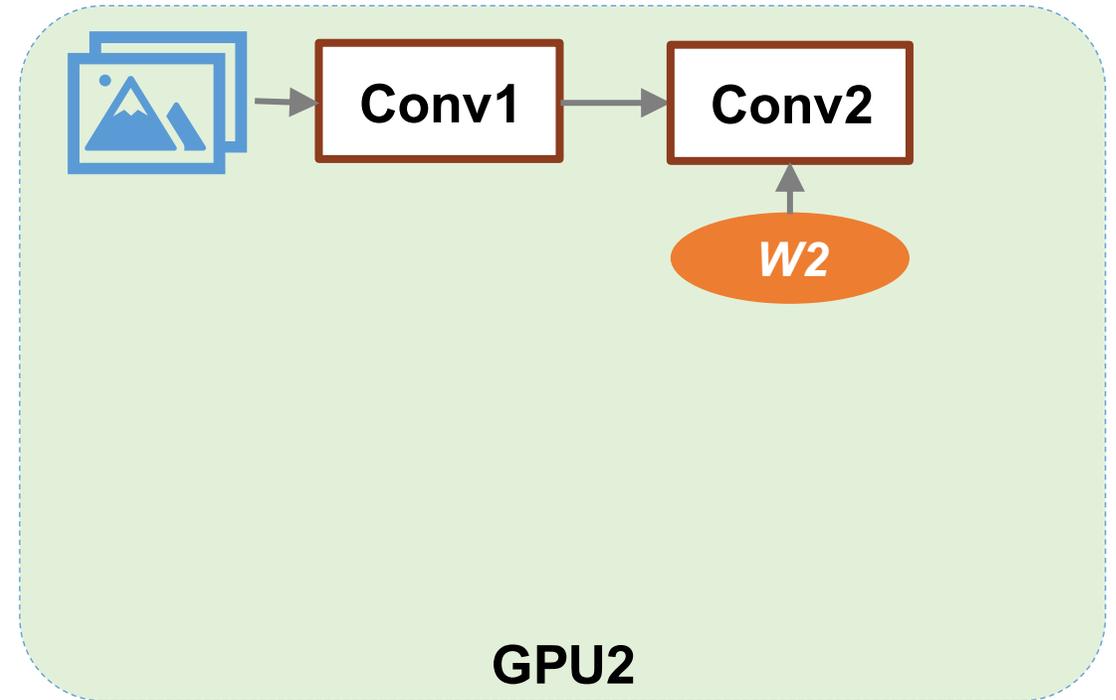
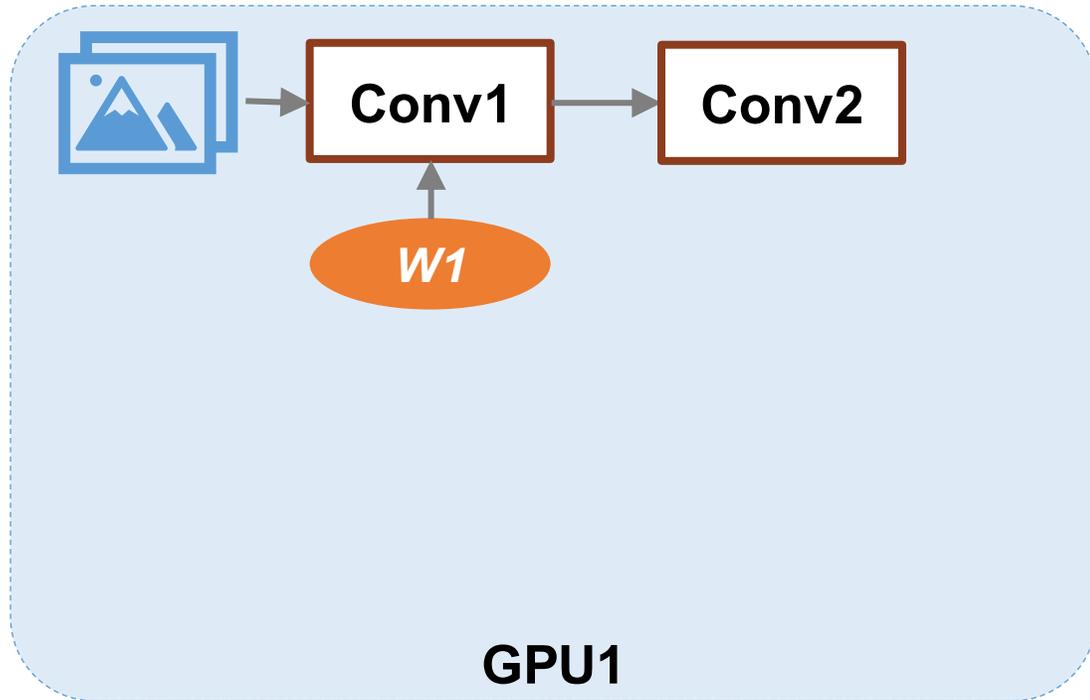
# ZeRO Stage 3: Partitioning Parameters

- In data parallel training, all GPUs keep **all** parameters during training



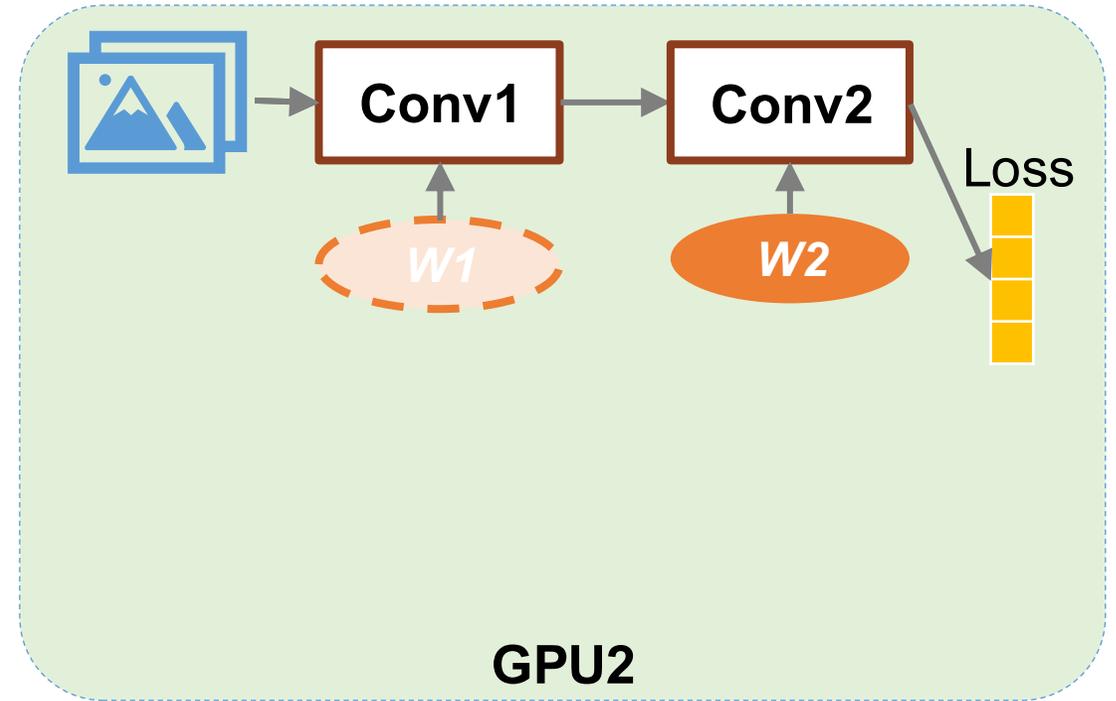
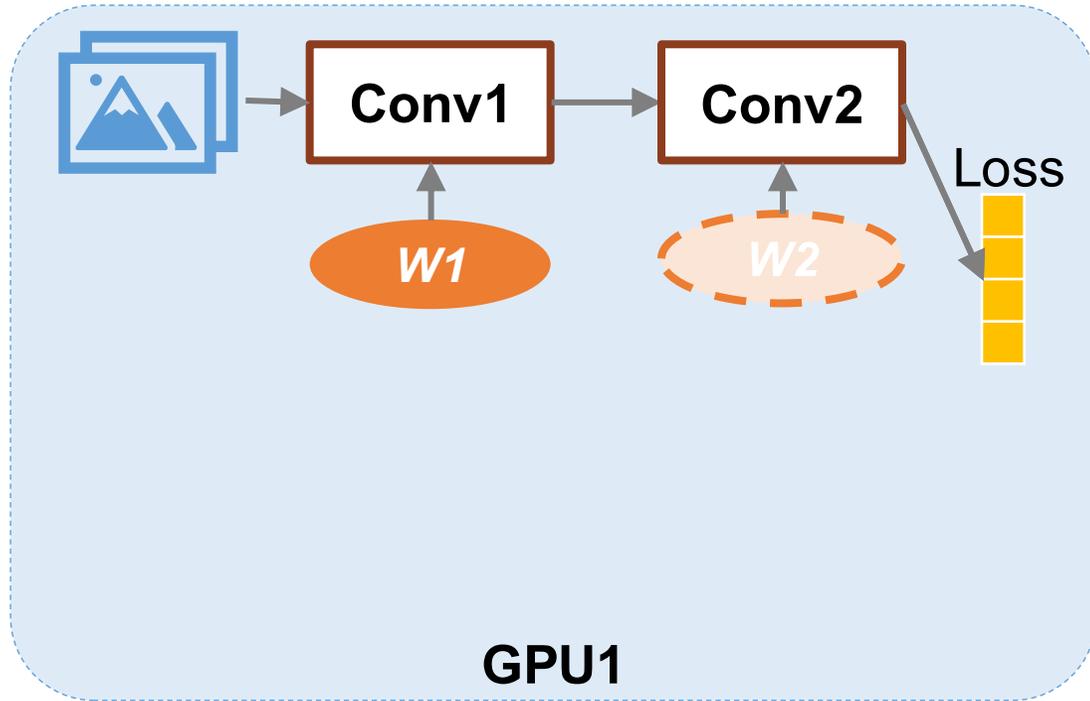
# ZeRO Stage 3: Partitioning Parameters

- In ZeRO, model parameters are partitioned across GPUs



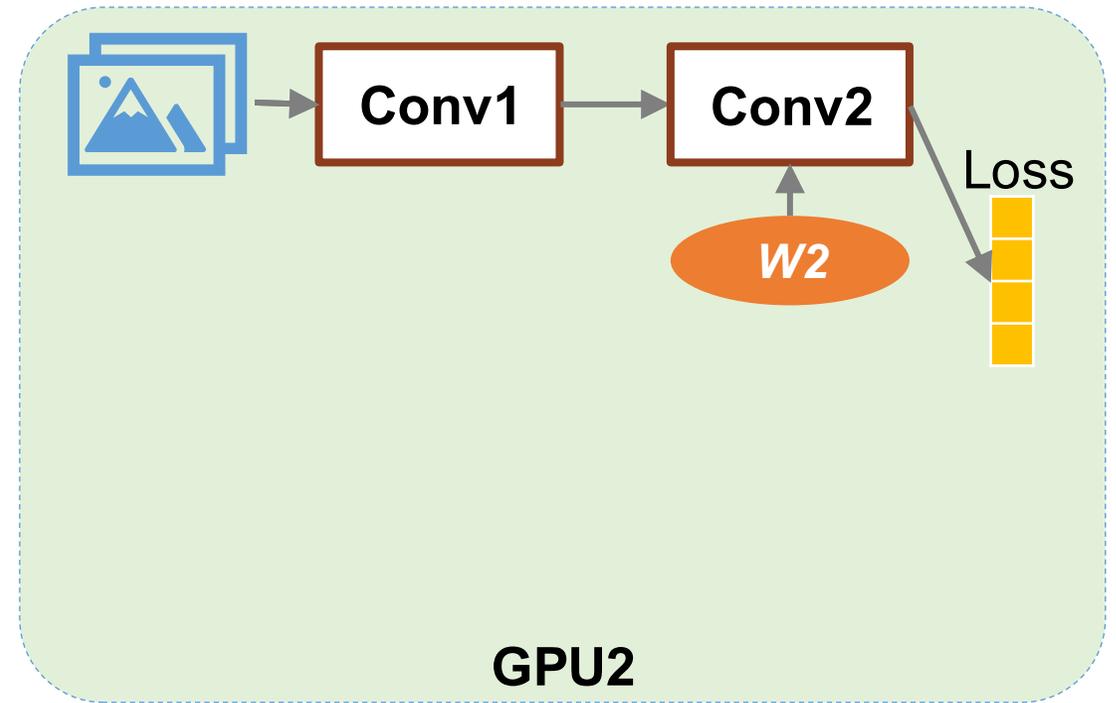
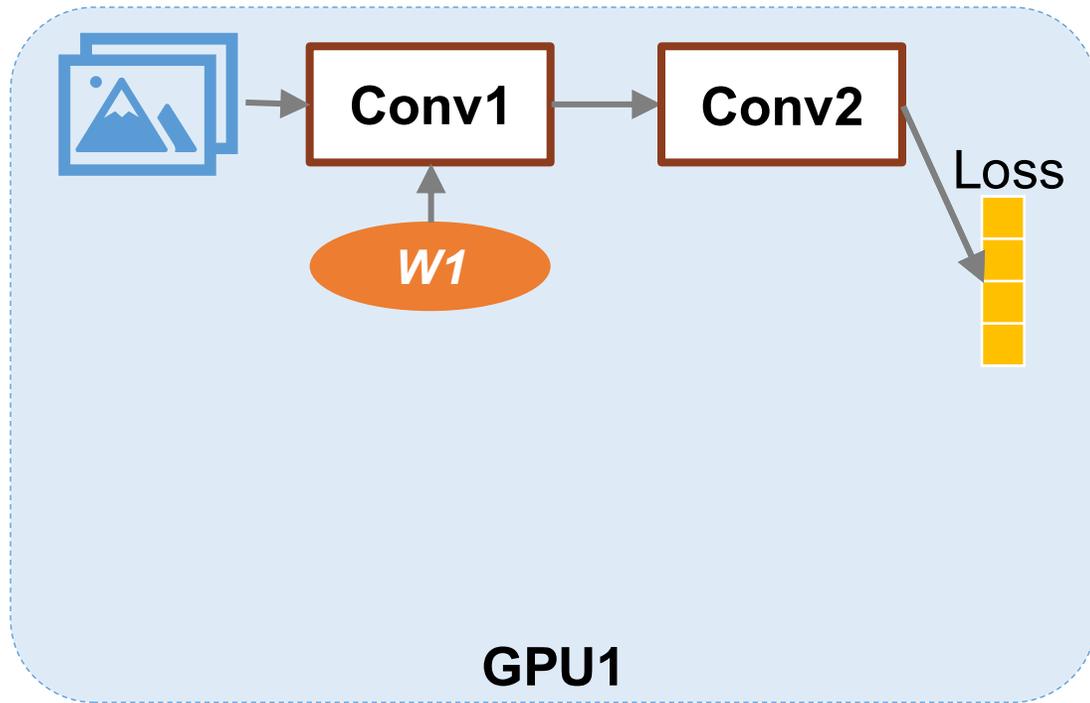
# ZeRO Stage 3: Partitioning Parameters

- In ZeRO, model parameters are partitioned across GPUs
- GPUs broadcast their parameters during forward



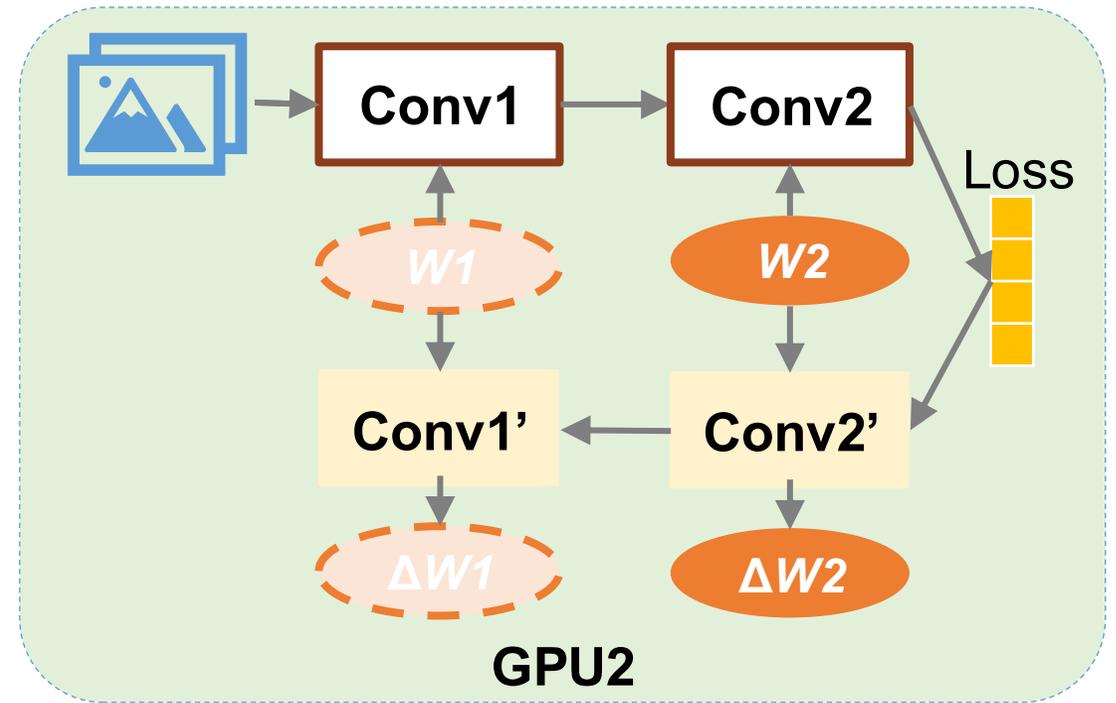
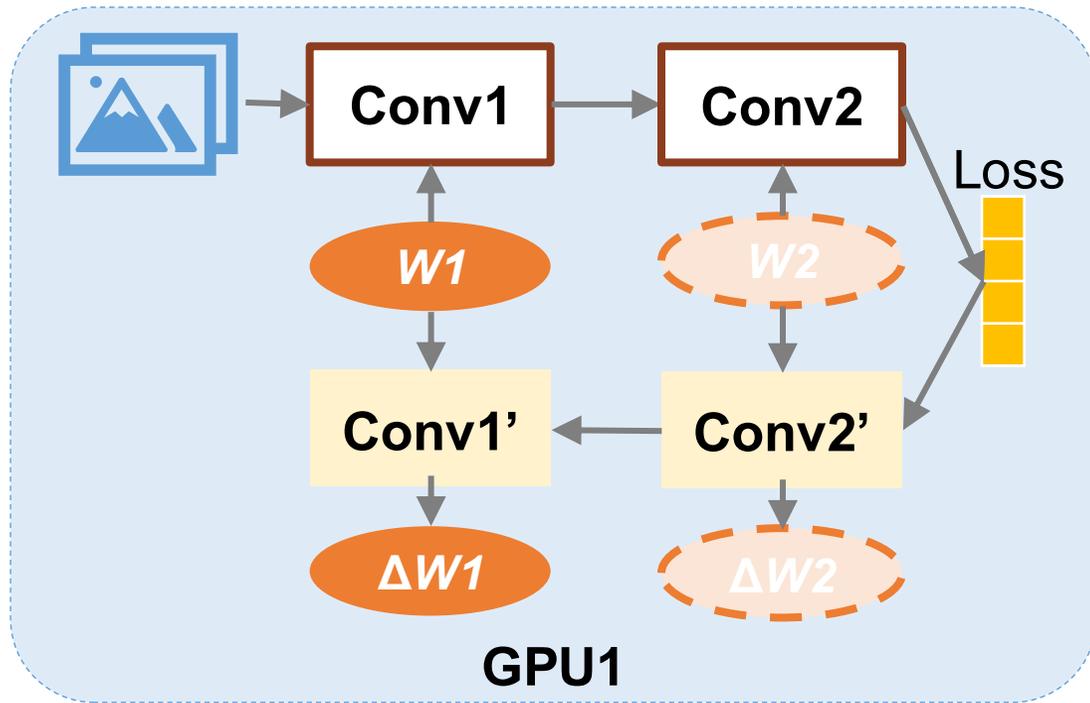
# ZeRO Stage 3: Partitioning Parameters

- In ZeRO, model parameters are partitioned across GPUs
- Parameters are discarded right after use



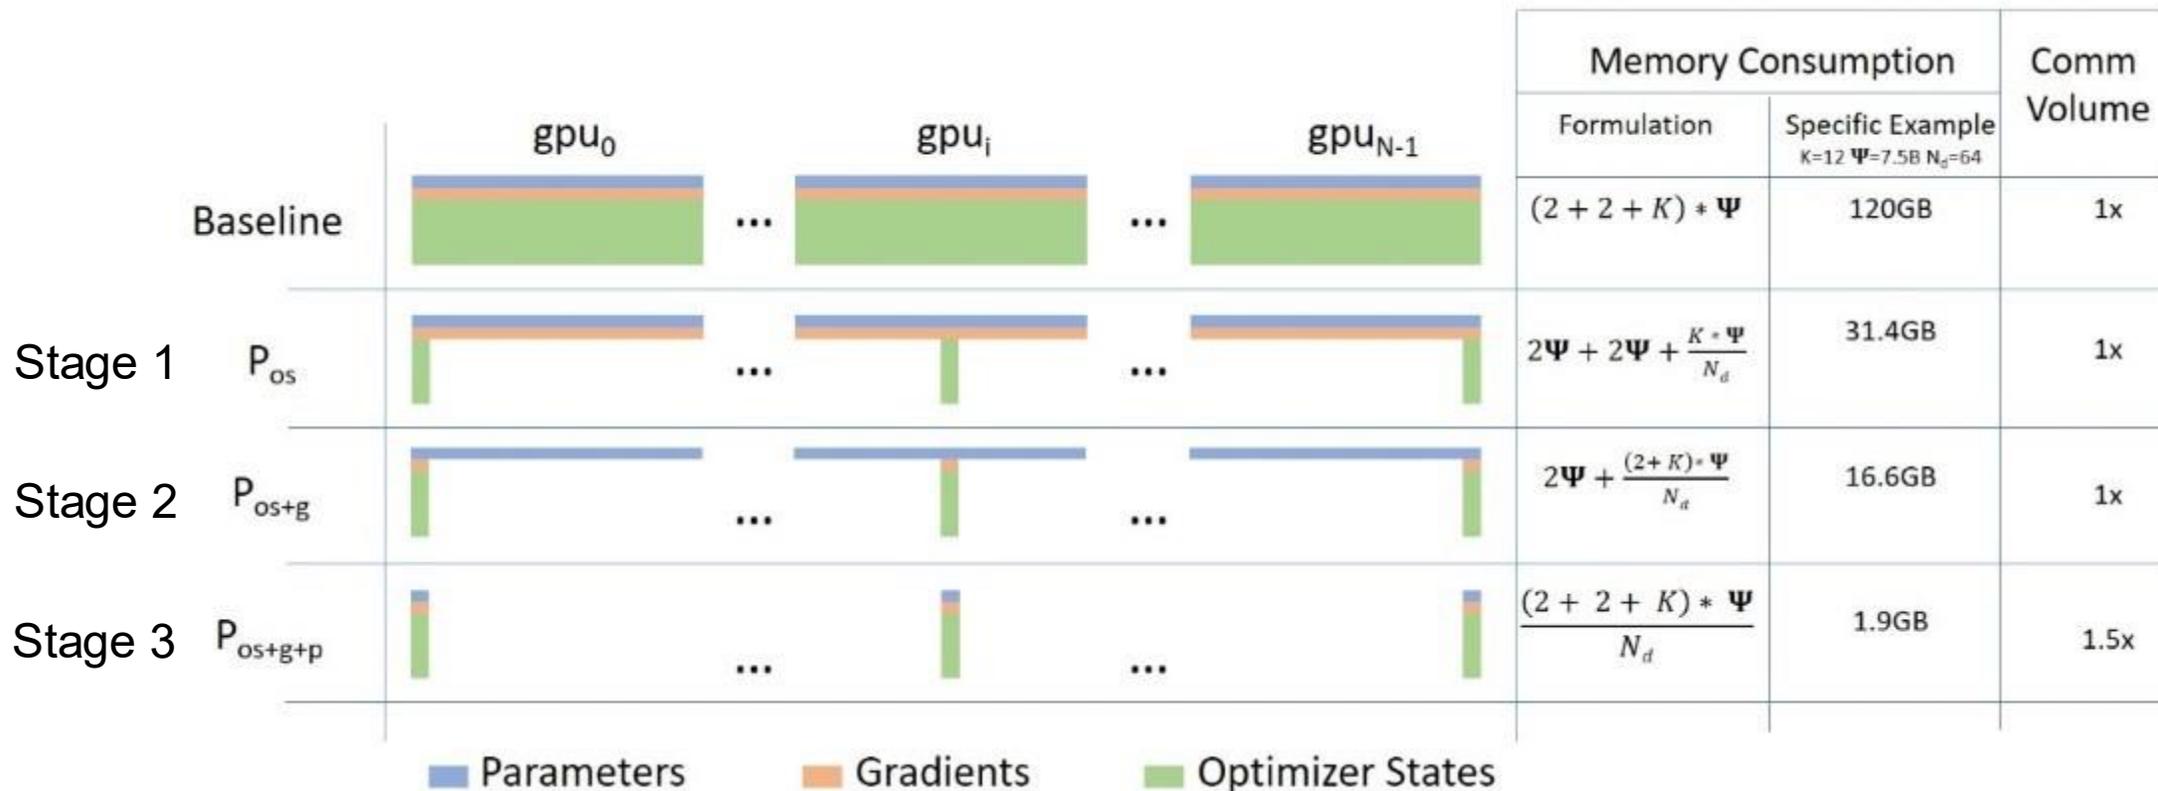
# ZeRO Stage 3: Partitioning Parameters

- In ZeRO, model parameters are partitioned across GPUs
- GPUs broadcast their parameters again during backward



# ZeRO: Zero Redundancy Optimizer

- ZeRO has three different stages
- Progressive memory savings and communication volume



# Summary

- Data-parallel training
  - Parameter server
  - Ring AllReduce
  - Tree AllReduce
  - Butterfly AllReduce
- ZeRO: zero redundancy optimizer