

# Agenda

What is NCCL?

Why use NCCL?

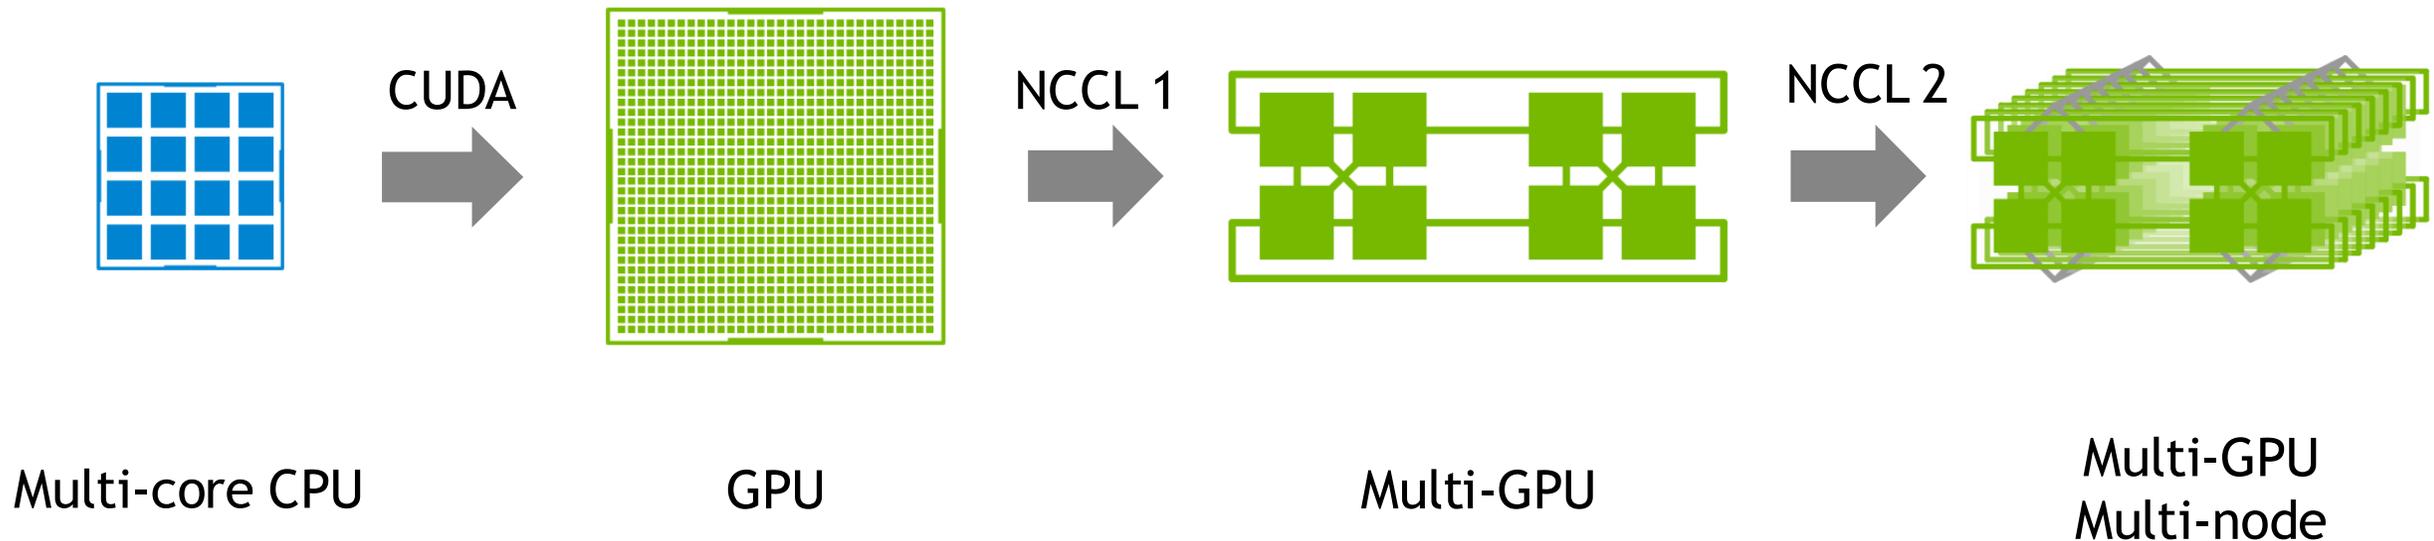
How to use NCCL?

**What is NCCL?**

# DEEP LEARNING ON GPUS

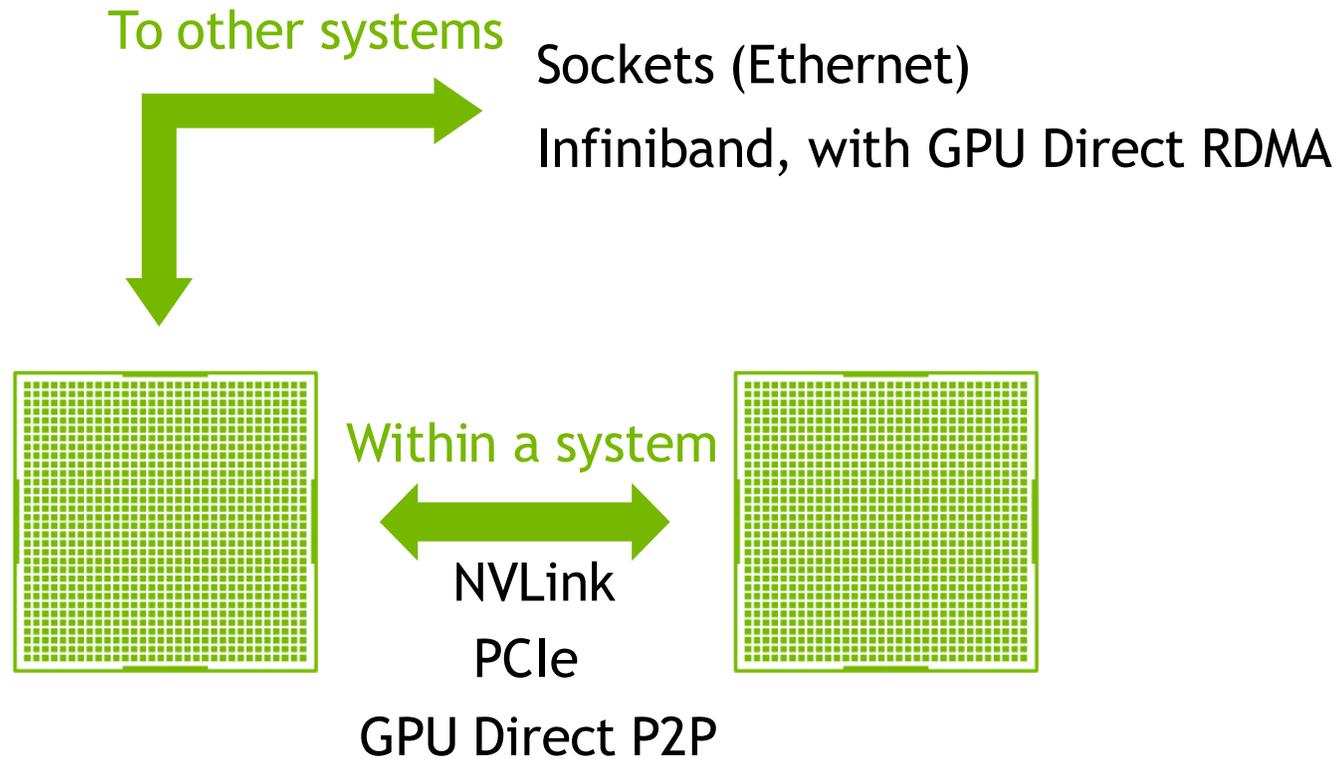
Making DL training times shorter

Deeper neural networks, larger data sets ... training is a very, very long operation !



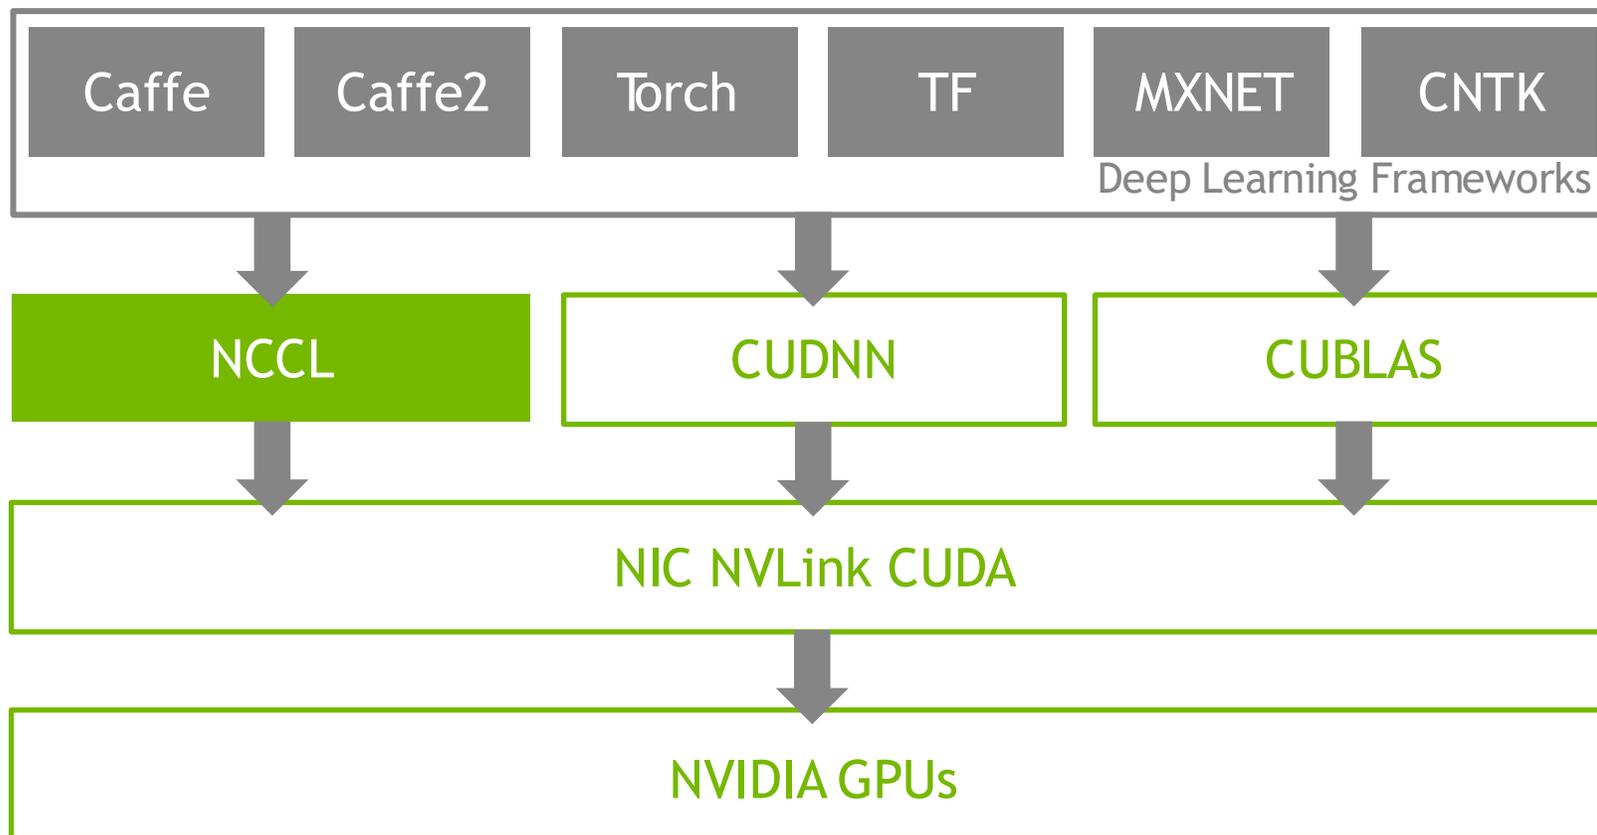
# NCCL

A multi-GPU communication library



# NCCL

## Architecture



**Why use NCCL?**

# DESIGN

Optimized collective communication library between CUDA devices.

Easy to integrate into any DL framework, as well as MPI.

Runs on the GPU using asynchronous **CUDA kernels**, for faster access to GPU traditional HPC apps using memory, parallel reductions, NVLink usage.

Operates on **CUDA pointers**. Operations are tied to a **CUDA stream**.

Uses as little threads as possible to **permit other computation** to progress simultaneously.

**What is collective communication?**

# Common Communication Patterns

## Point-to-point communication

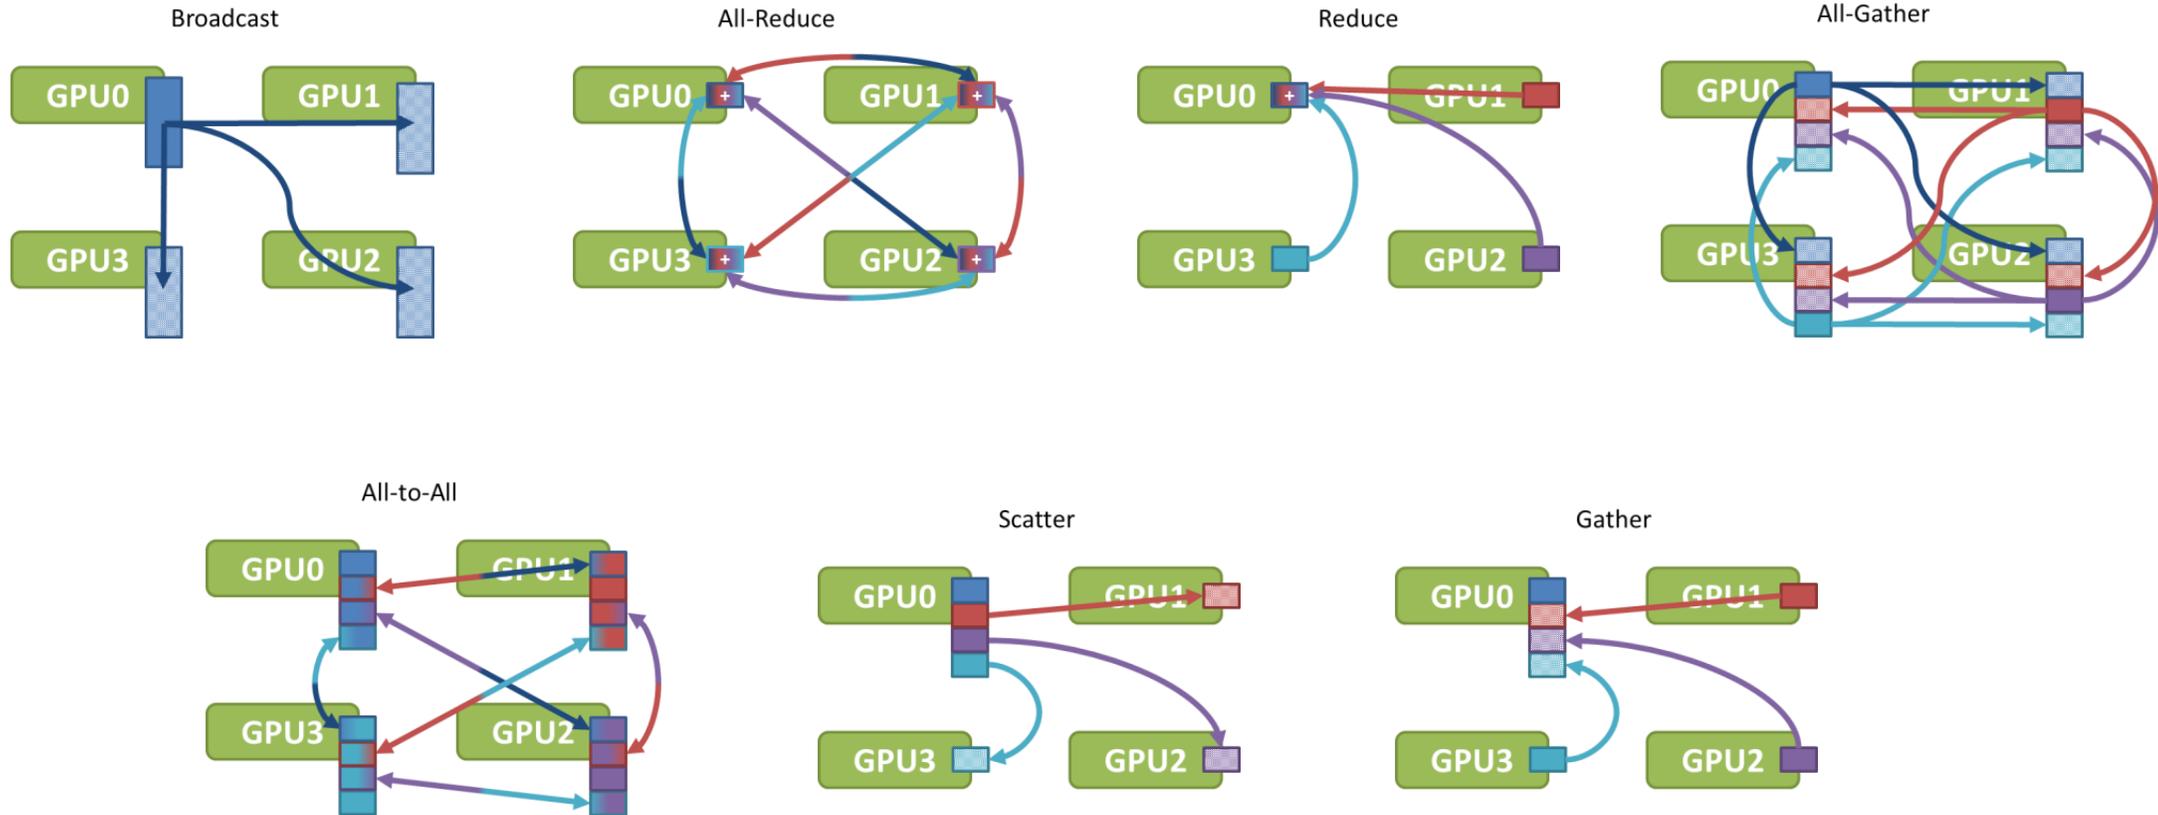
- Single sender, single receiver
- Relatively easy to implement efficiently

## Collective communication

- Multiple senders and/or receivers
- Patterns include broadcast, scatter, gather, reduce, all-to-all, ...
- Difficult to implement efficiently

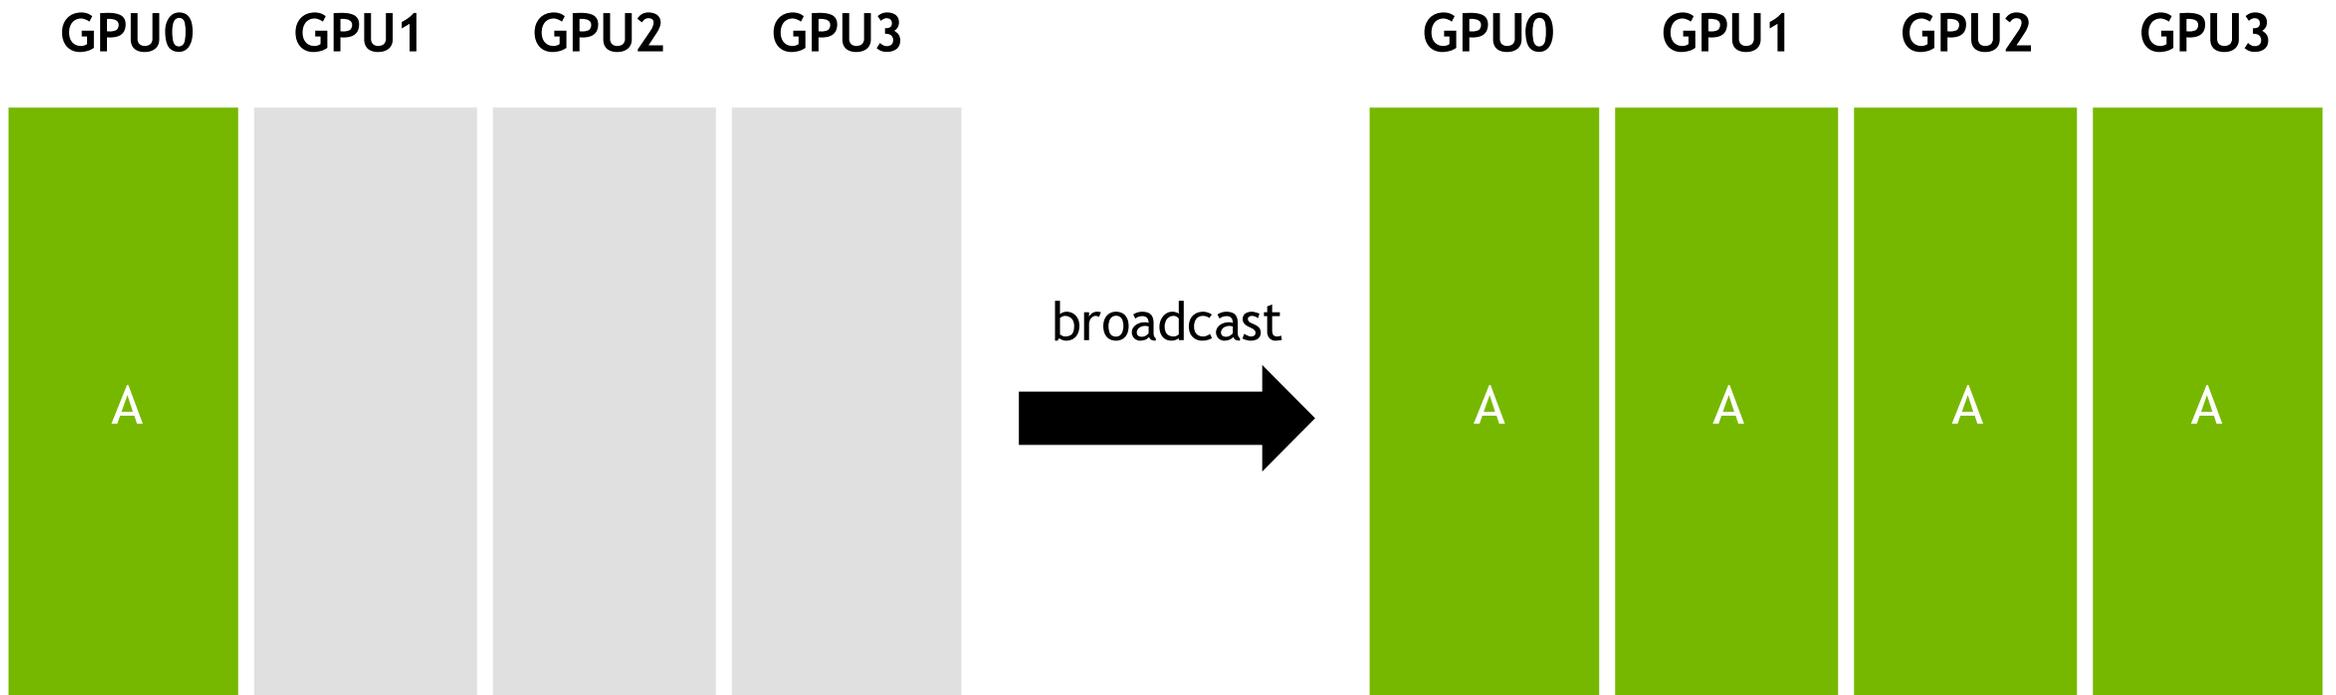
# COLLECTIVE COMMUNICATION

Multiple senders and/or receivers



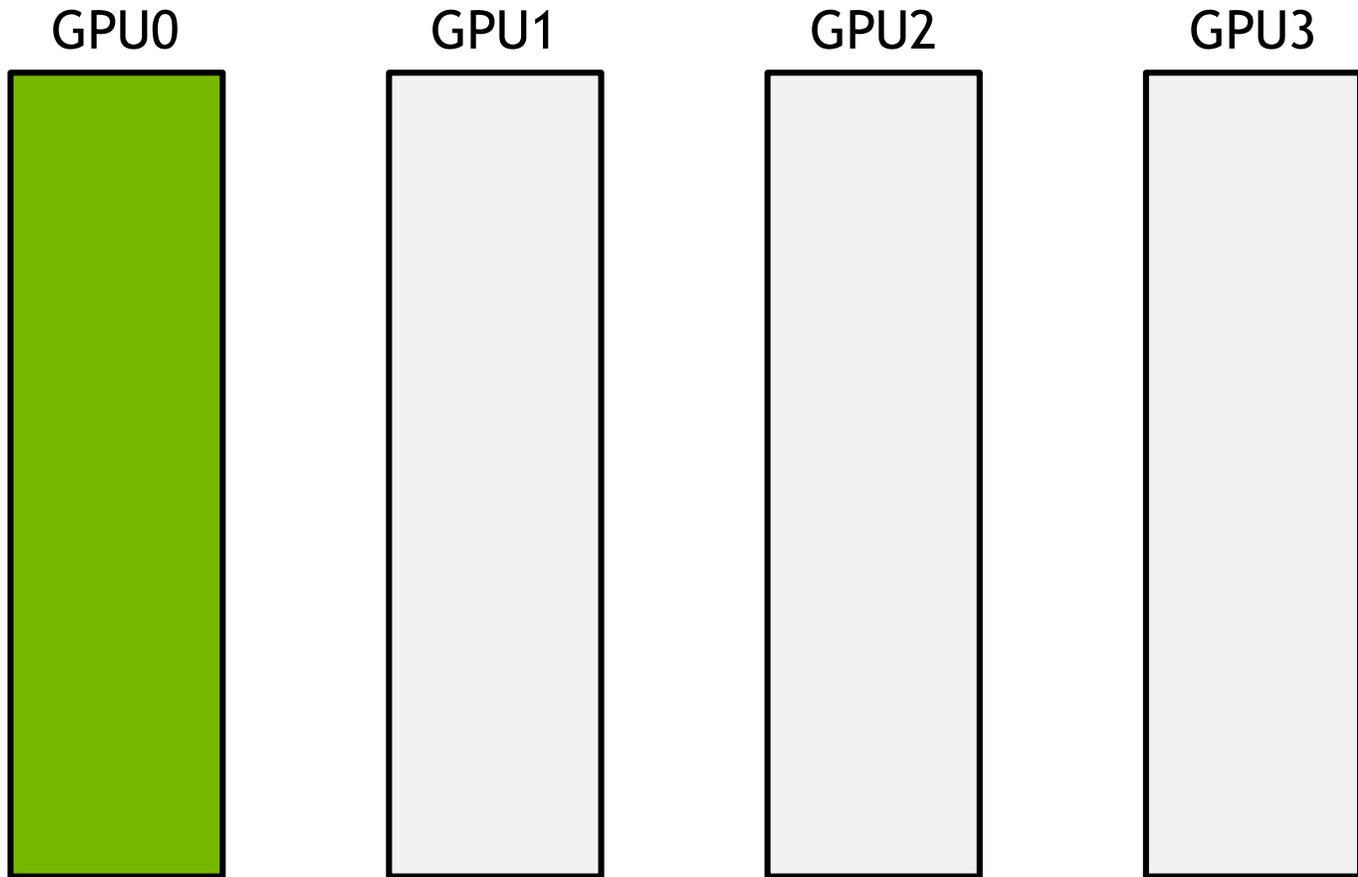
# BROADCAST

One sender, multiple receivers



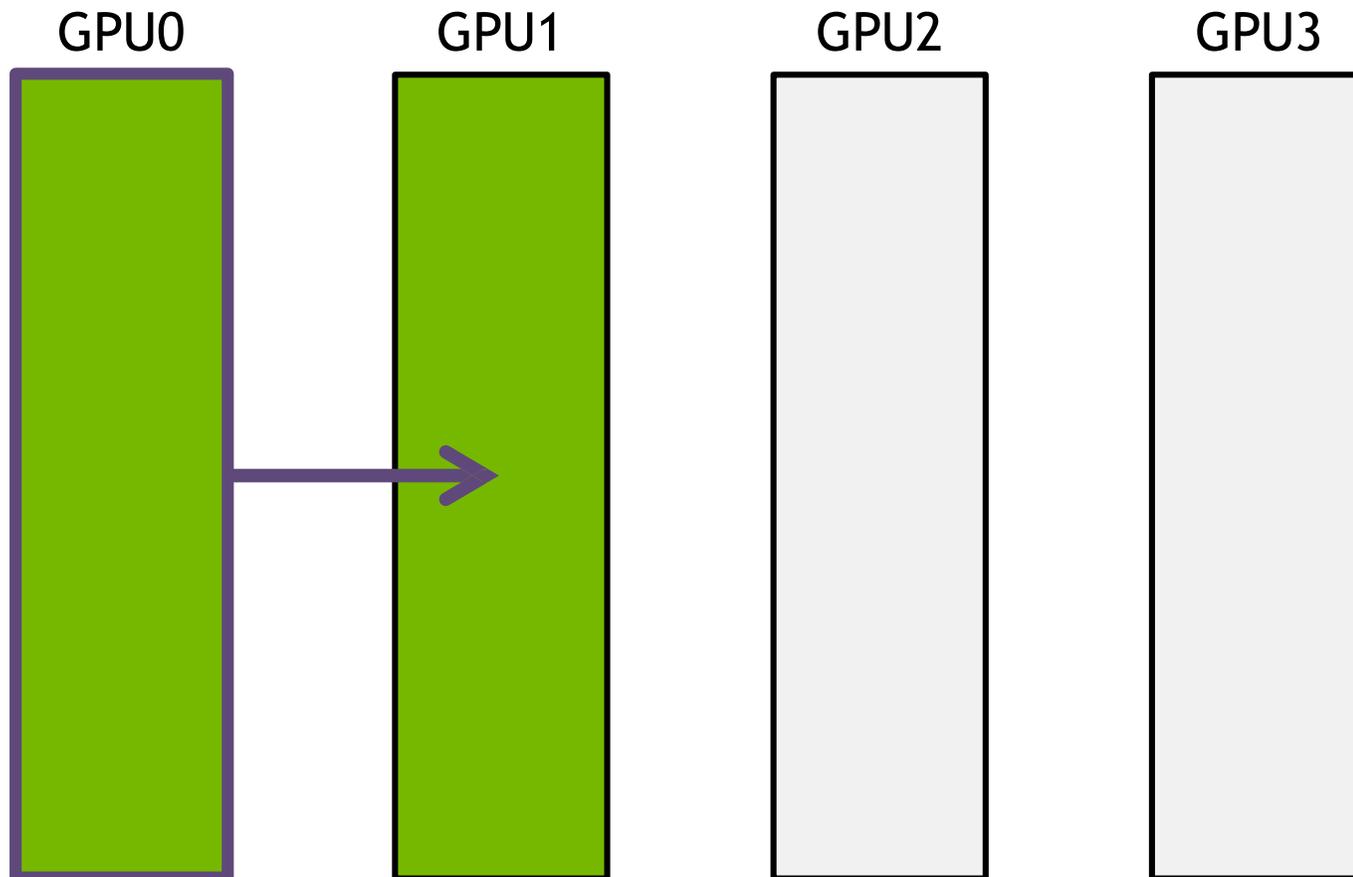
# BROADCAST

with unidirectional ring



# BROADCAST

with unidirectional ring



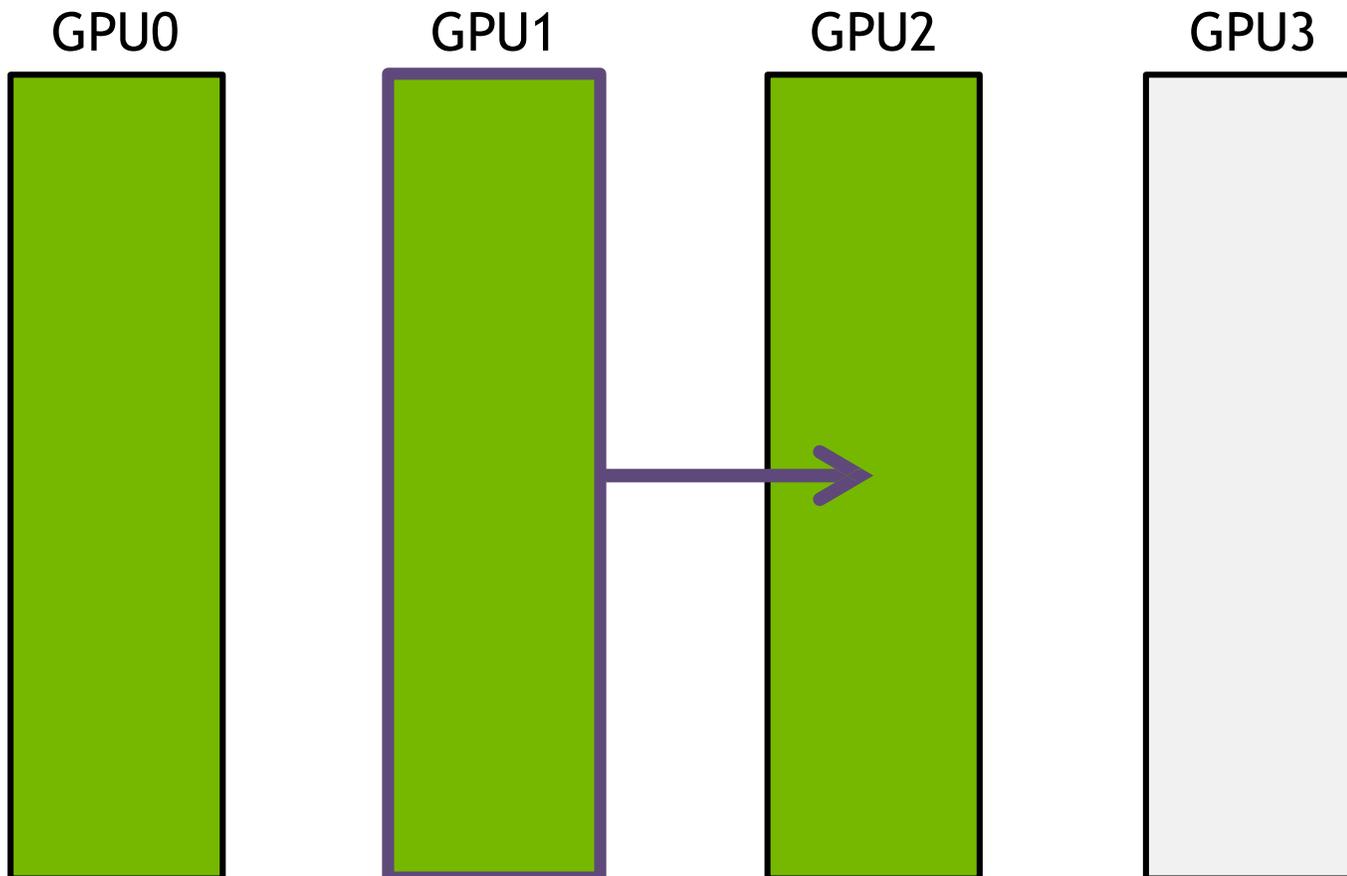
Step 1:  $\Delta t = N/B$

$N$ : bytes to broadcast

$B$ : bandwidth of each link

# BROADCAST

with unidirectional ring



Step 1:  $\Delta t = N/B$

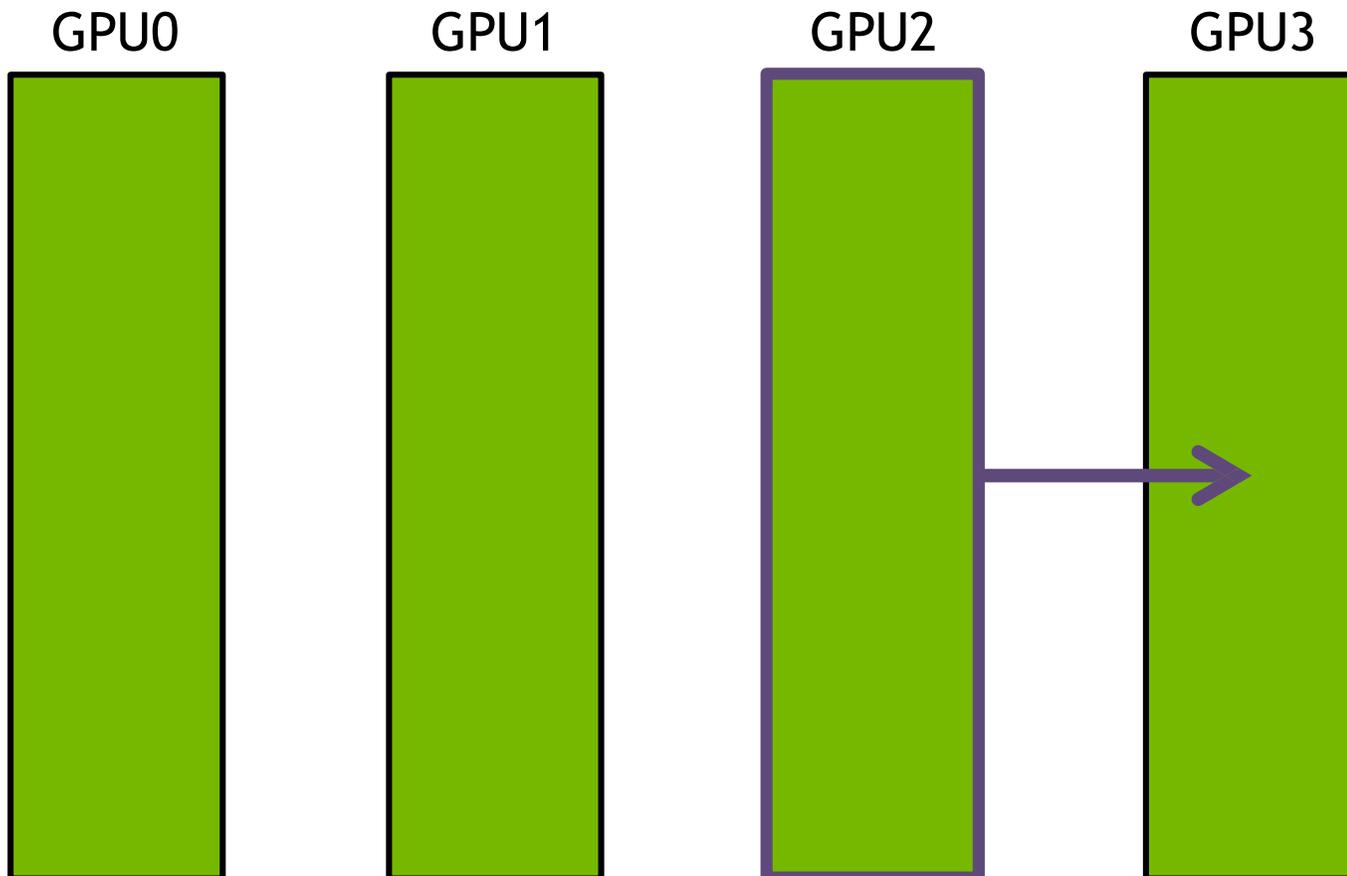
Step 2:  $\Delta t = N/B$

$N$ : bytes to broadcast

$B$ : bandwidth of each link

# BROADCAST

with unidirectional ring



Step 1:  $\Delta t = N/B$

Step 2:  $\Delta t = N/B$

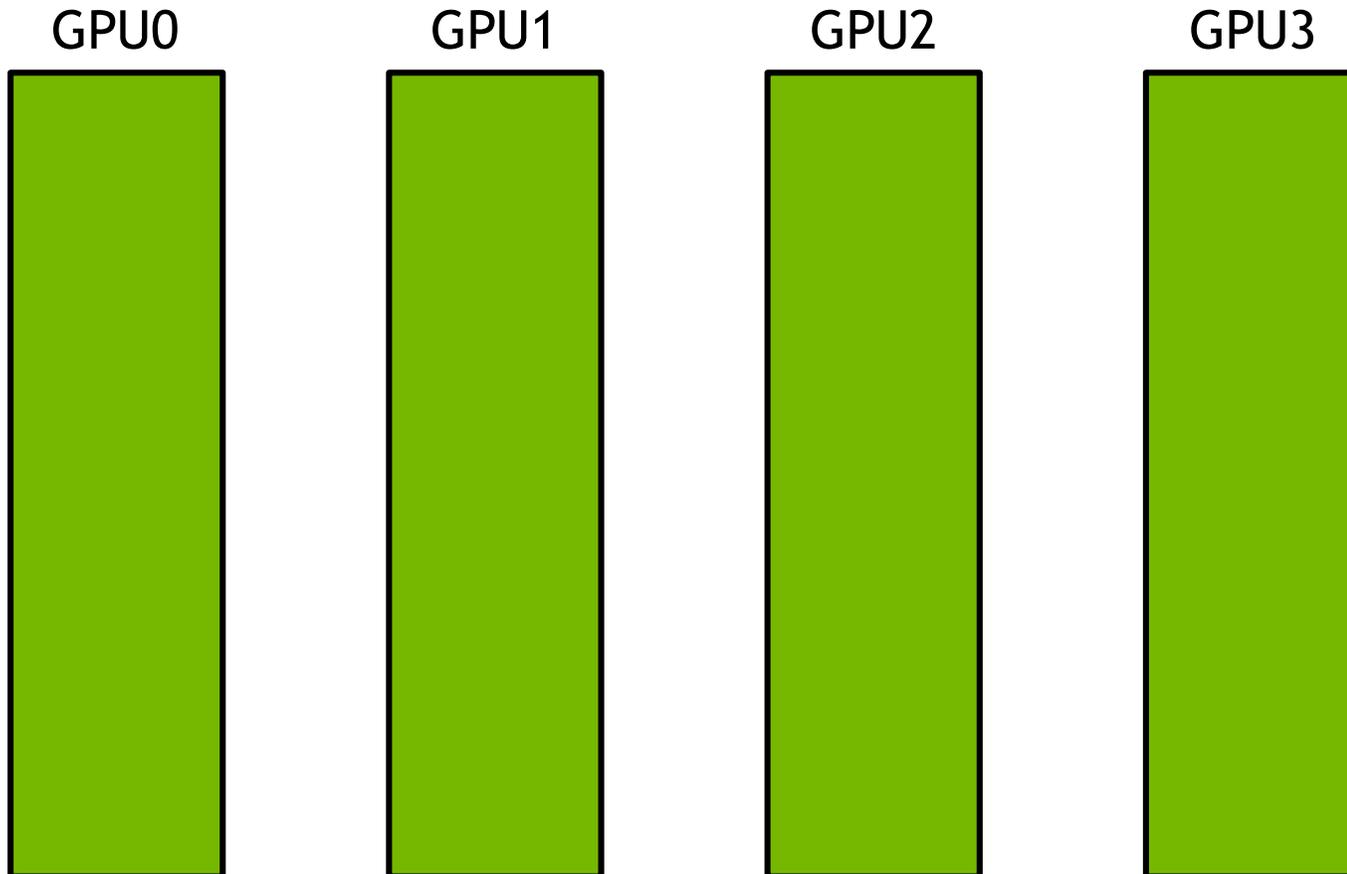
Step 3:  $\Delta t = N/B$

$N$ : bytes to broadcast

$B$ : bandwidth of each link

# BROADCAST

with unidirectional ring



Step 1:  $\Delta t = N/B$

Step 2:  $\Delta t = N/B$

Step 3:  $\Delta t = N/B$

Total time:  $(k - 1)N/B$

$N$ : bytes to broadcast

$B$ : bandwidth of each link

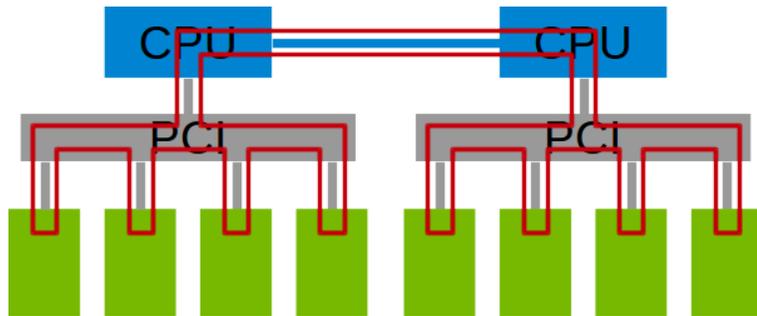
$k$ : number of GPUs

**How NCCL optimize it?**

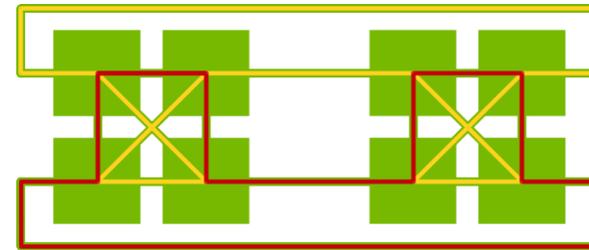
# DESIGN

## Rings

NCCL uses **rings** to move data across all GPUs.



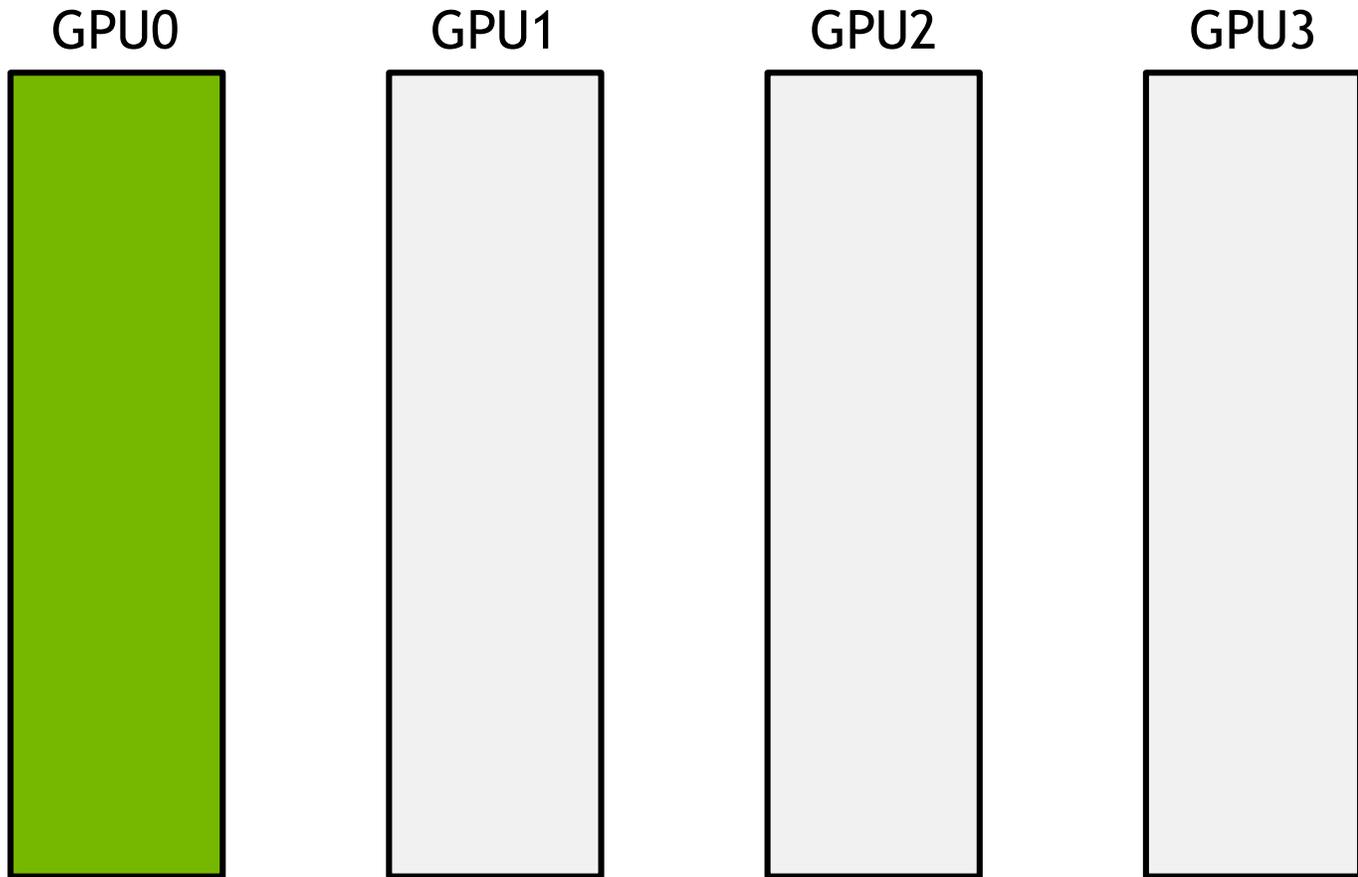
PCIe / QPI : 1 unidirectional ring



DGX-1 : 4 unidirectional rings

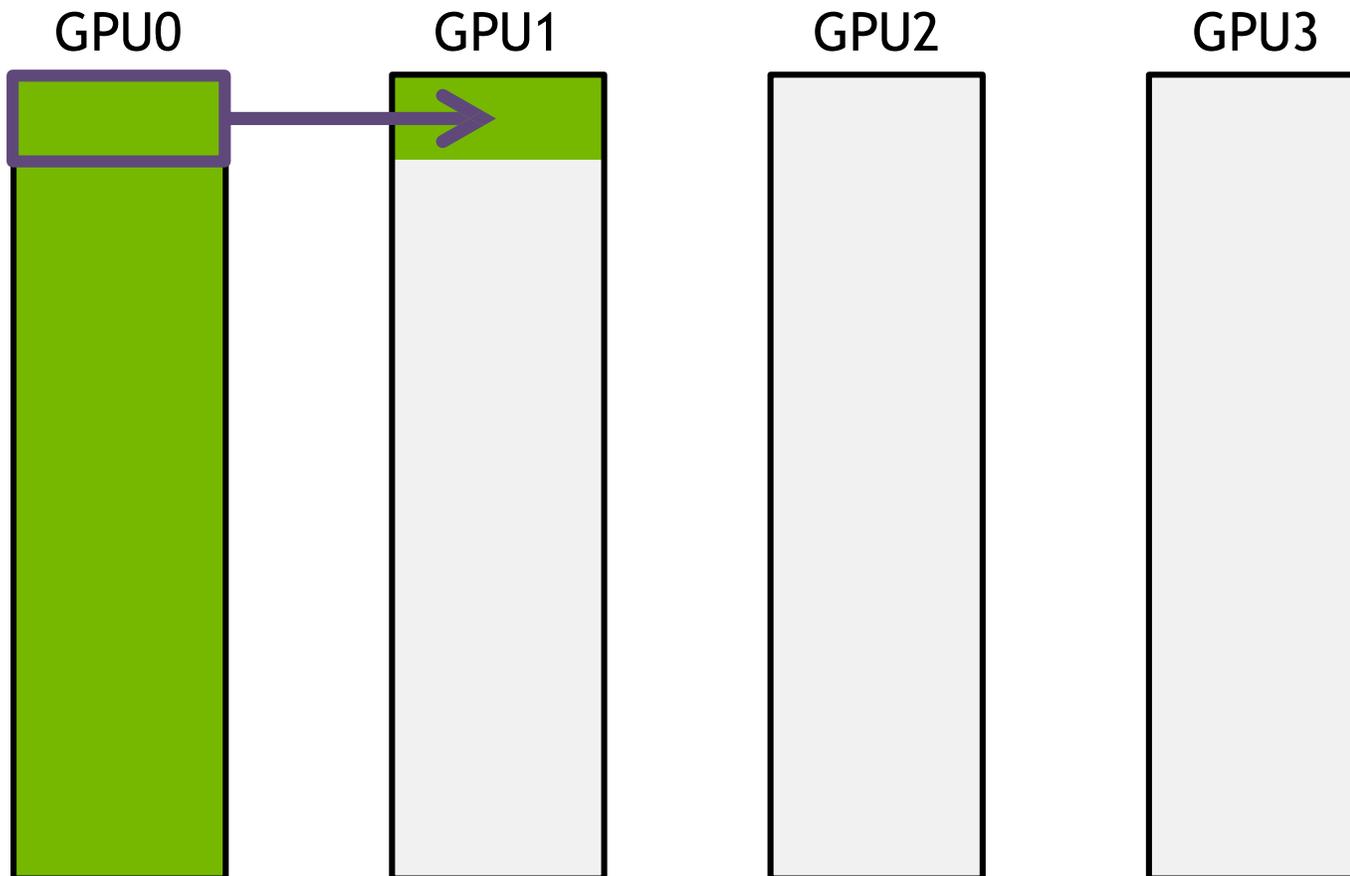
# BROADCAST

with unidirectional ring



# BROADCAST

with unidirectional ring

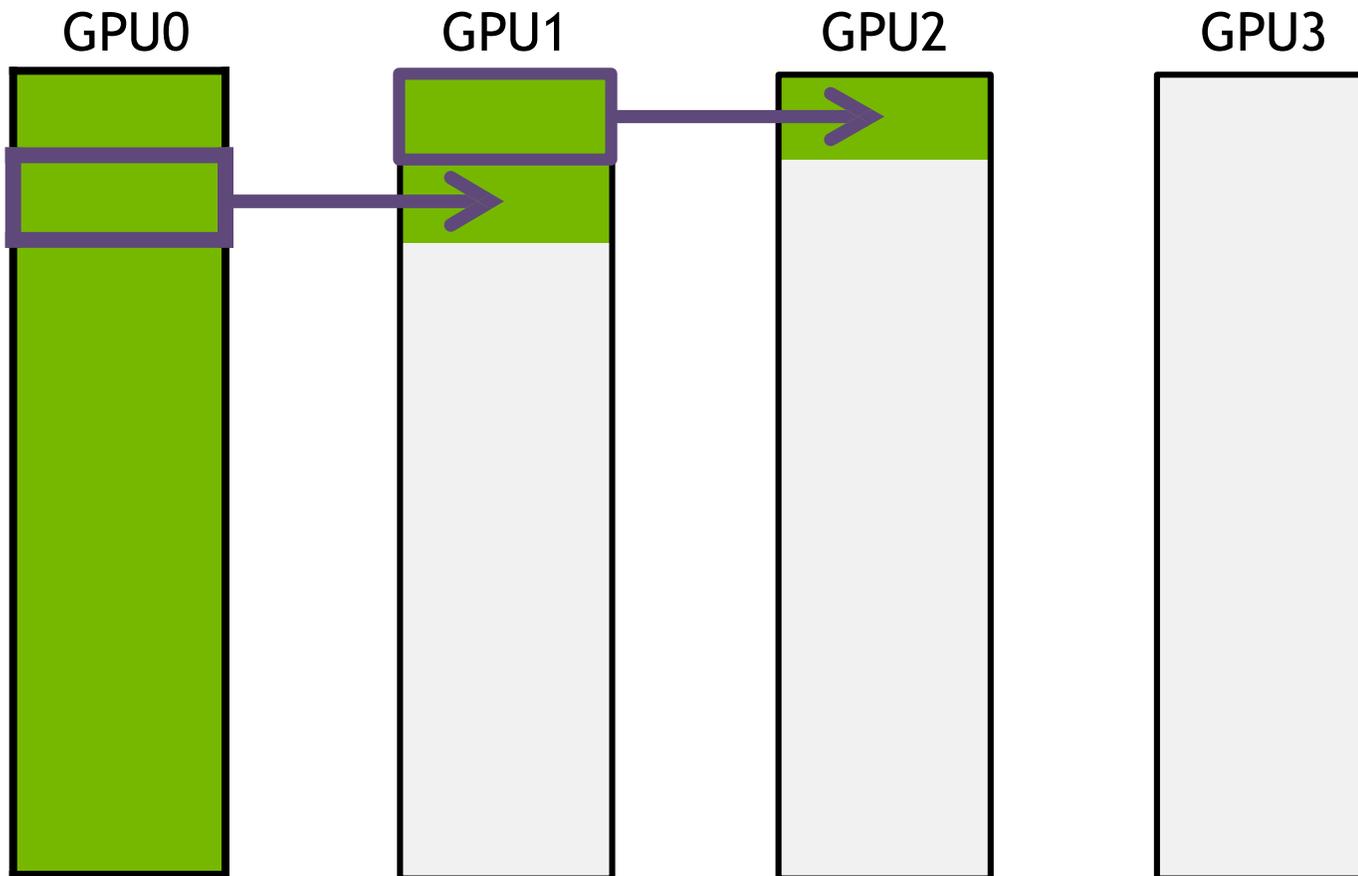


Split data into  $S$  messages

Step 1:  $\Delta t = N/(SB)$

# BROADCAST

with unidirectional ring



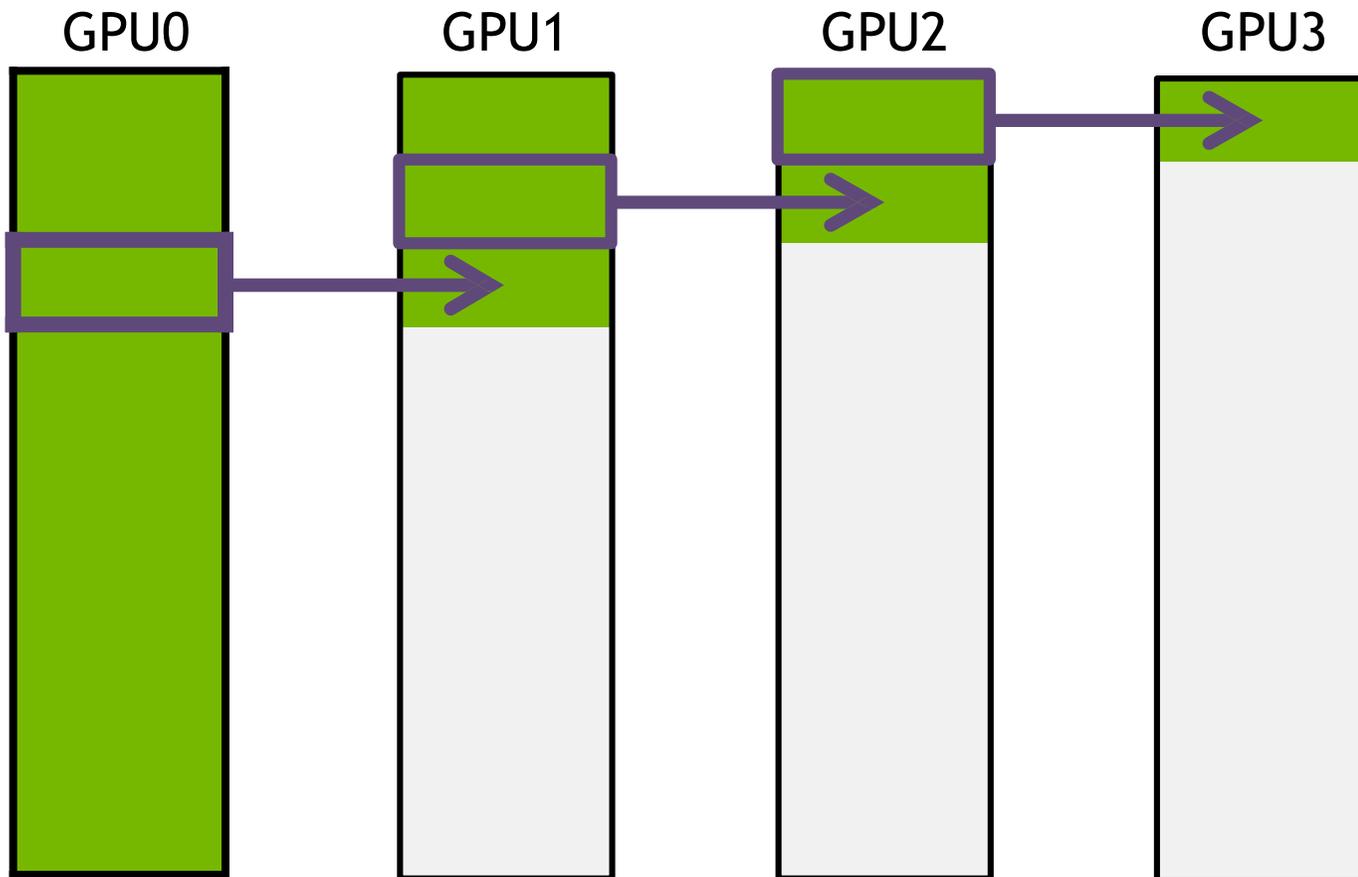
Split data into  $S$  messages

Step 1:  $\Delta t = N/(SB)$

Step 2:  $\Delta t = N/(SB)$

# BROADCAST

with unidirectional ring



Split data into  $S$  messages

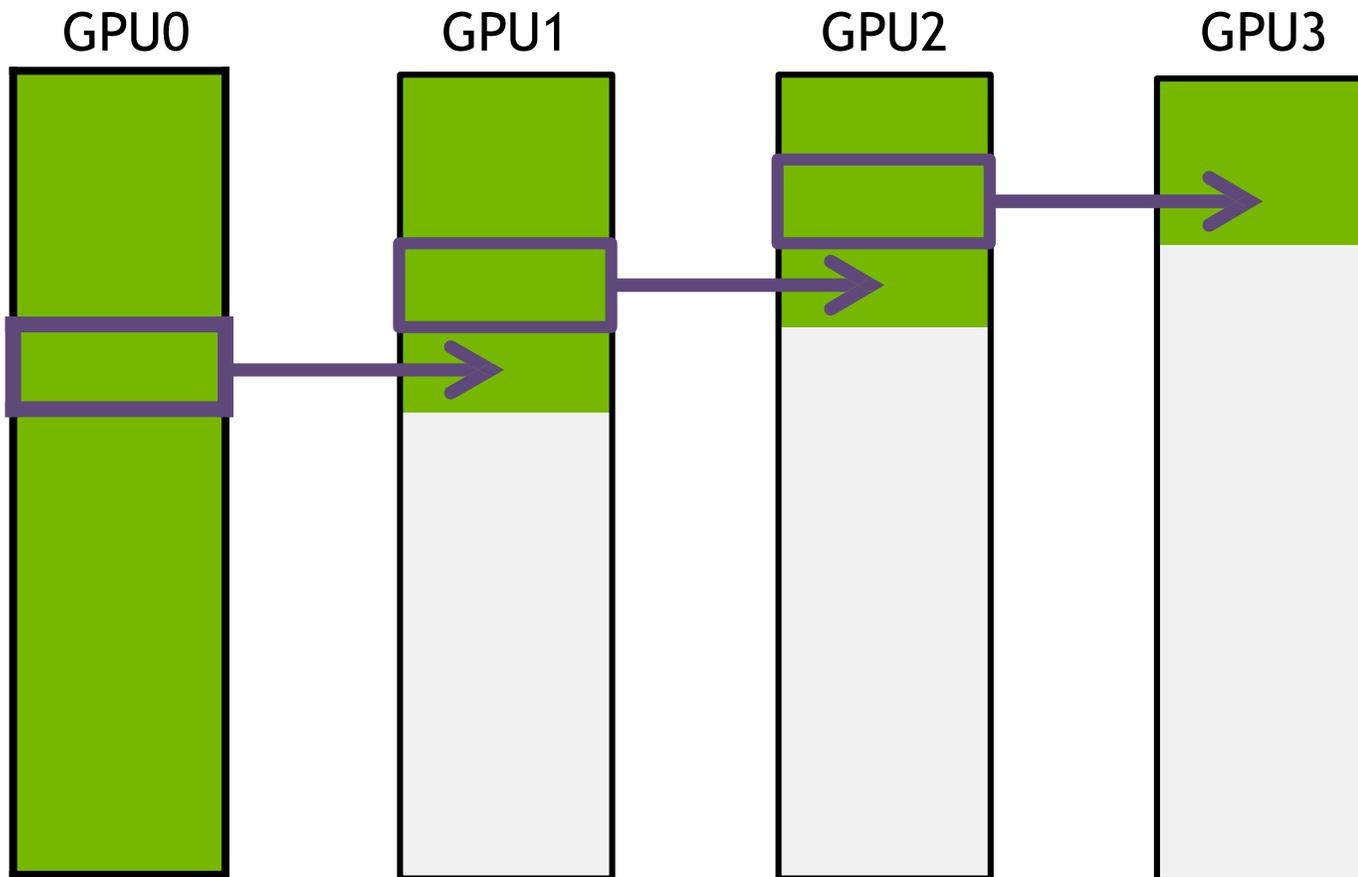
Step 1:  $\Delta t = N/(SB)$

Step 2:  $\Delta t = N/(SB)$

Step 3:  $\Delta t = N/(SB)$

# BROADCAST

with unidirectional ring



Split data into  $S$  messages

Step 1:  $\Delta t = N/(SB)$

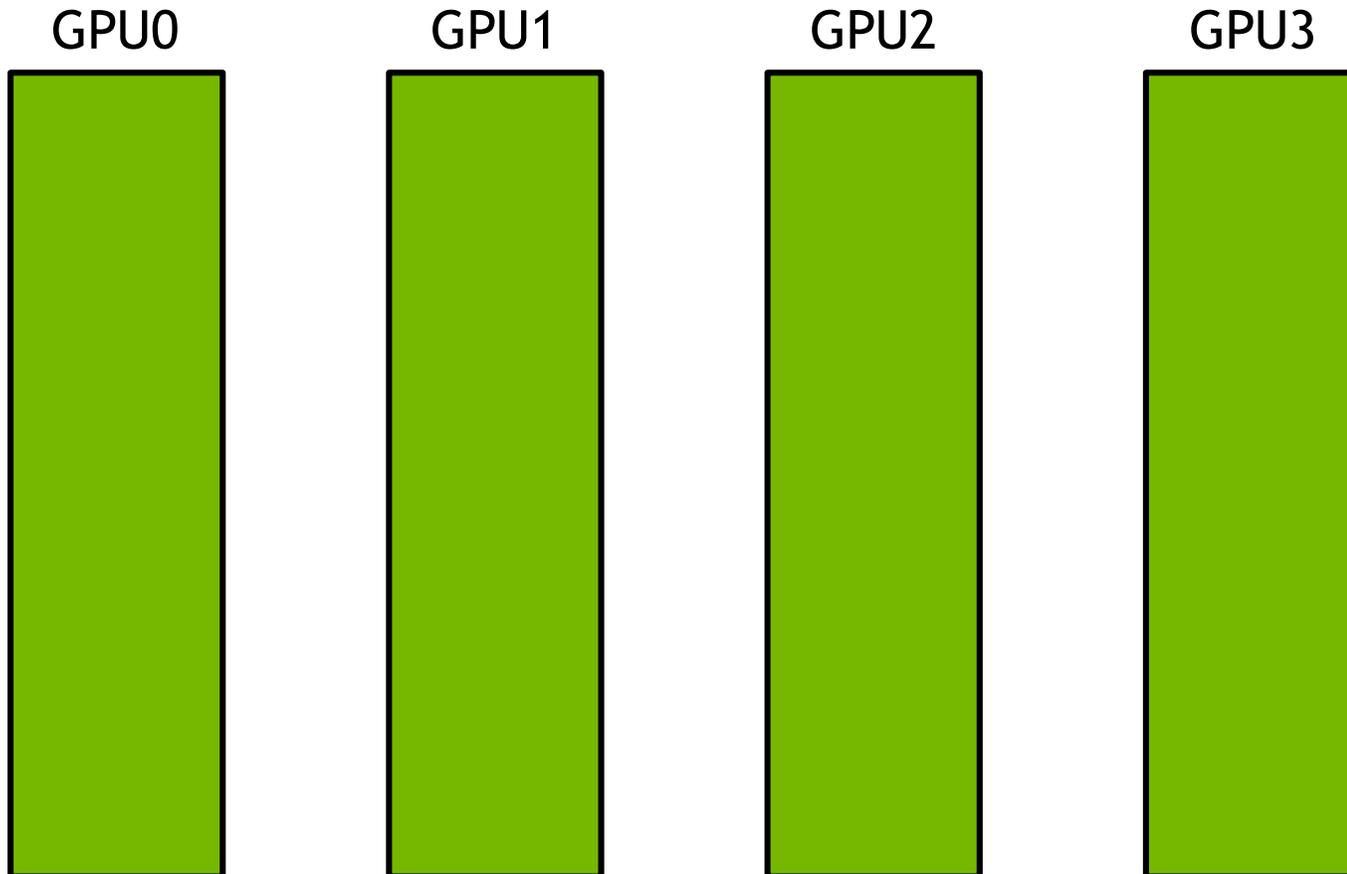
Step 2:  $\Delta t = N/(SB)$

Step 3:  $\Delta t = N/(SB)$

Step 4:  $\Delta t = N/(SB)$

# BROADCAST

with unidirectional ring



Split data into  $S$  messages

Step 1:  $\Delta t = N/(SB)$

Step 2:  $\Delta t = N/(SB)$

Step 3:  $\Delta t = N/(SB)$

Step 4:  $\Delta t = N/(SB)$

...

Total time:

$$SN/(SB) + (k - 2) N/(SB)$$

$$= N(S + k - 2)/(SB) \rightarrow N/B$$

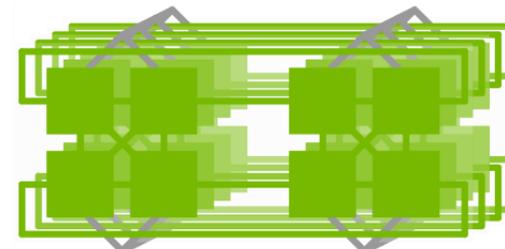
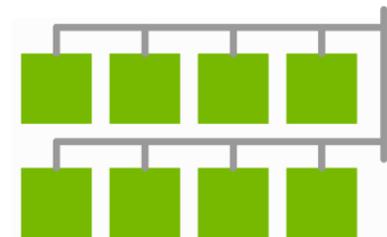
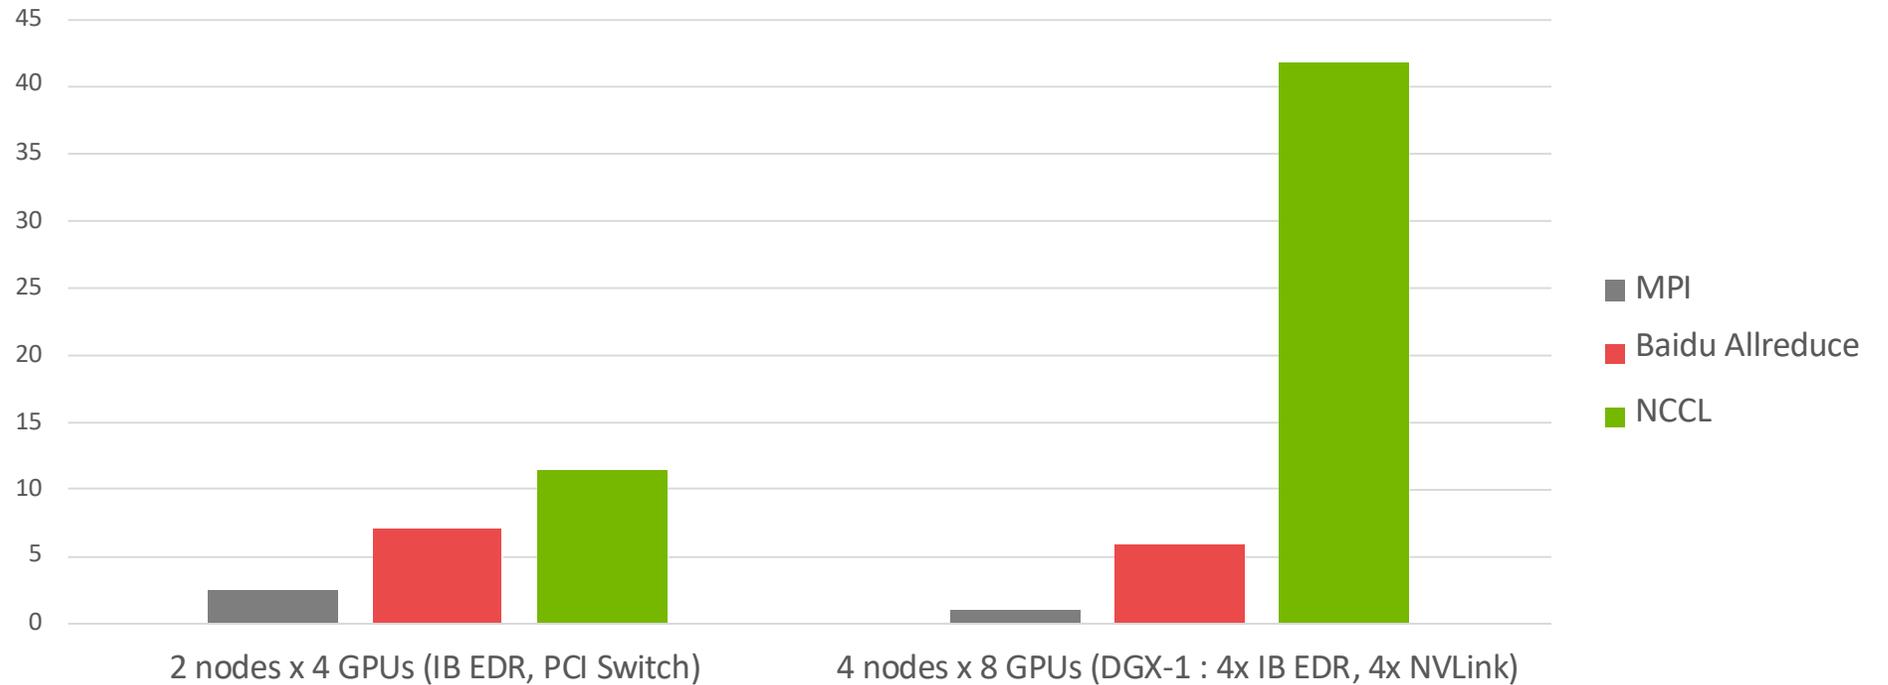
# NCCL

- Automatic topology detection.
  - Better utilization of NVLink topology
- Automatic algorithm selection and optimization.
  - Ring, Multi-ring.
  - Tree.

# PERFORMANCE

## Inter-node performance

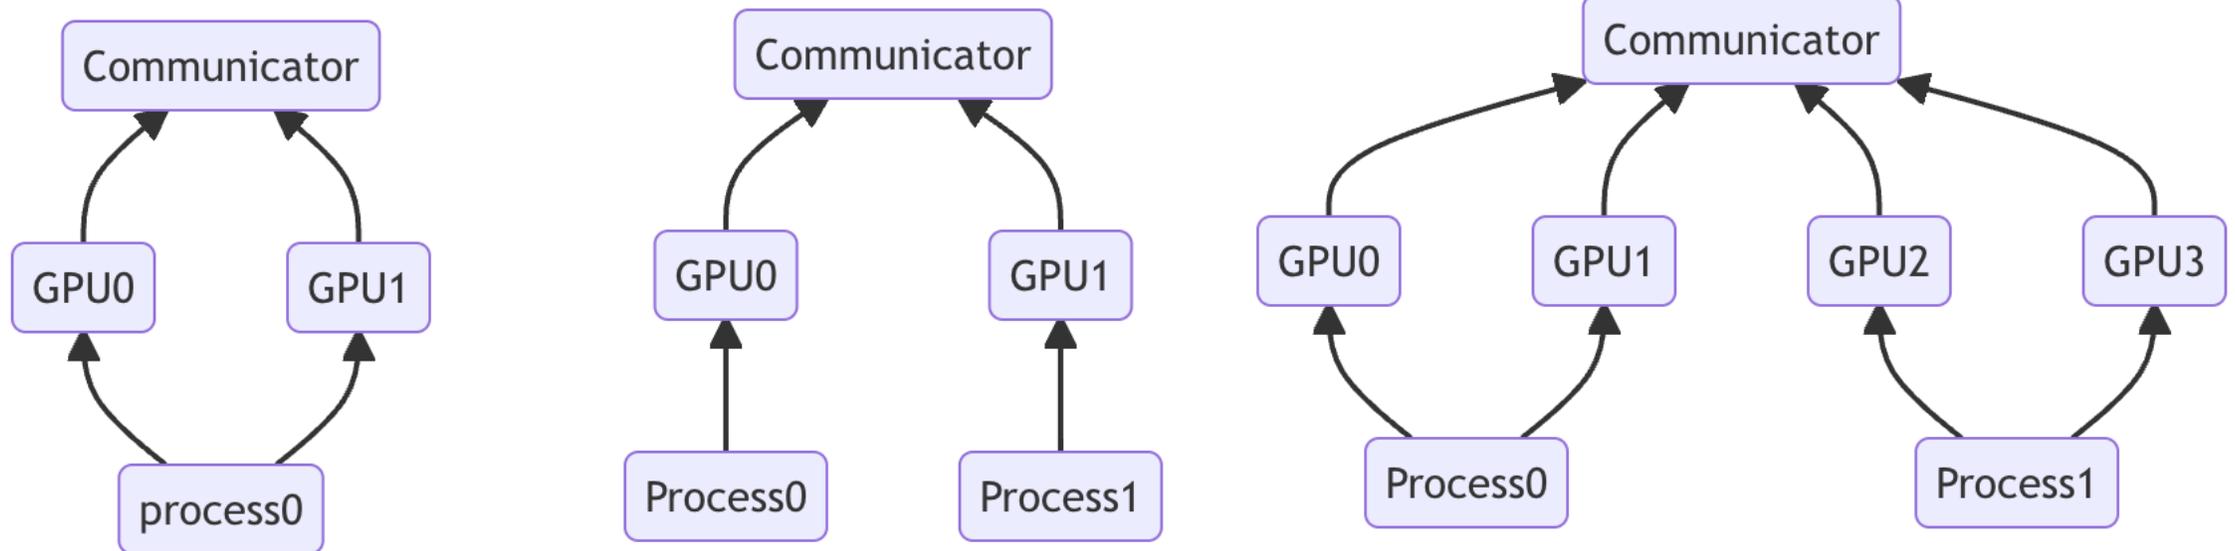
AllReduce bandwidth (OMB, size=128MB, in GB/s)



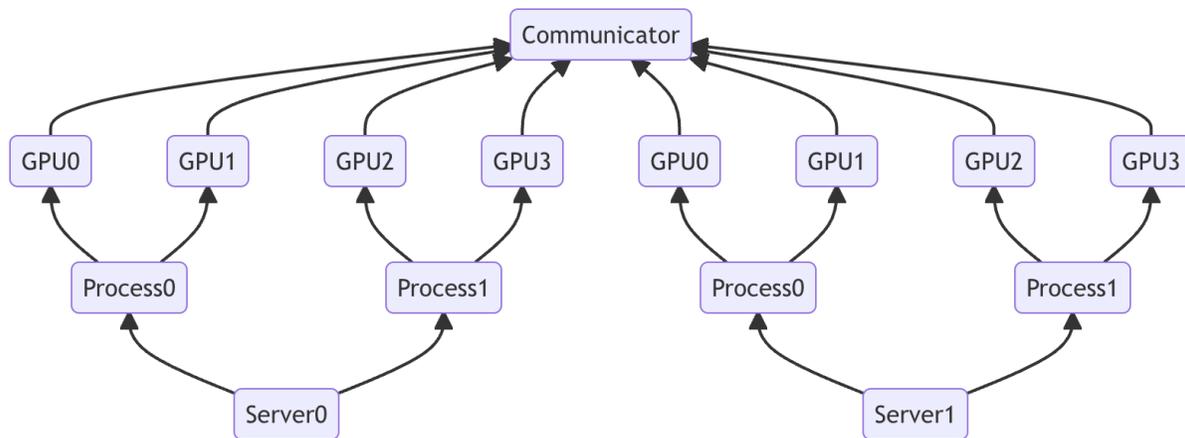
**How to use NCCL?**

# Common Cases

1. Single process/thread multiple devices.
2. One Device per process/thread.
3. Multiple devices per process/thread.



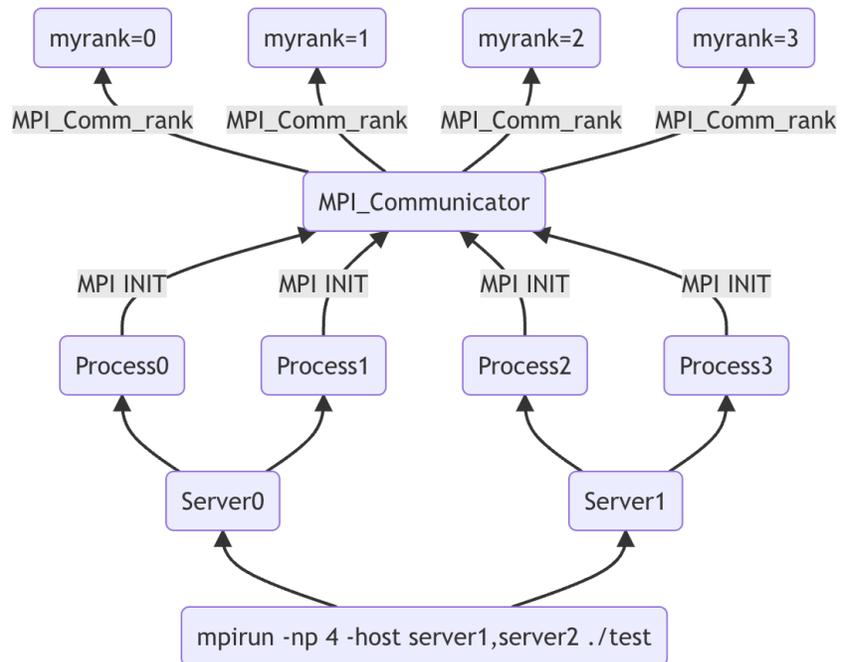
# M servers, M process each with M devices



► Source Code: Multiple Devices Per Thread Example.

	command	notes
Compile	<code>g++ -o test test.cpp -lncc1 -lcudart -lmpi</code>	
Run 1 host	<code>mpirun -np n ./test</code>	run n process, need 2nGPU
Run M host	<code>mpirun -np n -host server1, server2 ./test</code>	need ssh config

# Step 1: Initialize MPI

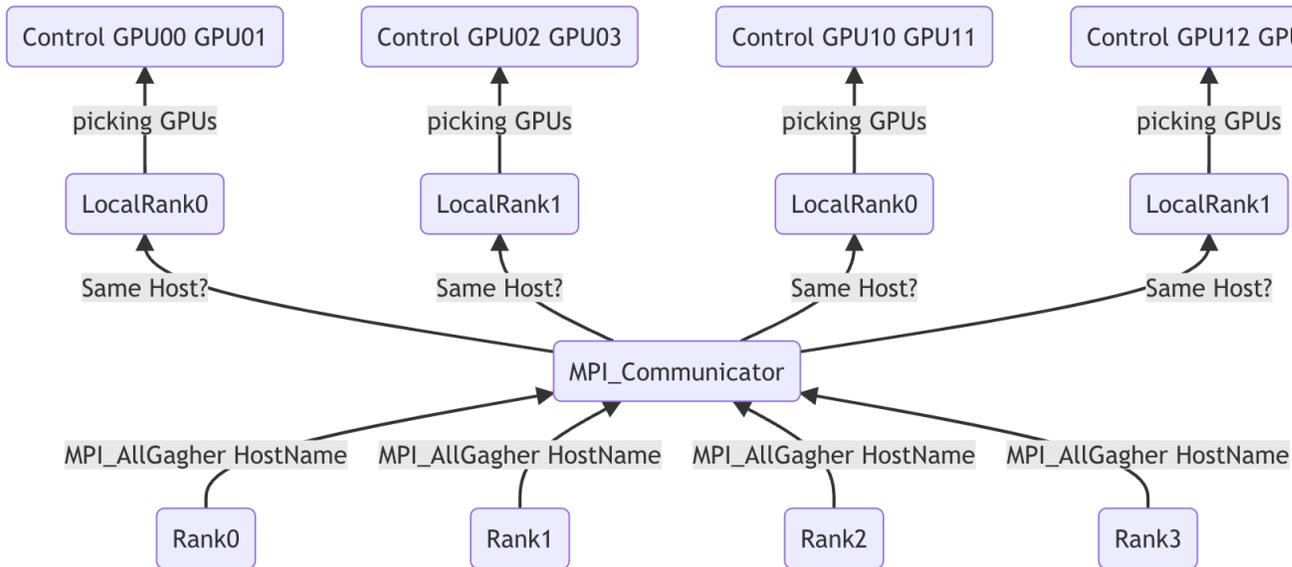


- Local process gets its rank of global.
- Local process gets total process size.

1. `MPI_Init` : Init MPI environment and core datastructure (communicator).
2. `MPI_COMM_WORLD` : Return the address of MPI core datastructure.
3. `MPI_Comm_rank` : Get the current process ID rank in `MPI_COMM_WORLD`.
4. `MPI_Comm_size` : Get the size of the communicator in `MPI_COMM_WORLD`.

```
1 int myRank, nRanks, localRank = 0;
2
3 // initializing MPI
4 MPICHECK(MPI_Init(&argc, &argv));
5 MPICHECK(MPI_Comm_rank(MPI_COMM_WORLD, &myRank));
6 MPICHECK(MPI_Comm_size(MPI_COMM_WORLD, &nRanks));
```

# Step2: Get Local Rank

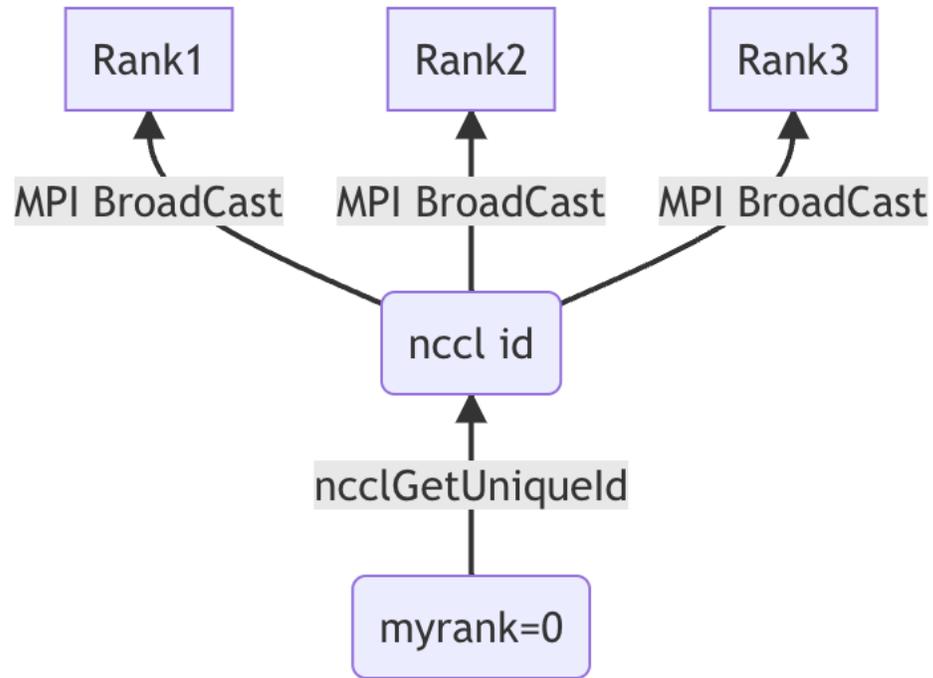


- Local process uses MPI All\_gather to get global hostHashs.
- Local process gets localRank of the same host.

1. `MPI_Allgather` : Gather data from all processes.
2. `MPI_IN_PLACE` : Use input buffer as output buffer.
3. `MPI_DATATYPE_NULL` : No datatype specified.
4. `MPI_BYTE` : Byte-sized data type.

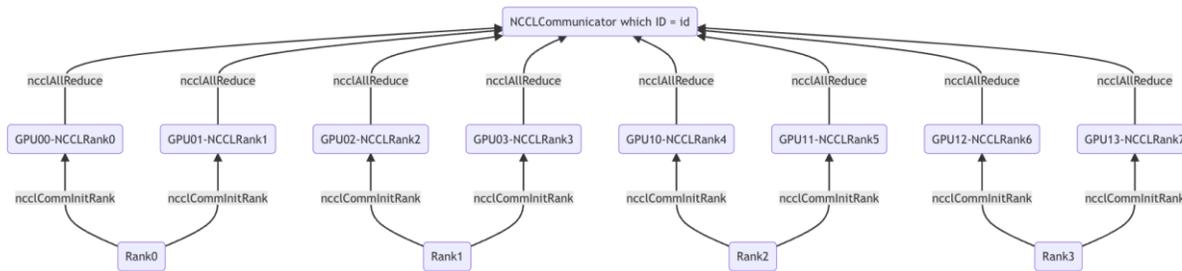
```
1  int size = 32*1024*1024;
2
3  int myRank, nRanks, localRank = 0;
4
5  // calculating localRank based on hostname which is used in selecting a GPU
6  uint64_t hostHashs[nRanks];
7  char hostname[1024];
8  getHostName(hostname, 1024);
9  hostHashs[myRank] = getHostHash(hostname);
10 MPICHECK(MPI_Allgather(MPI_IN_PLACE, 0, MPI_DATATYPE_NULL, hostHashs,
11                       sizeof(uint64_t), MPI_BYTE, MPI_COMM_WORLD));
12 for (int p = 0; p < nRanks; p++) {
13     if (p == myRank) break;
14     if (hostHashs[p] == hostHashs[myRank]) localRank++;
15 }
```

# Step3: Broadcast NCCL ID



```
1 //each process is using two GPUs
2 int nDev = 2;
3
4 float** sendbuff = (float**)malloc(nDev * sizeof(float*));
5 float** recvbuff = (float**)malloc(nDev * sizeof(float*));
6 cudaStream_t* s = (cudaStream_t*)malloc(sizeof(cudaStream_t)*nDev);
7
8 //picking GPUs based on localRank
9 for (int i = 0; i < nDev; ++i) {
10     CUDACHECK(cudaSetDevice(localRank*nDev + i));
11     CUDACHECK(cudaMalloc(sendbuff + i, size * sizeof(float)));
12     CUDACHECK(cudaMalloc(recvbuff + i, size * sizeof(float)));
13     CUDACHECK(cudaMemset(sendbuff[i], 1, size * sizeof(float)));
14     CUDACHECK(cudaMemset(recvbuff[i], 0, size * sizeof(float)));
15     CUDACHECK(cudaStreamCreate(s+i));
16 }
17
18 nccUniqueId id;
19 nccComm_t comms[nDev];
20
21 //generating NCCL unique ID at one process and broadcasting it to all
22 if (myRank == 0) nccGetUniqueId(&id);
23 MPICHECK(MPI_Bcast((void *)&id, sizeof(id), MPI_BYTE, 0, MPI_COMM_WORLD));
```

# Step4: NCCL Communicate



1. `ncclGroupStart` : Begin grouping NCCL calls to perform collective operations on multiple devices.
2. `ncclAllReduce` : Prepare an all-reduce collective operation, store in stream.
3. `ncclGroupEnd` : End the grouping of NCCL calls. Trigger for these operations to commence.
4. `cudaStreamSynchronize` : Synchronize on CUDA streams to ensure completion of NCCL operations.

```
1 //initializing NCCL, group API is required around ncclCommInitRank as it is
2 //called across multiple GPUs in each thread/process
3 NCCLCHECK(ncclGroupStart());
4 for (int i=0; i<nDev; i++) {
5     CUDAHECK(cudaSetDevice(localRank*nDev + i));
6     NCCLCHECK(ncclCommInitRank(comms+i, nRanks*nDev, id, myRank*nDev + i));
7 }
8 NCCLCHECK(ncclGroupEnd());
9
10 //calling NCCL communication API. Group API is required when using
11 //multiple devices per thread/process
12 NCCLCHECK(ncclGroupStart());
13 for (int i=0; i<nDev; i++)
14     NCCLCHECK(ncclAllReduce((const void*)sendbuff[i], (void*)recvbuff[i], size, ncclFloat, n
15         comms[i], s[i]));
16 NCCLCHECK(ncclGroupEnd());
```

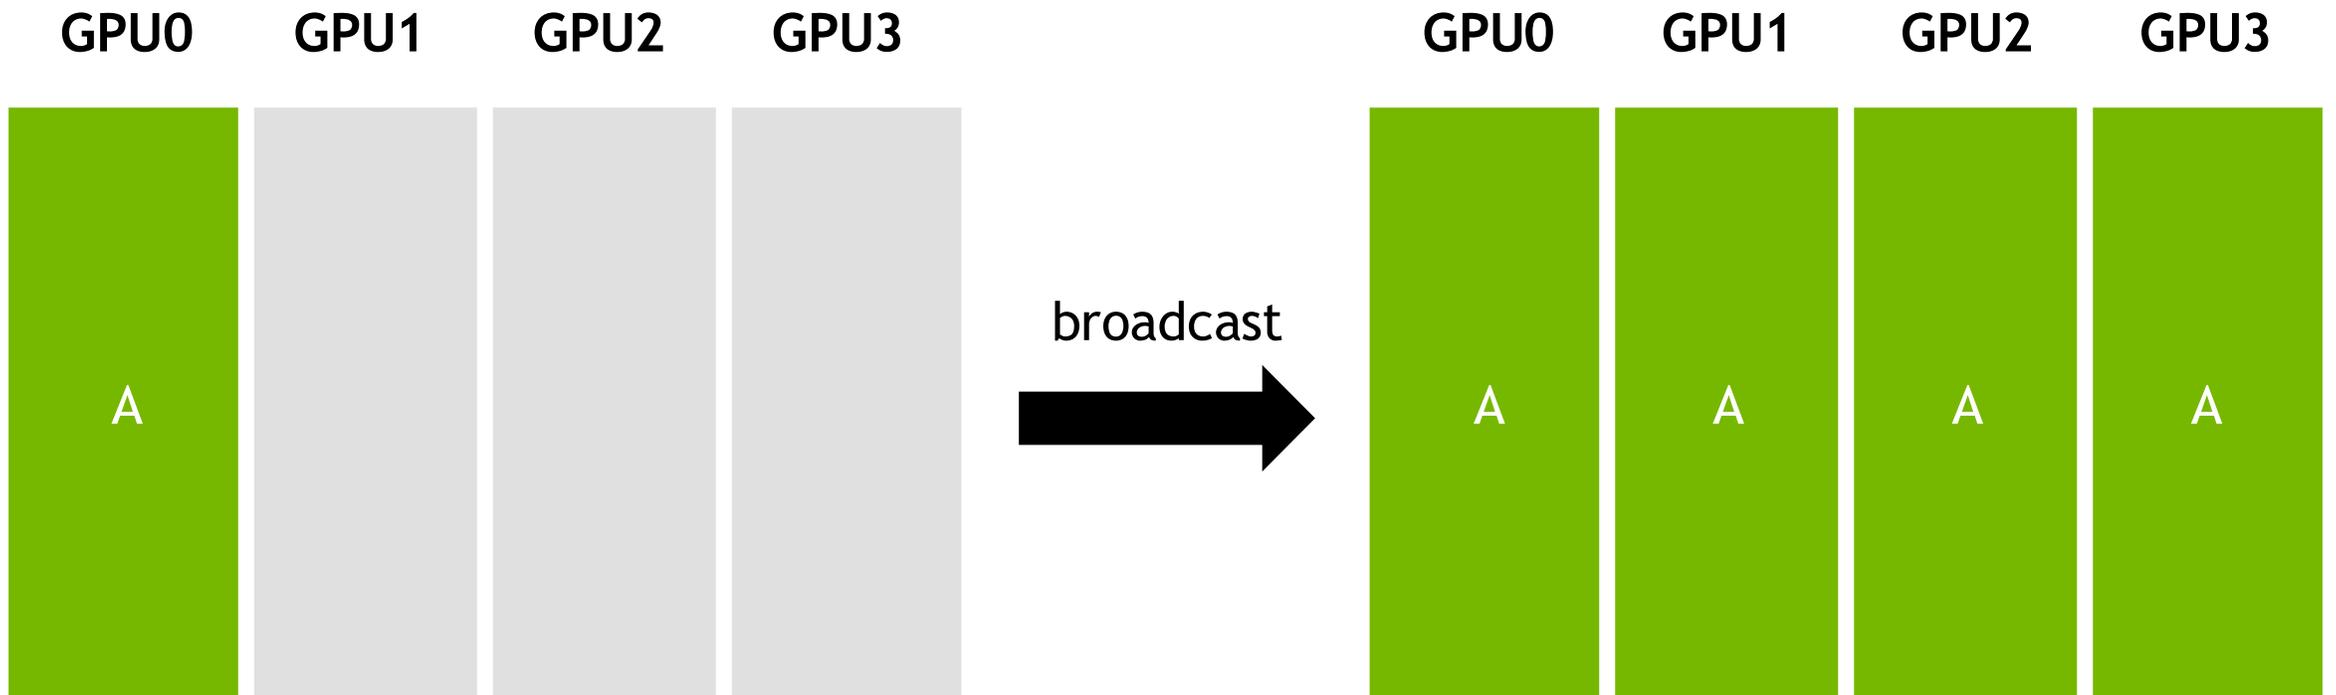


**Thanks!**

**Backup**

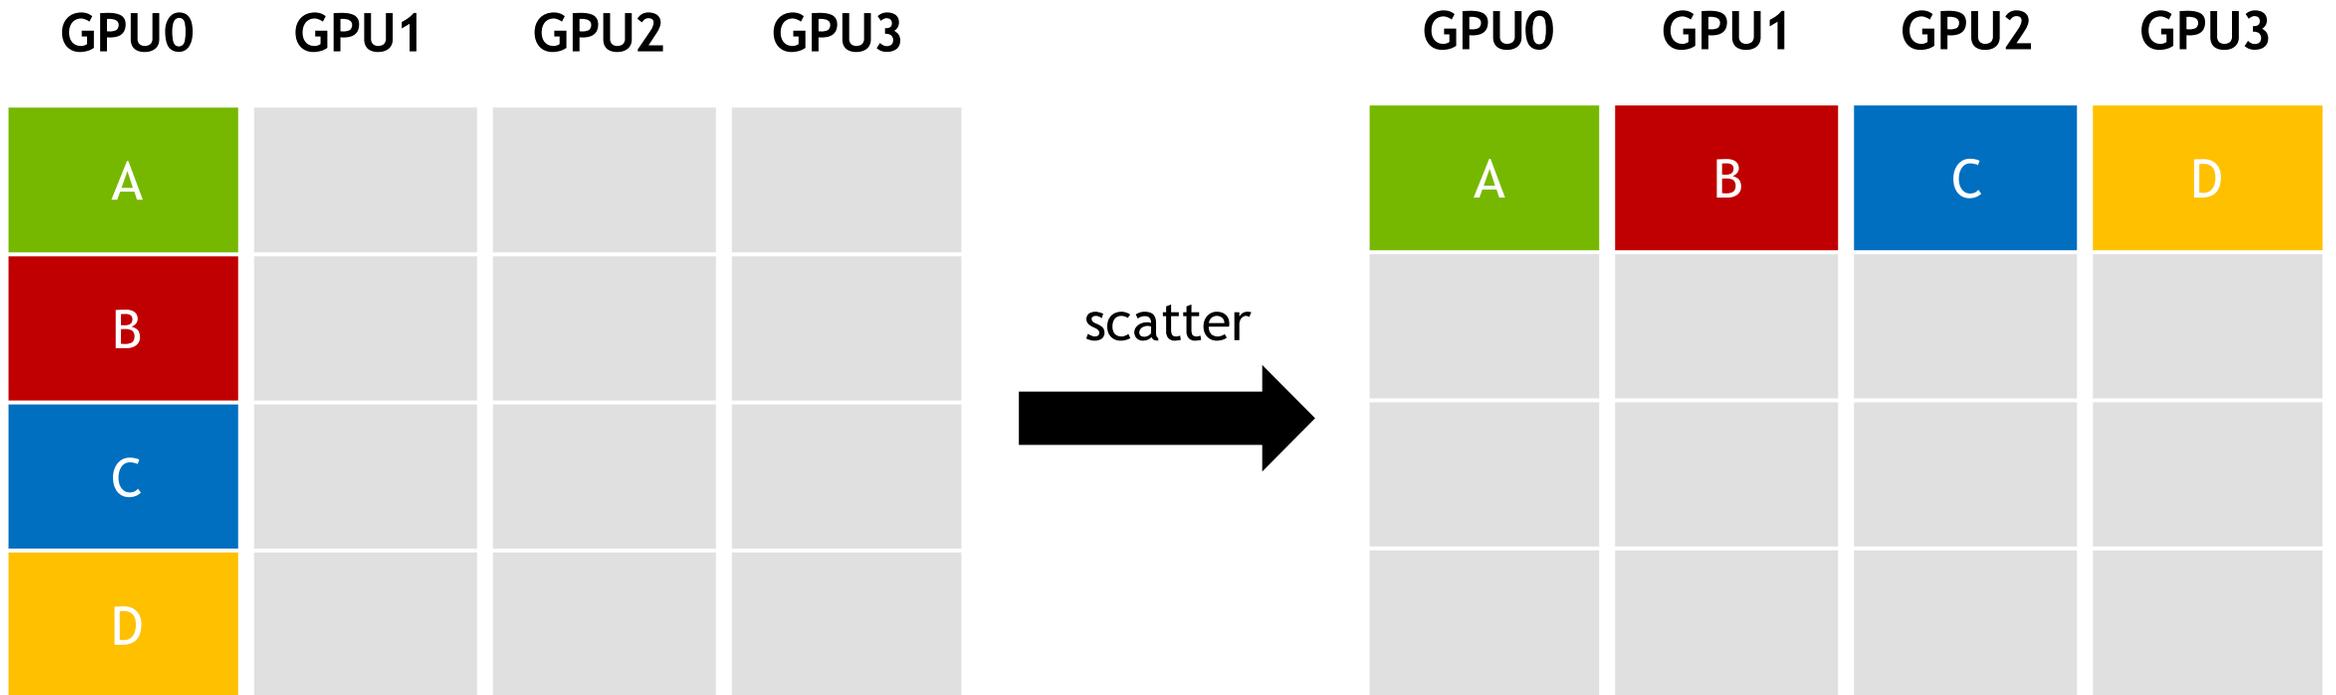
# BROADCAST

One sender, multiple receivers



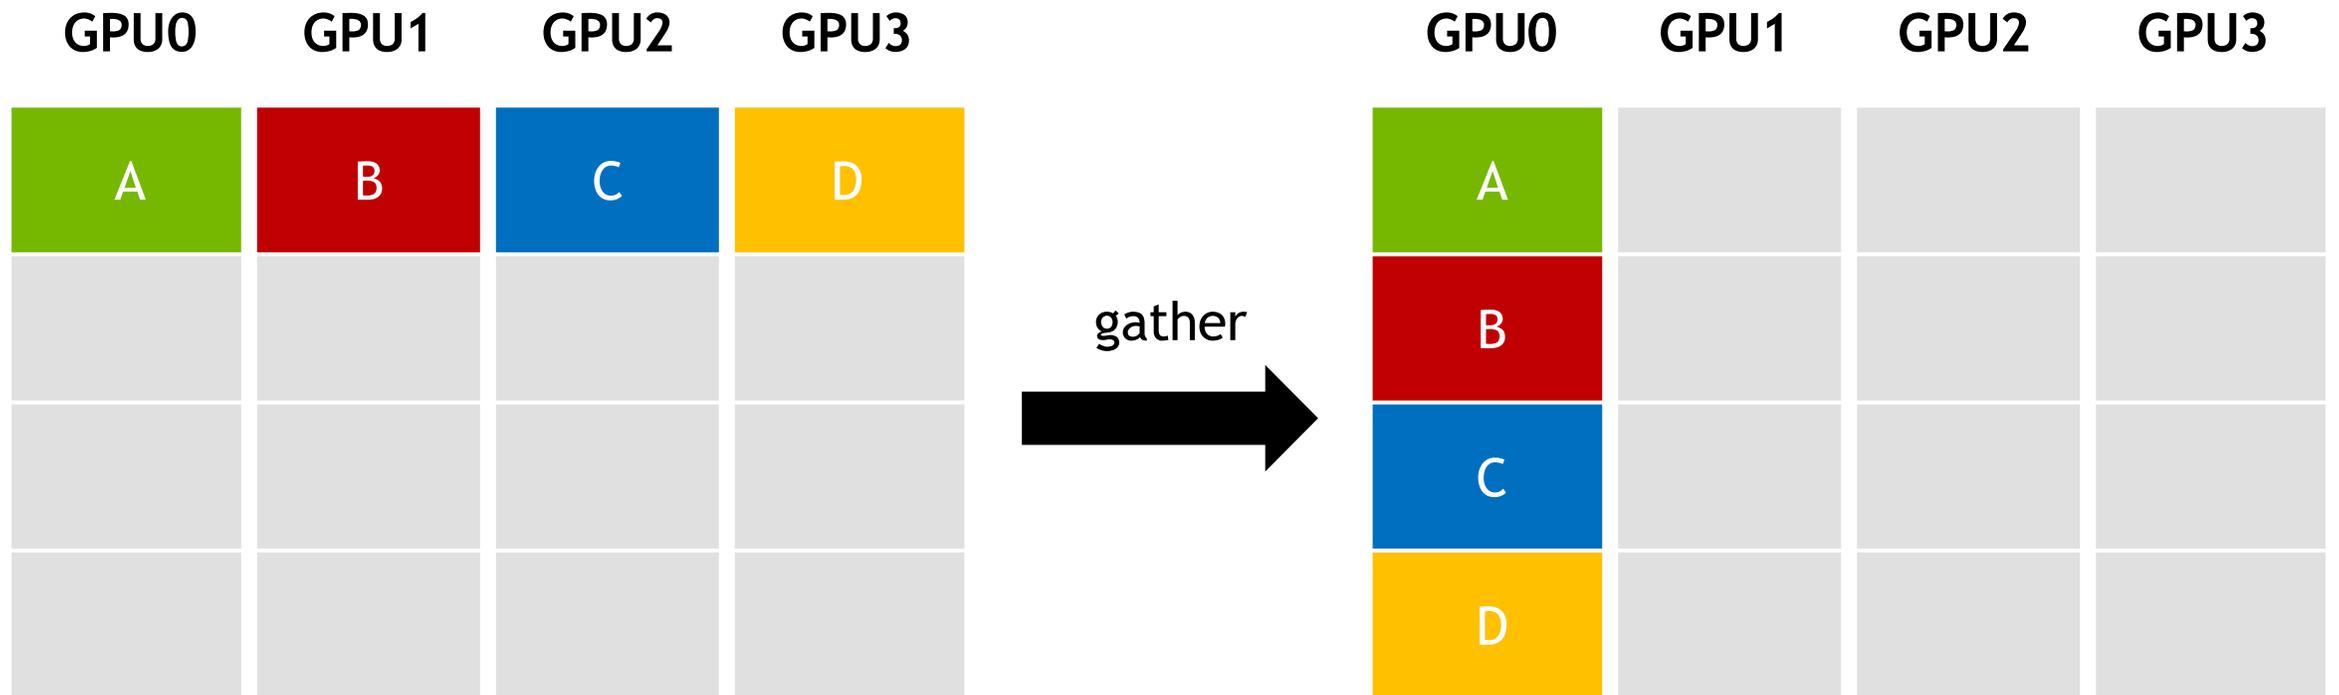
# SCATTER

One sender; data is distributed among multiple receivers



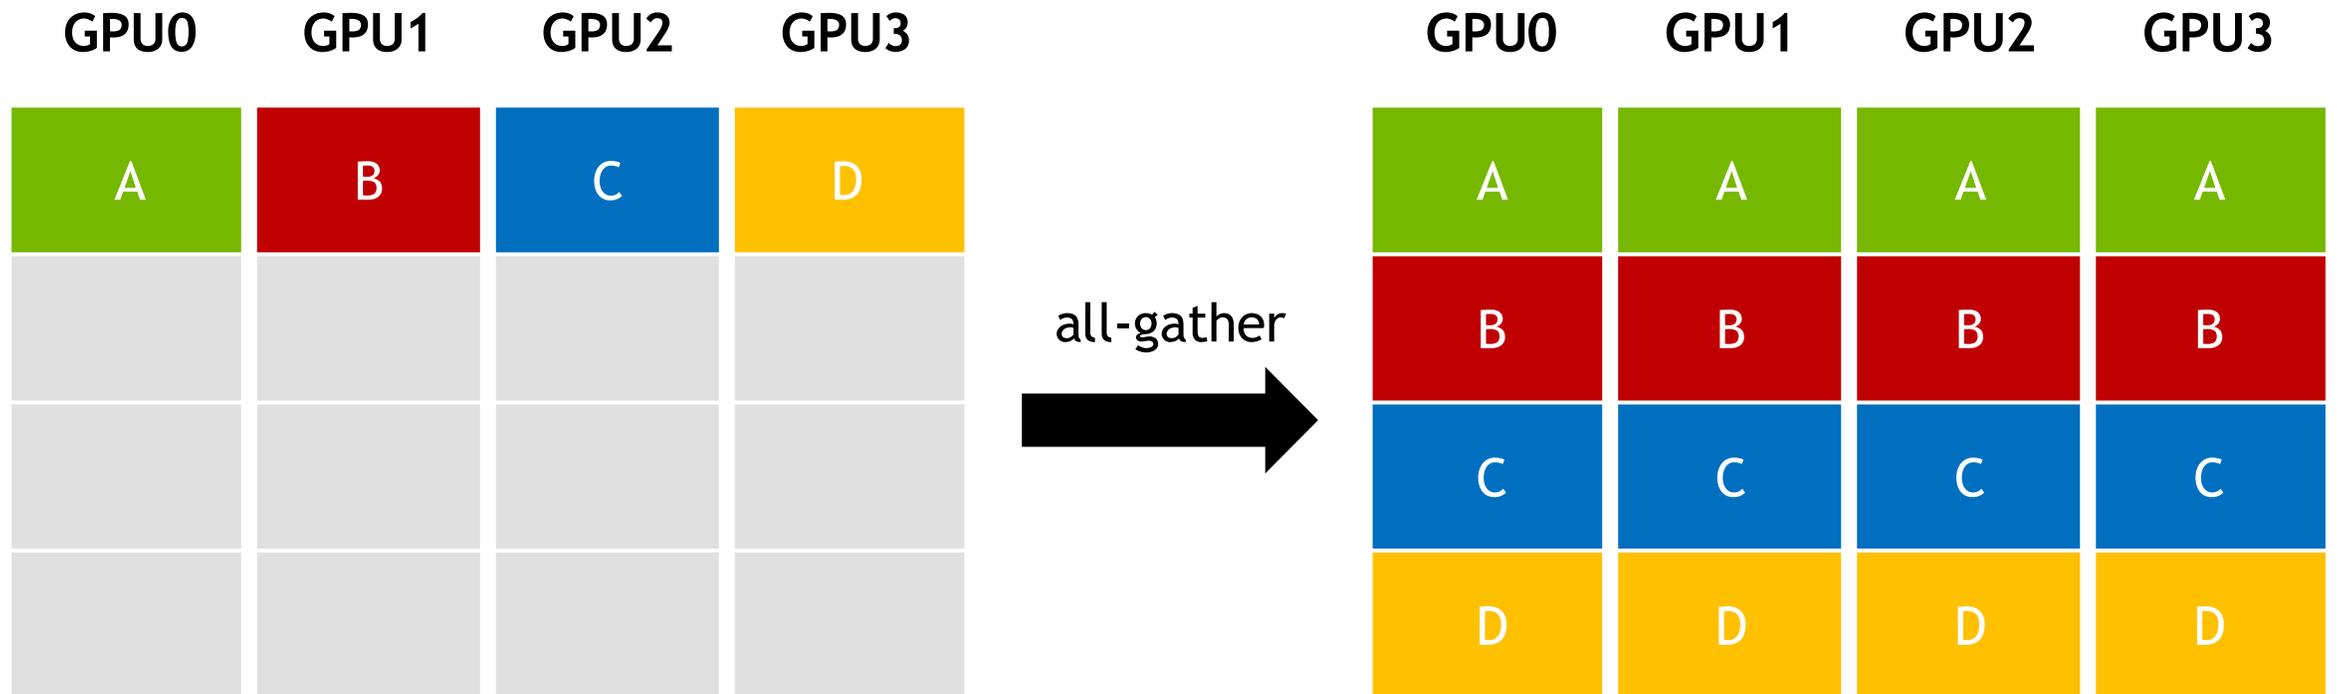
# GATHER

Multiple senders, one receiver



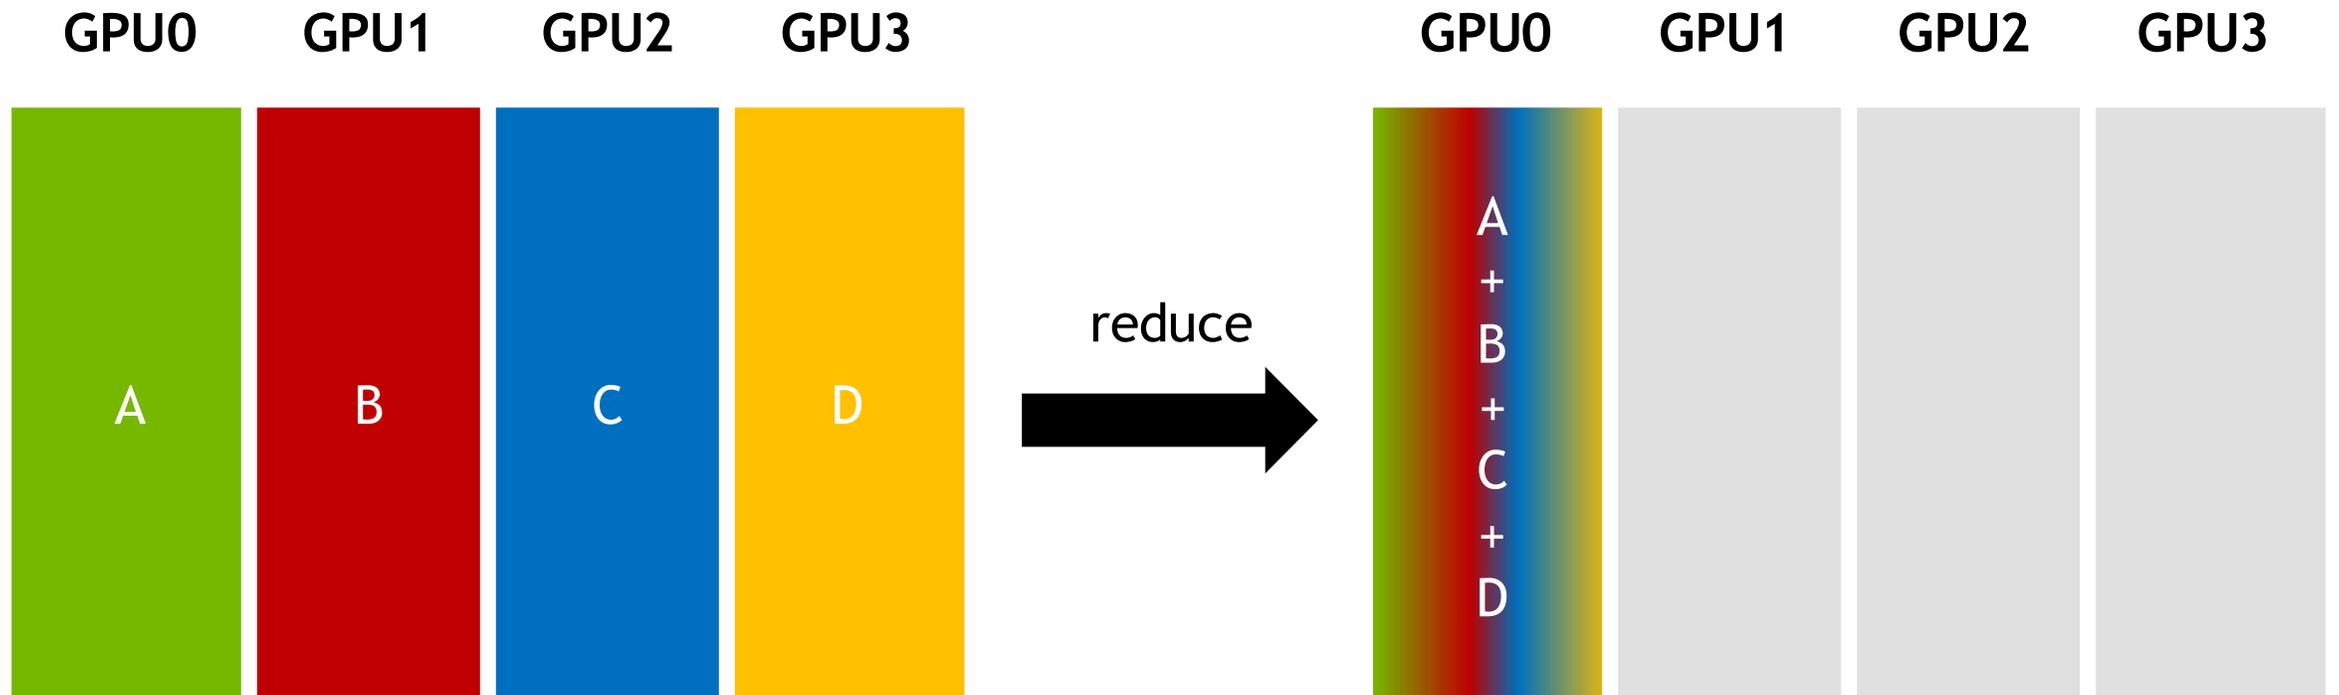
# ALL-GATHER

Gather messages from all; deliver gathered data to all participants



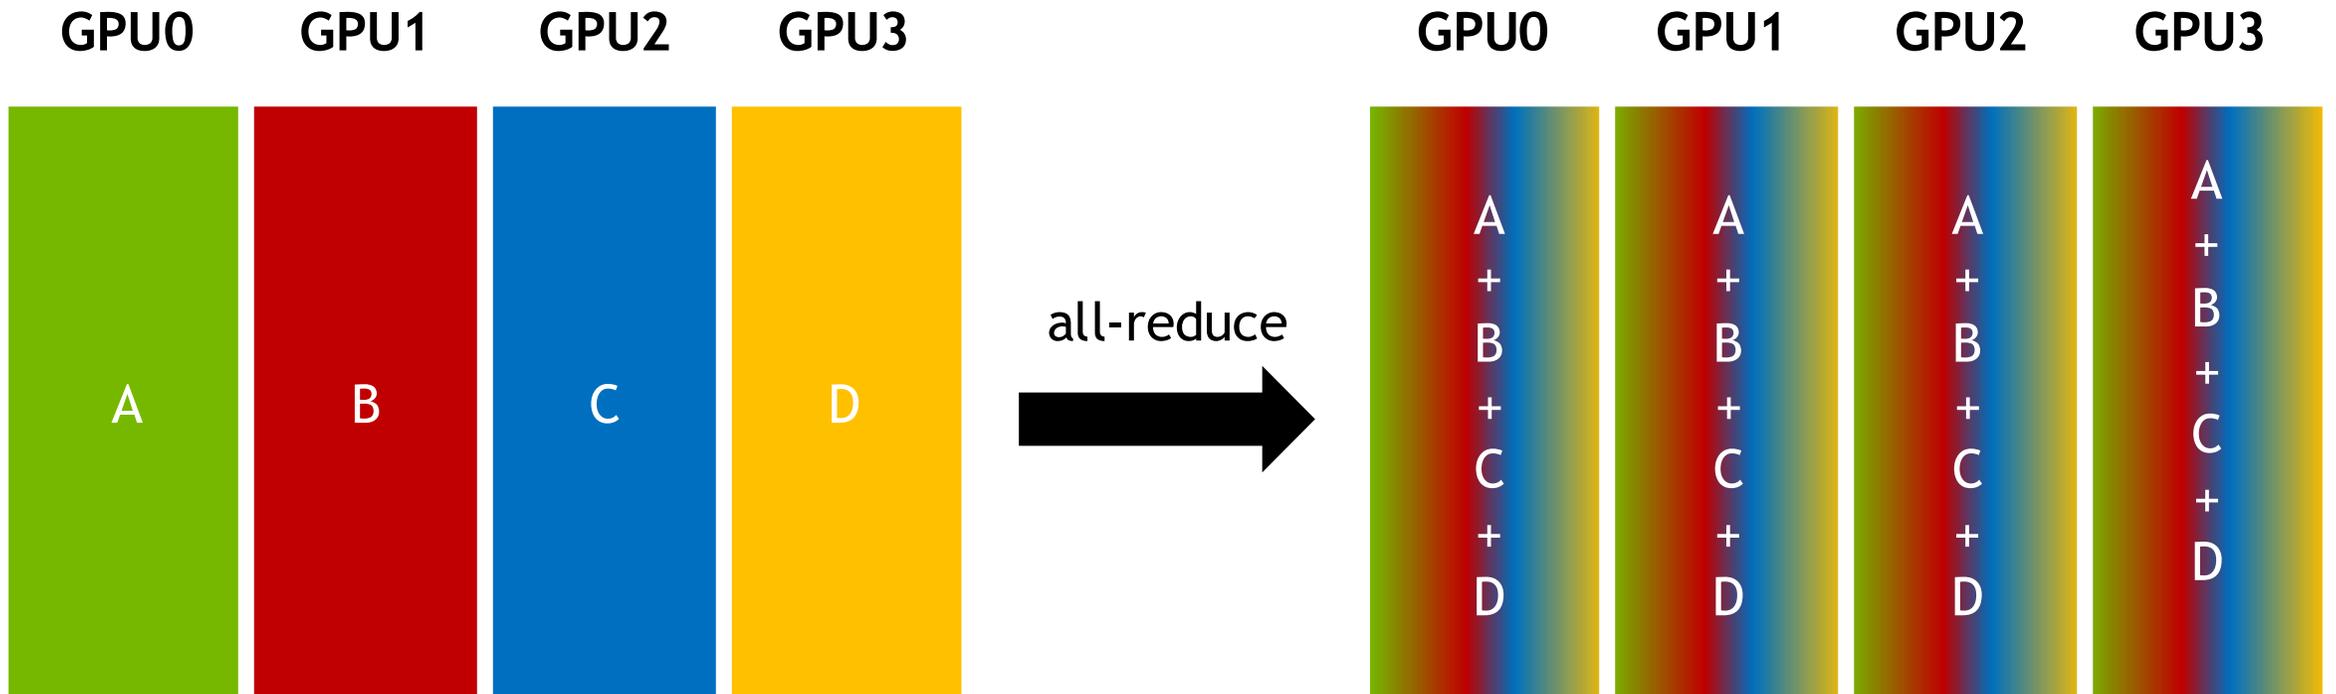
# REDUCE

Combine data from all senders; deliver the result to one receiver



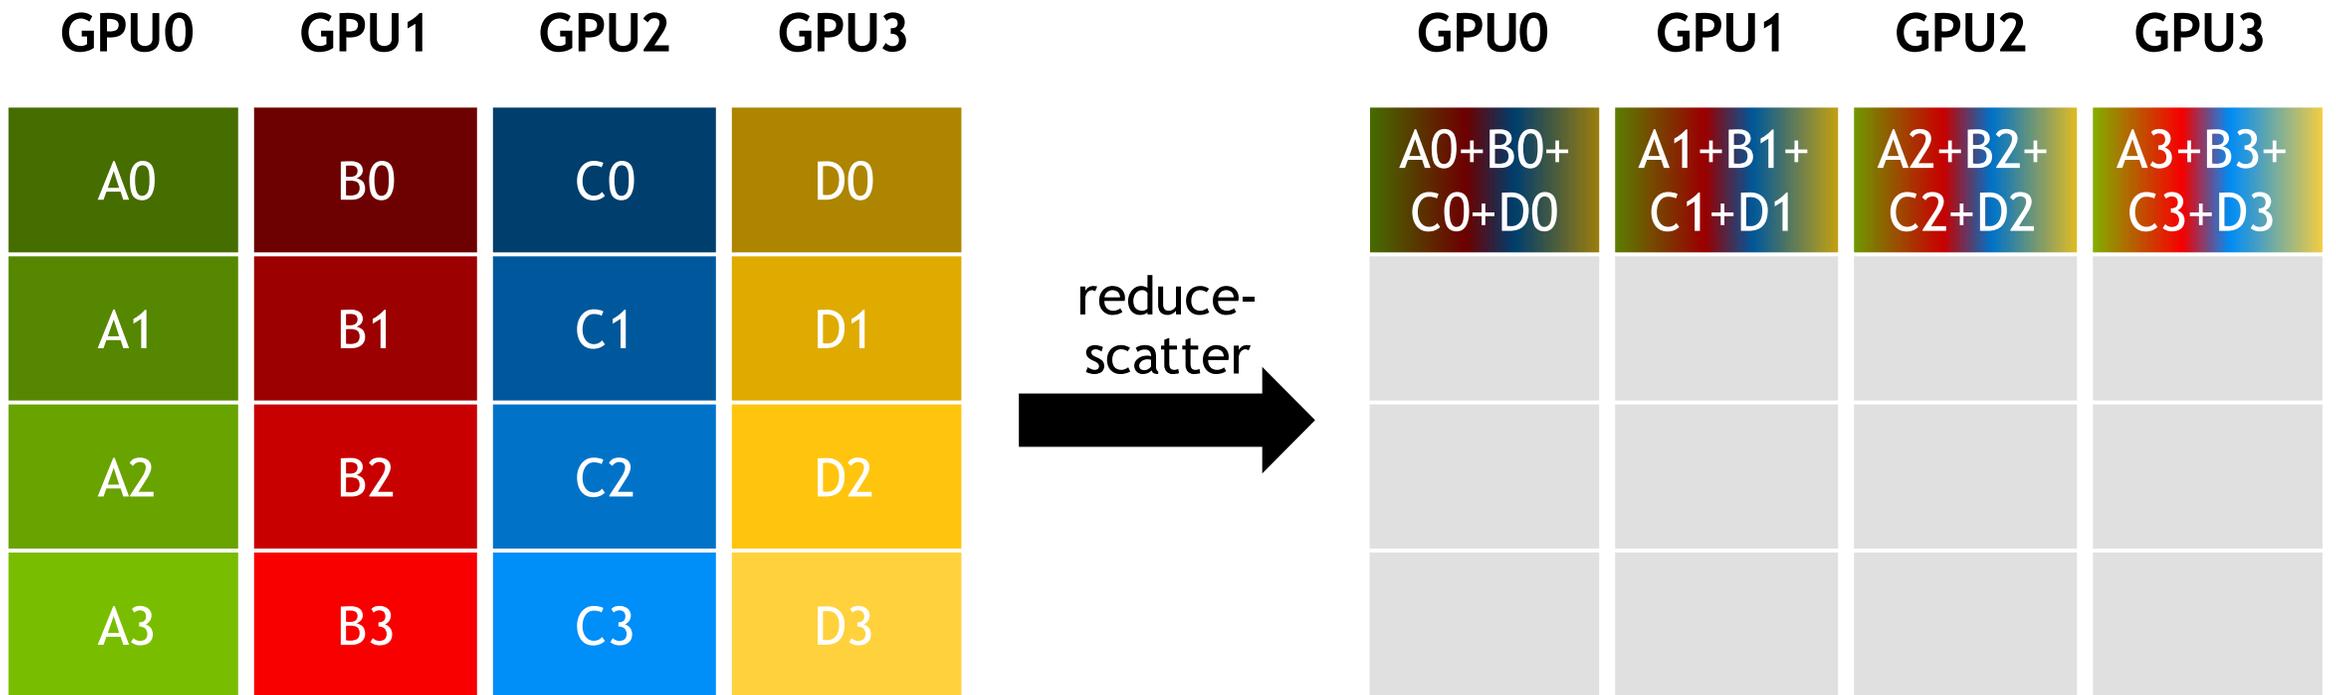
# ALL-REDUCE

Combine data from all senders; deliver the result to all participants



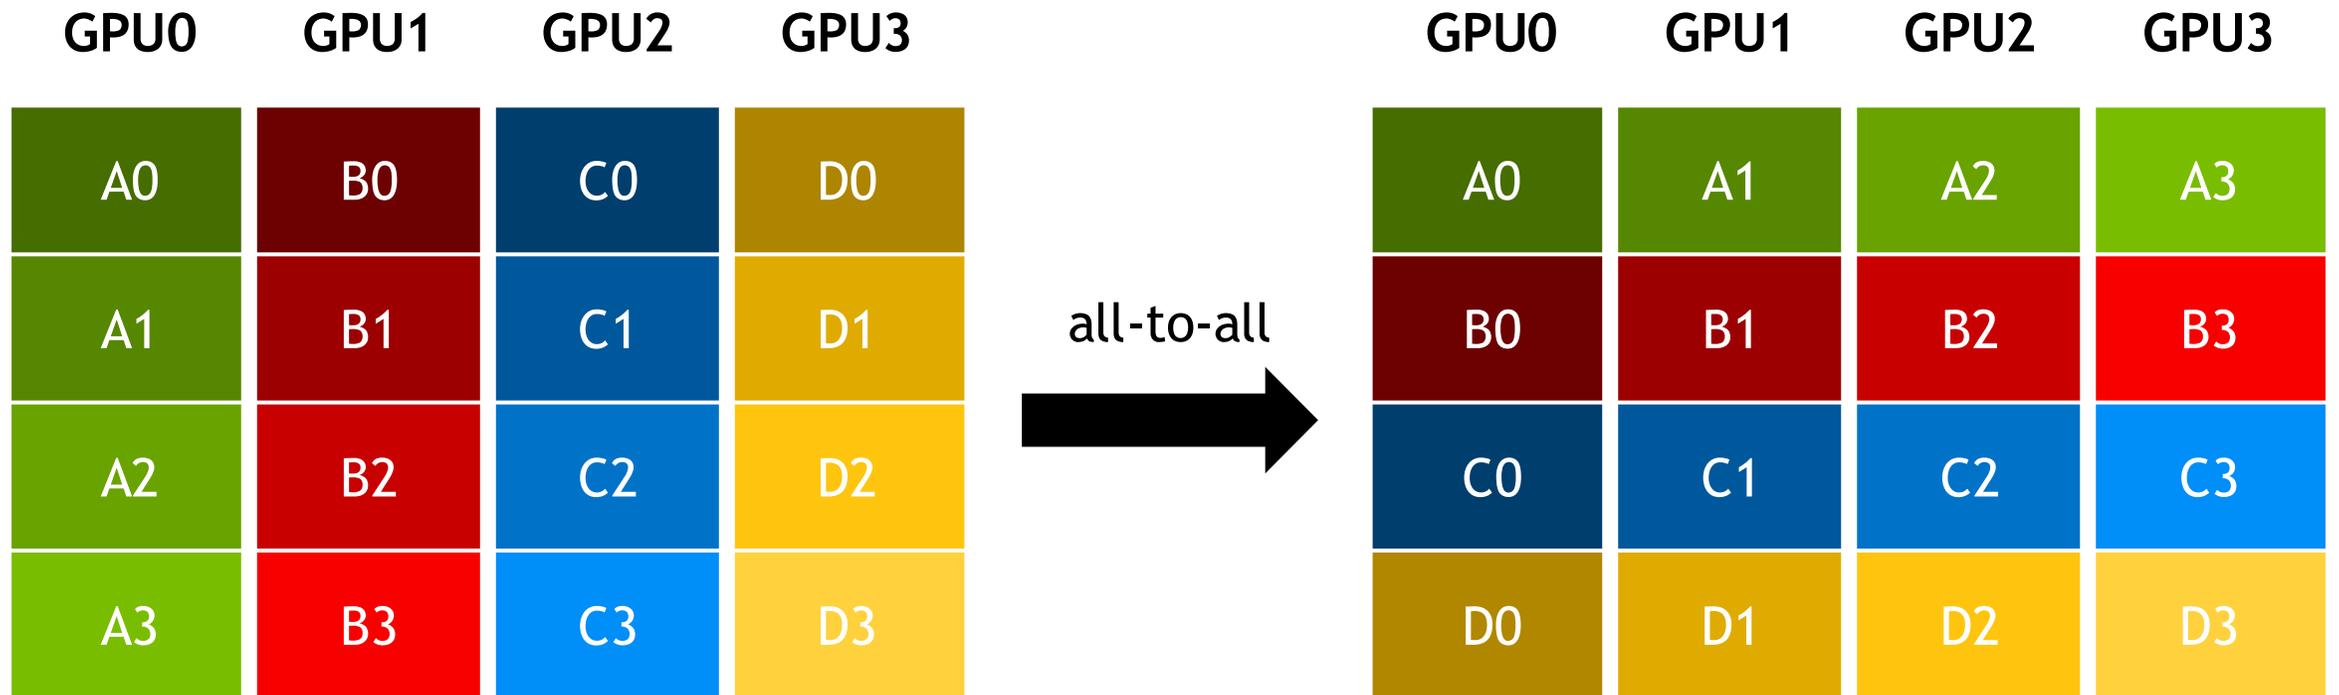
# REDUCE-SCATTER

Combine data from all senders; distribute result across participants



# ALL-TO-ALL

Scatter/Gather distinct messages from each participant to every other

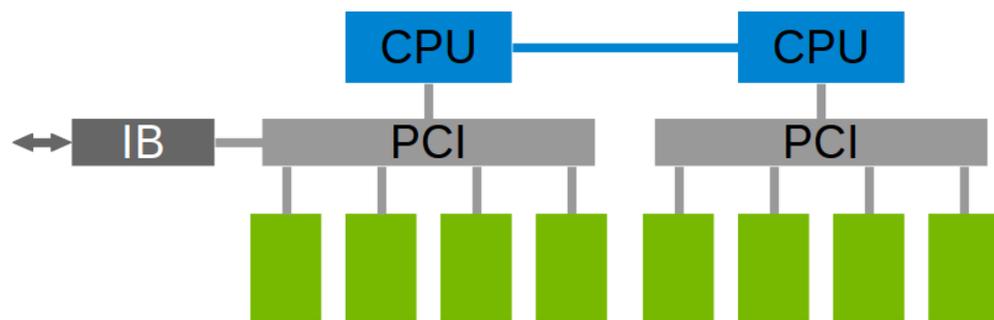


# THE CHALLENGE OF COLLECTIVES

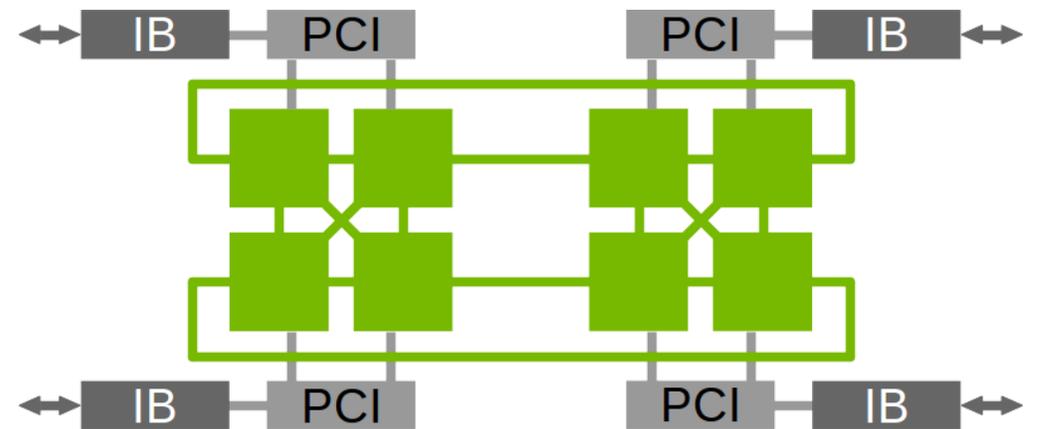
# NCCL

Inter-node communication using Sockets or Infiniband verbs, with multi-rail support, topology detection and automatic use of GPU Direct RDMA.

Optimal. combination of NVLink, PCI and network interfaces to maximize bandwidth and create rings across nodes



PCIe, Infiniband



DGX-1 : NVLink, 4x Infiniband

# THE CHALLENGE OF COLLECTIVES

Collectives are often avoided because they are expensive. Why?

Having multiple senders and/or receivers compounds communication inefficiencies

- For small transfers, latencies dominate; more participants increase latency
- For large transfers, bandwidth is key; bottlenecks are easily exposed
- May require topology-aware implementation for high performance
- Collectives are often blocking/non-overlapped

# THE CHALLENGE OF COLLECTIVES

If collectives are so expensive, do they actually get used? YES!

Collectives are central to scalability in a variety of key applications:

- Deep Learning (All-reduce, broadcast, gather)
- Parallel FFT (Transposition is all-to-all)
- Molecular Dynamics (All-reduce)
- Graph Analytics (All-to-all)
- ...

# THE CHALLENGE OF COLLECTIVES

Many implementations seen in the wild are suboptimal

Scaling requires efficient communication algorithms and careful implementation

Communication algorithms are topology-dependent

Topologies can be complex - not every system is a fat tree

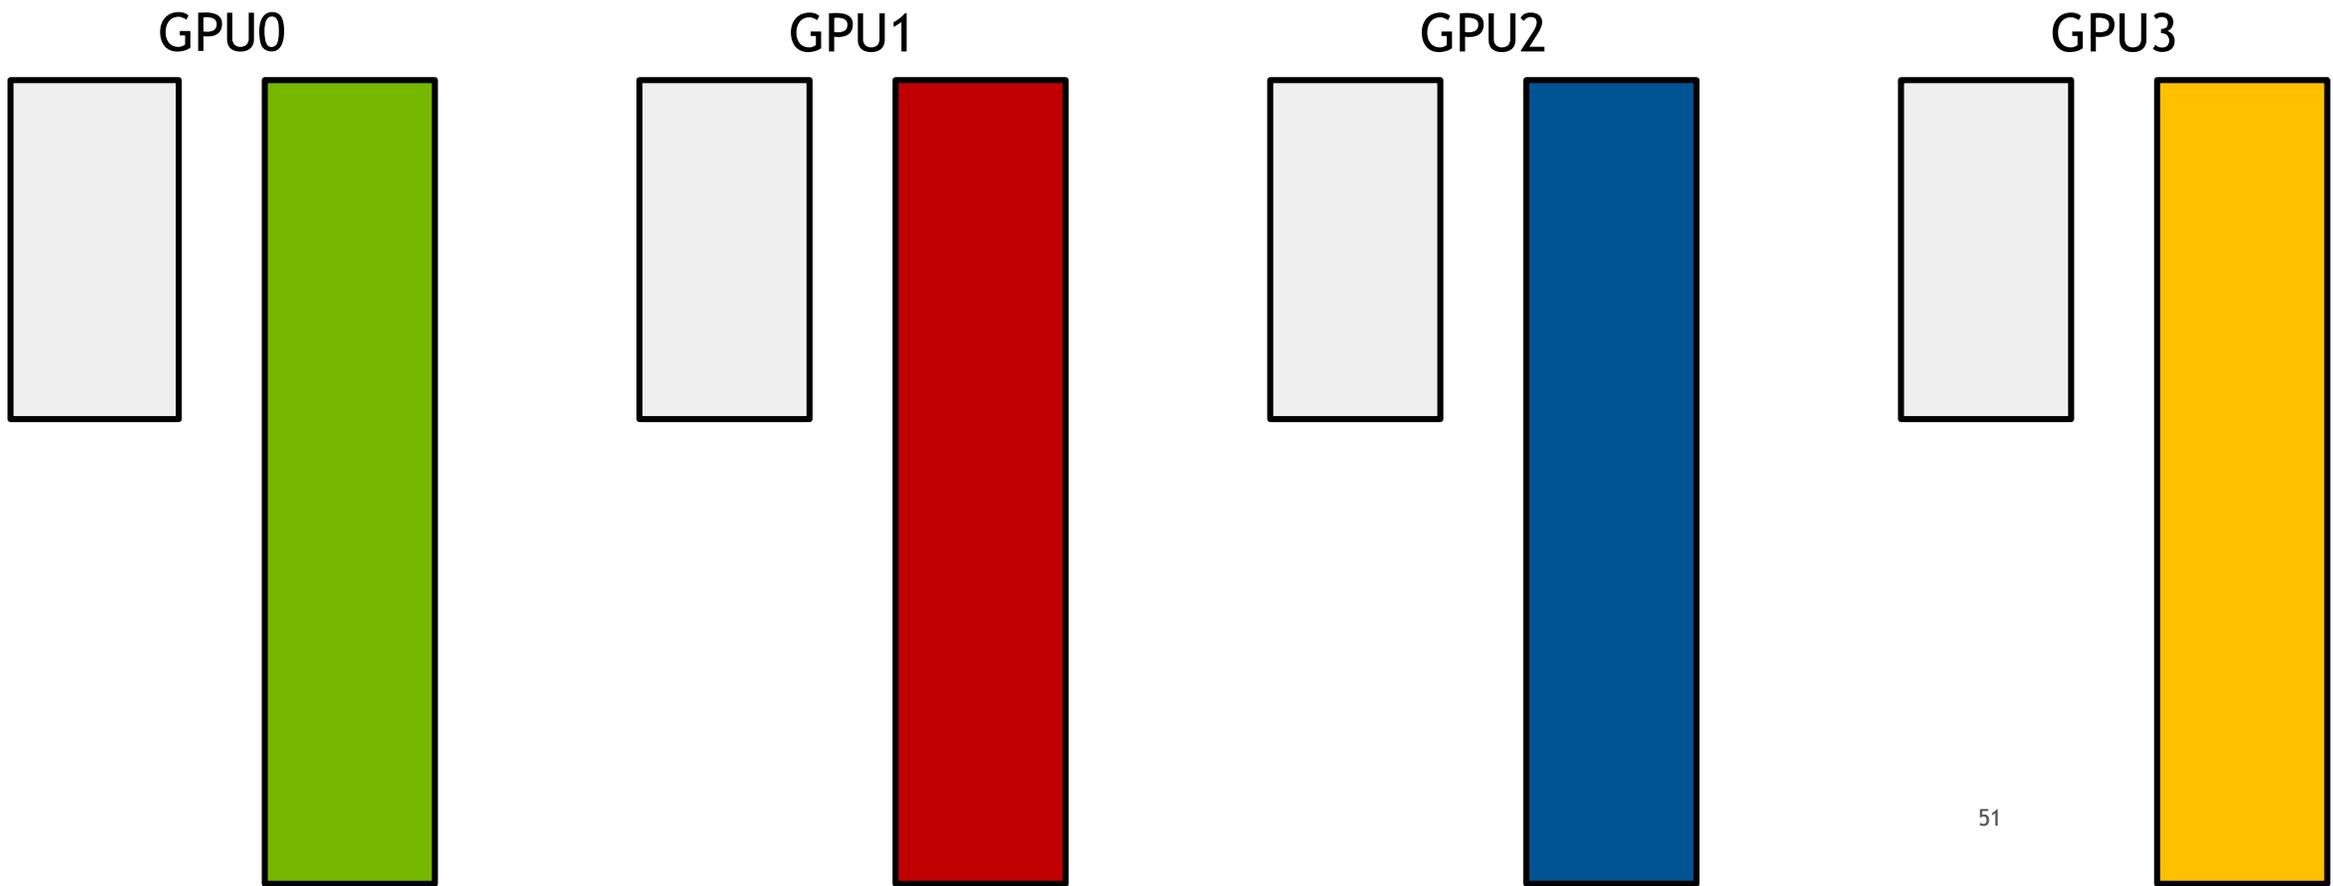
Most collectives amenable to bandwidth-optimal implementation on rings, and many topologies can be interpreted as one or more rings [P. Patarasuk and X. Yuan]

**All-reduce**

# ALL-REDUCE

with unidirectional ring

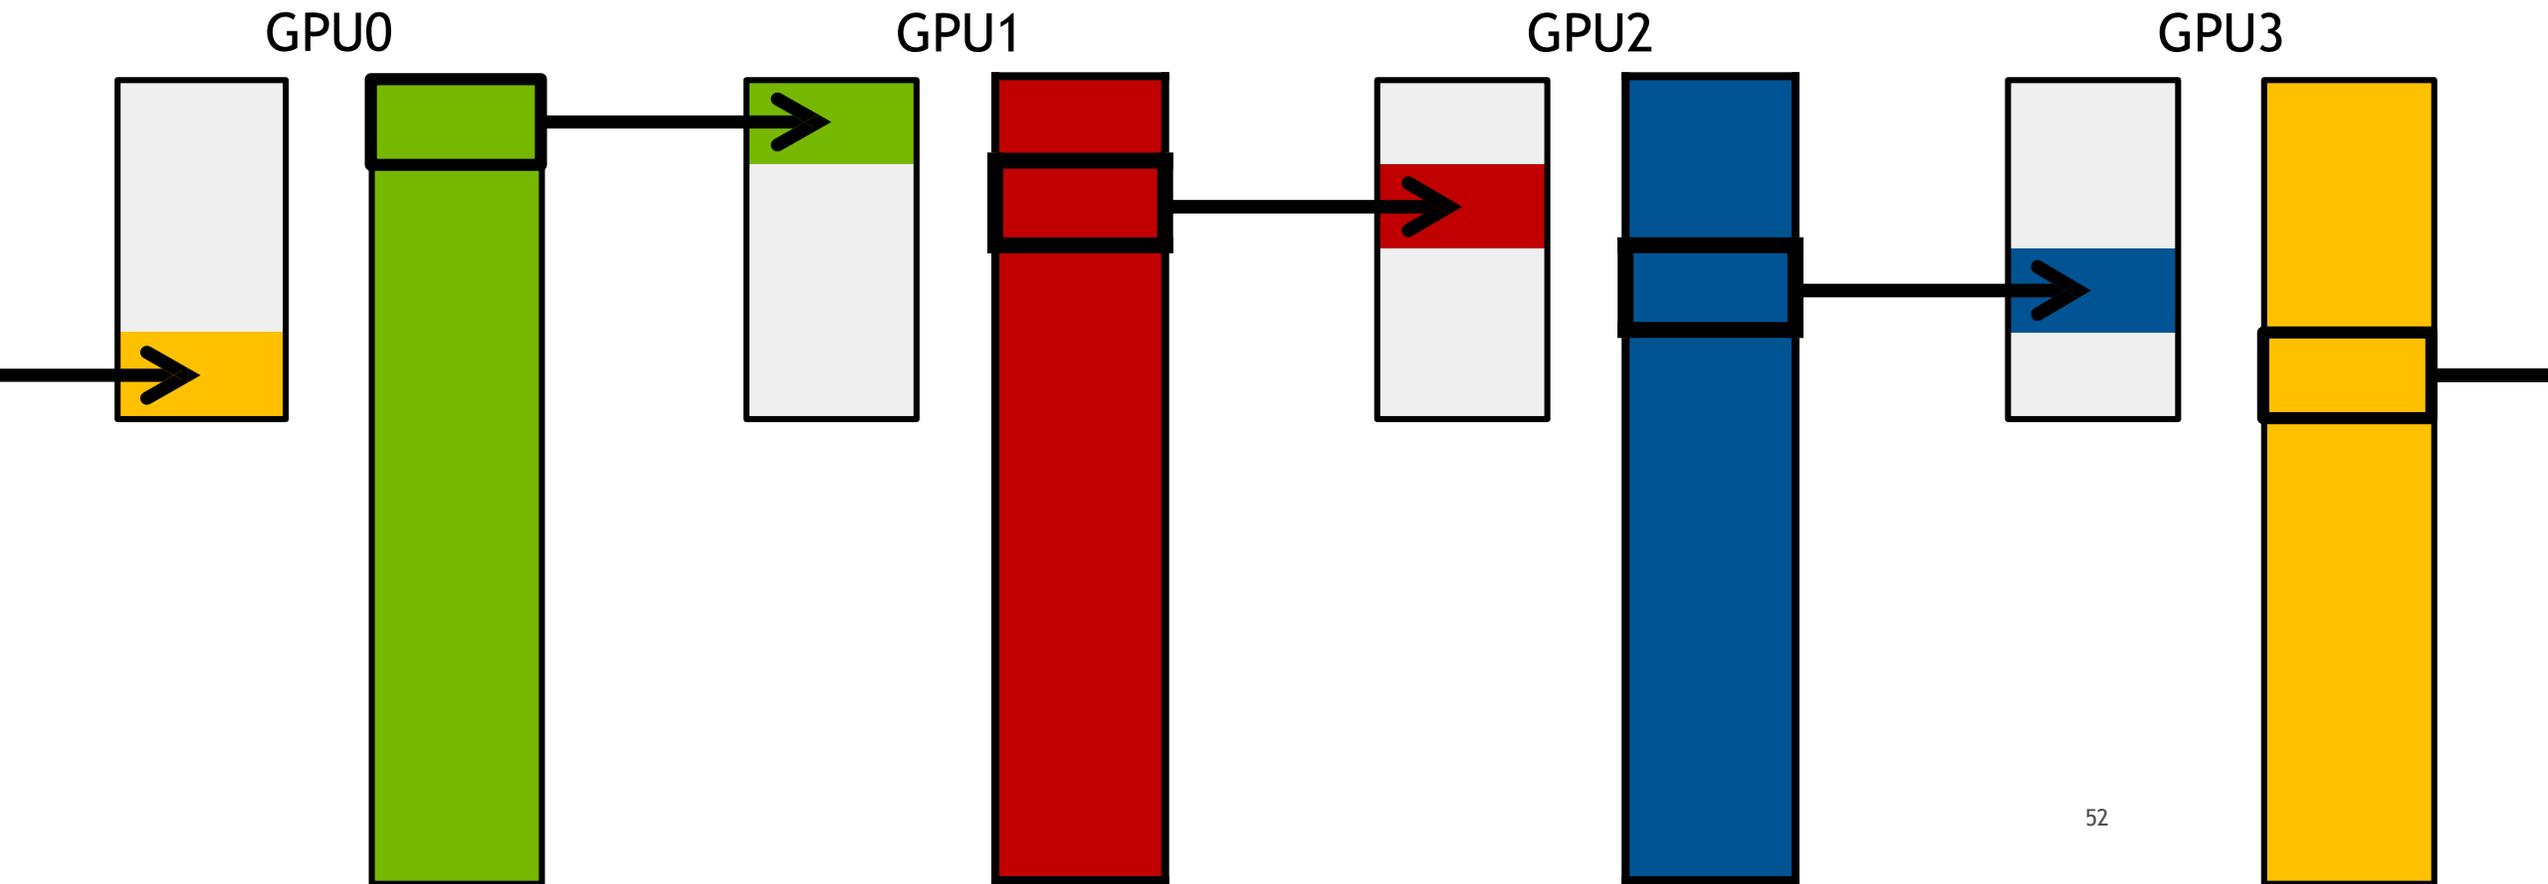
Chunk: 1  
Step:



# ALL-REDUCE

with unidirectional ring

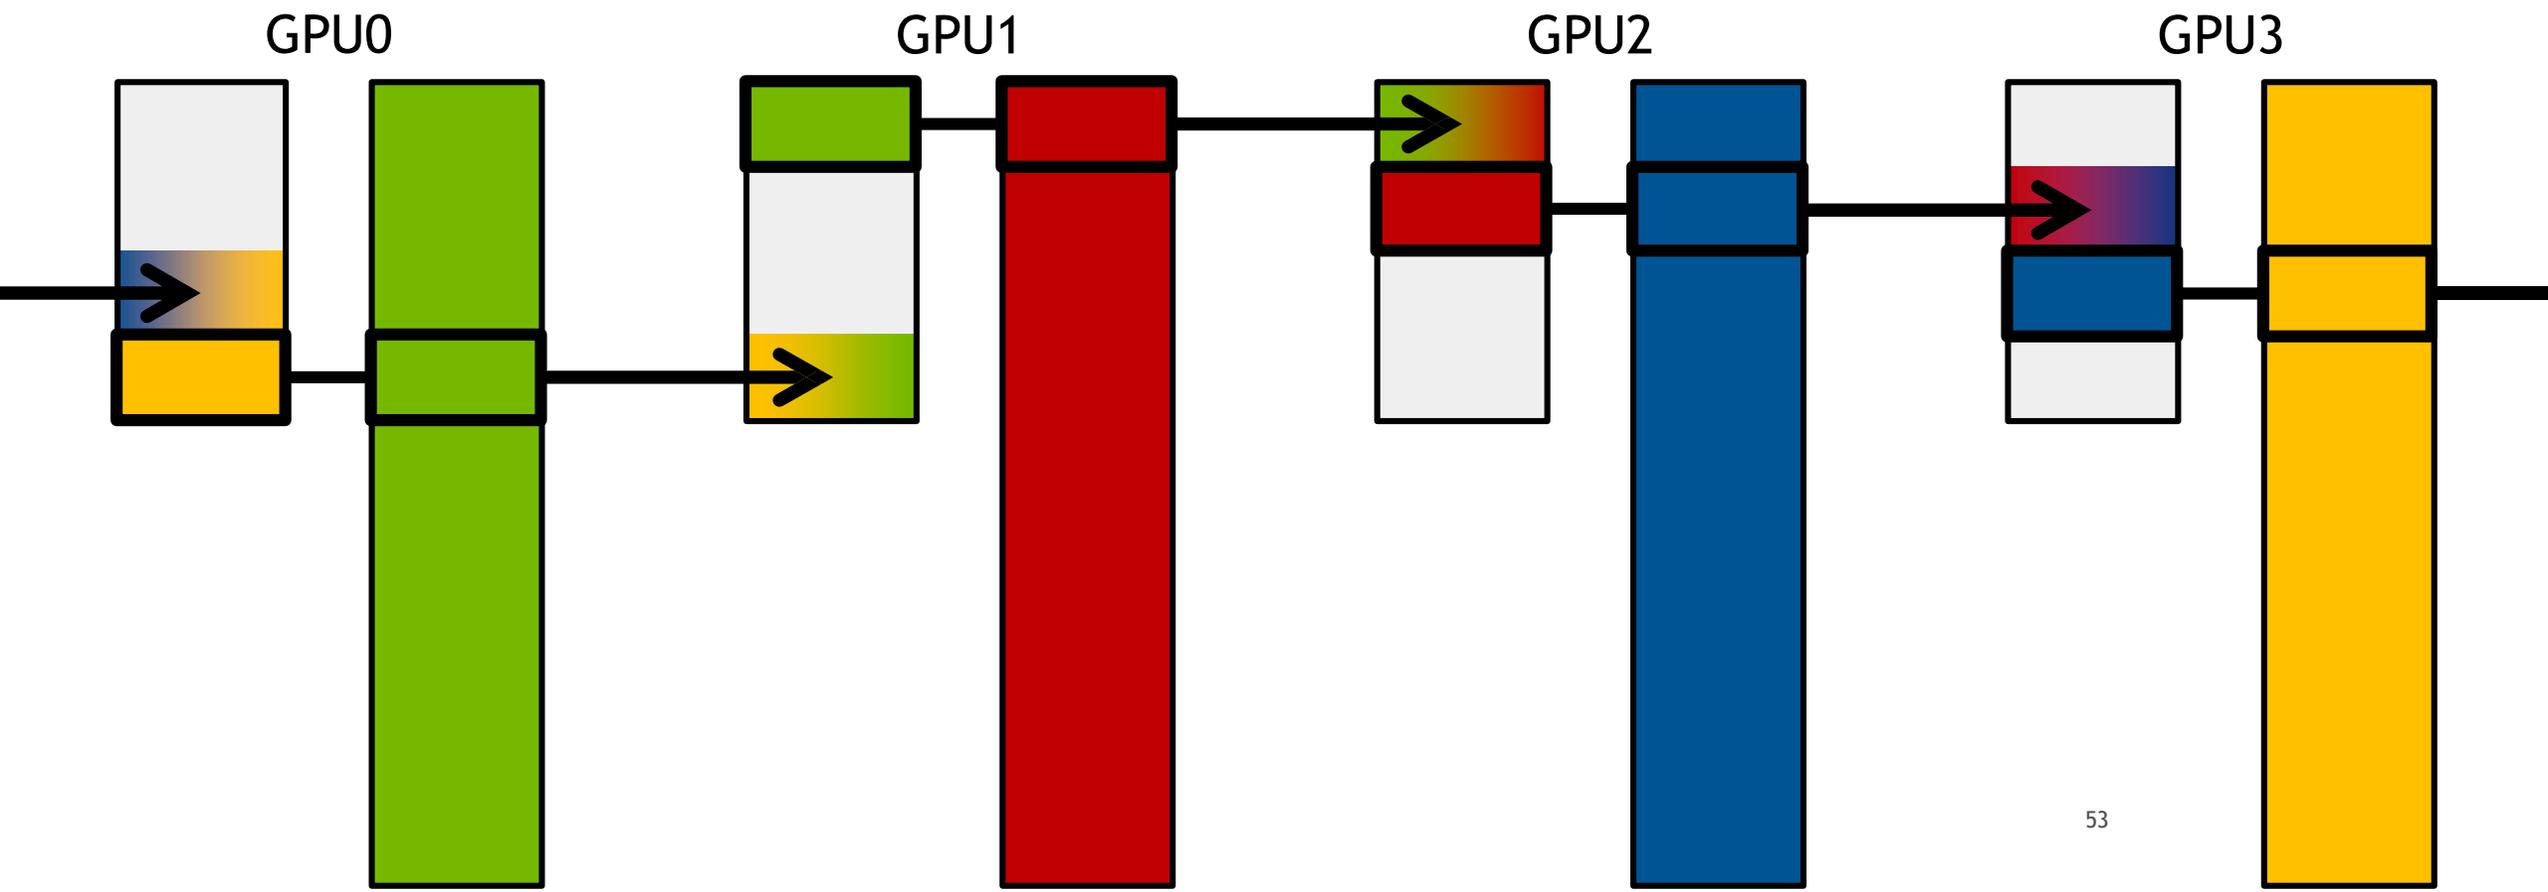
Chunk: 1  
Step: 1



# ALL-REDUCE

with unidirectional ring

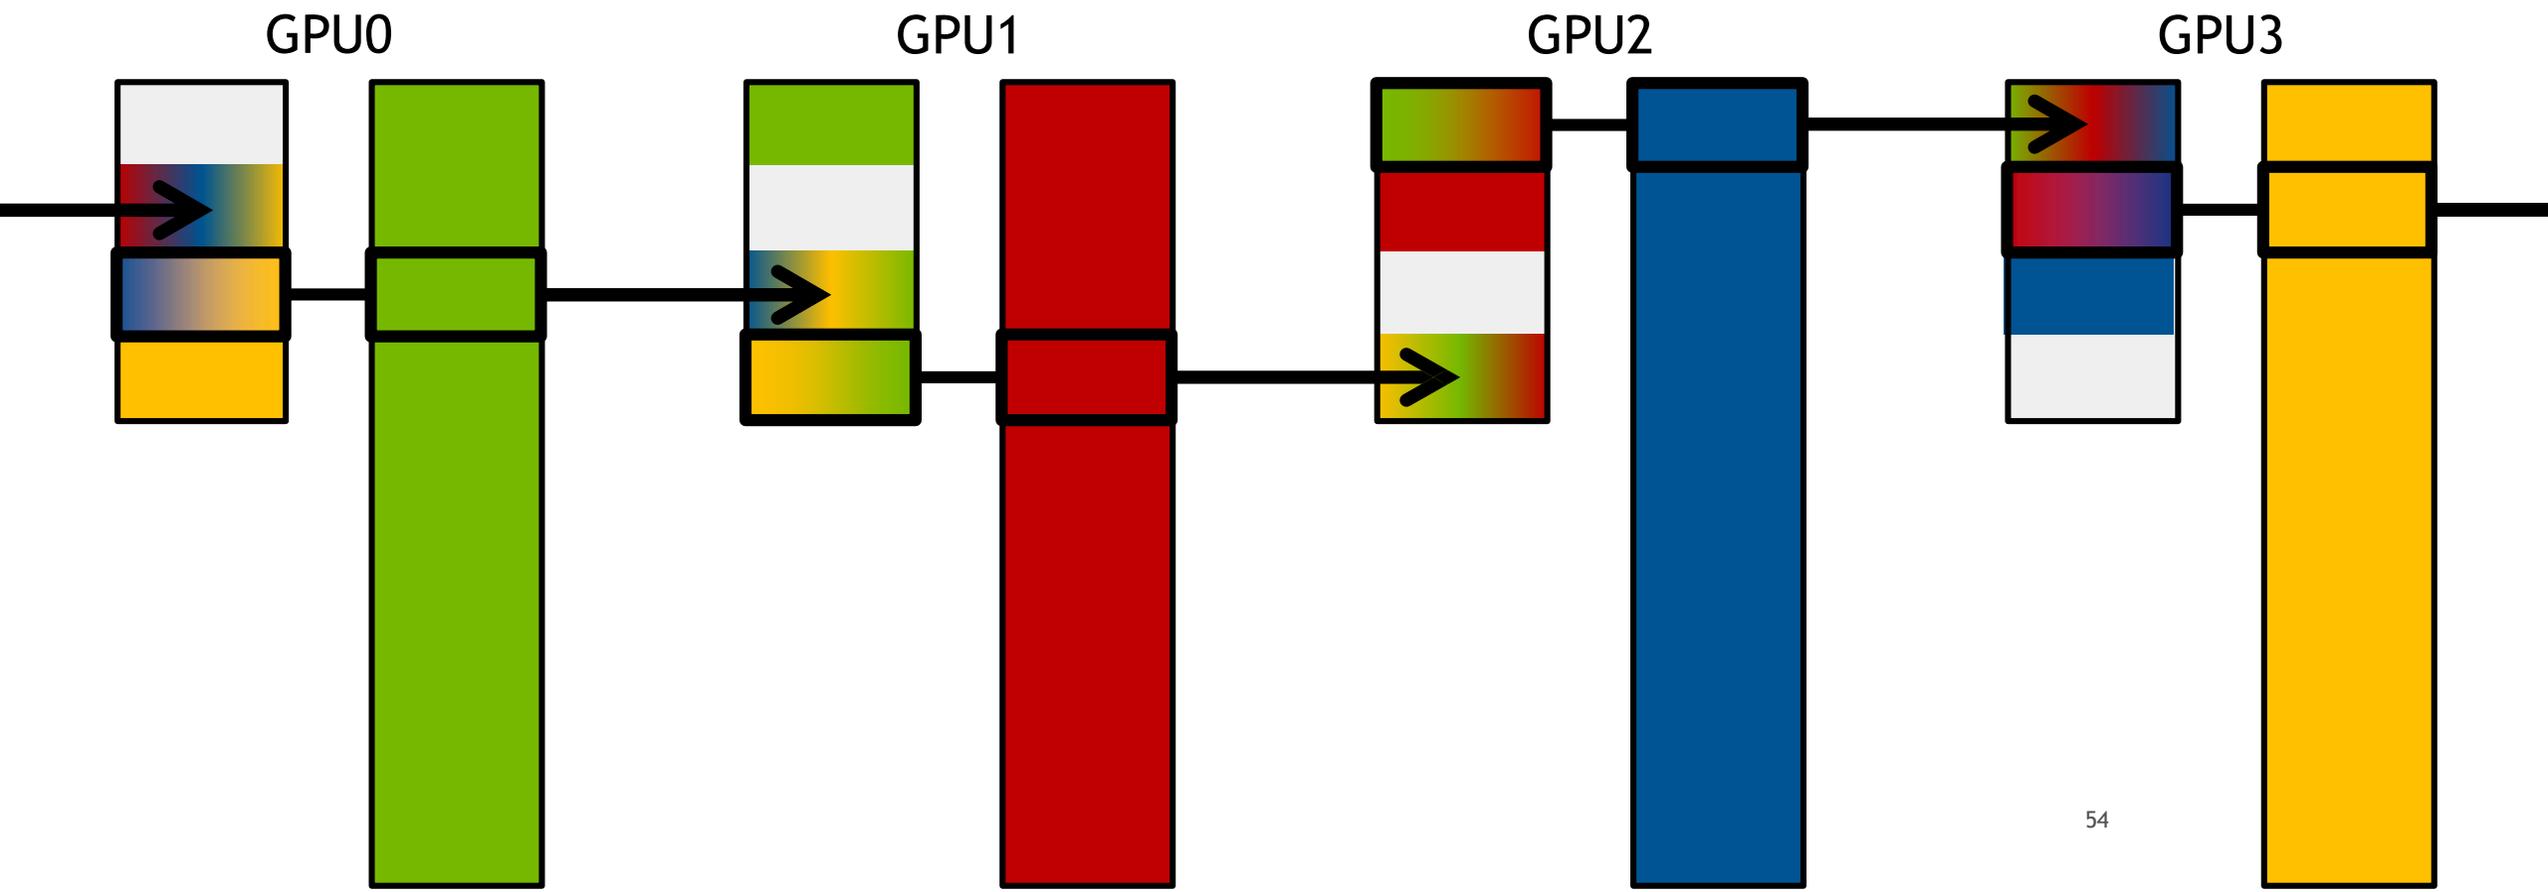
Chunk: 1  
Step: 2



# ALL-REDUCE

with unidirectional ring

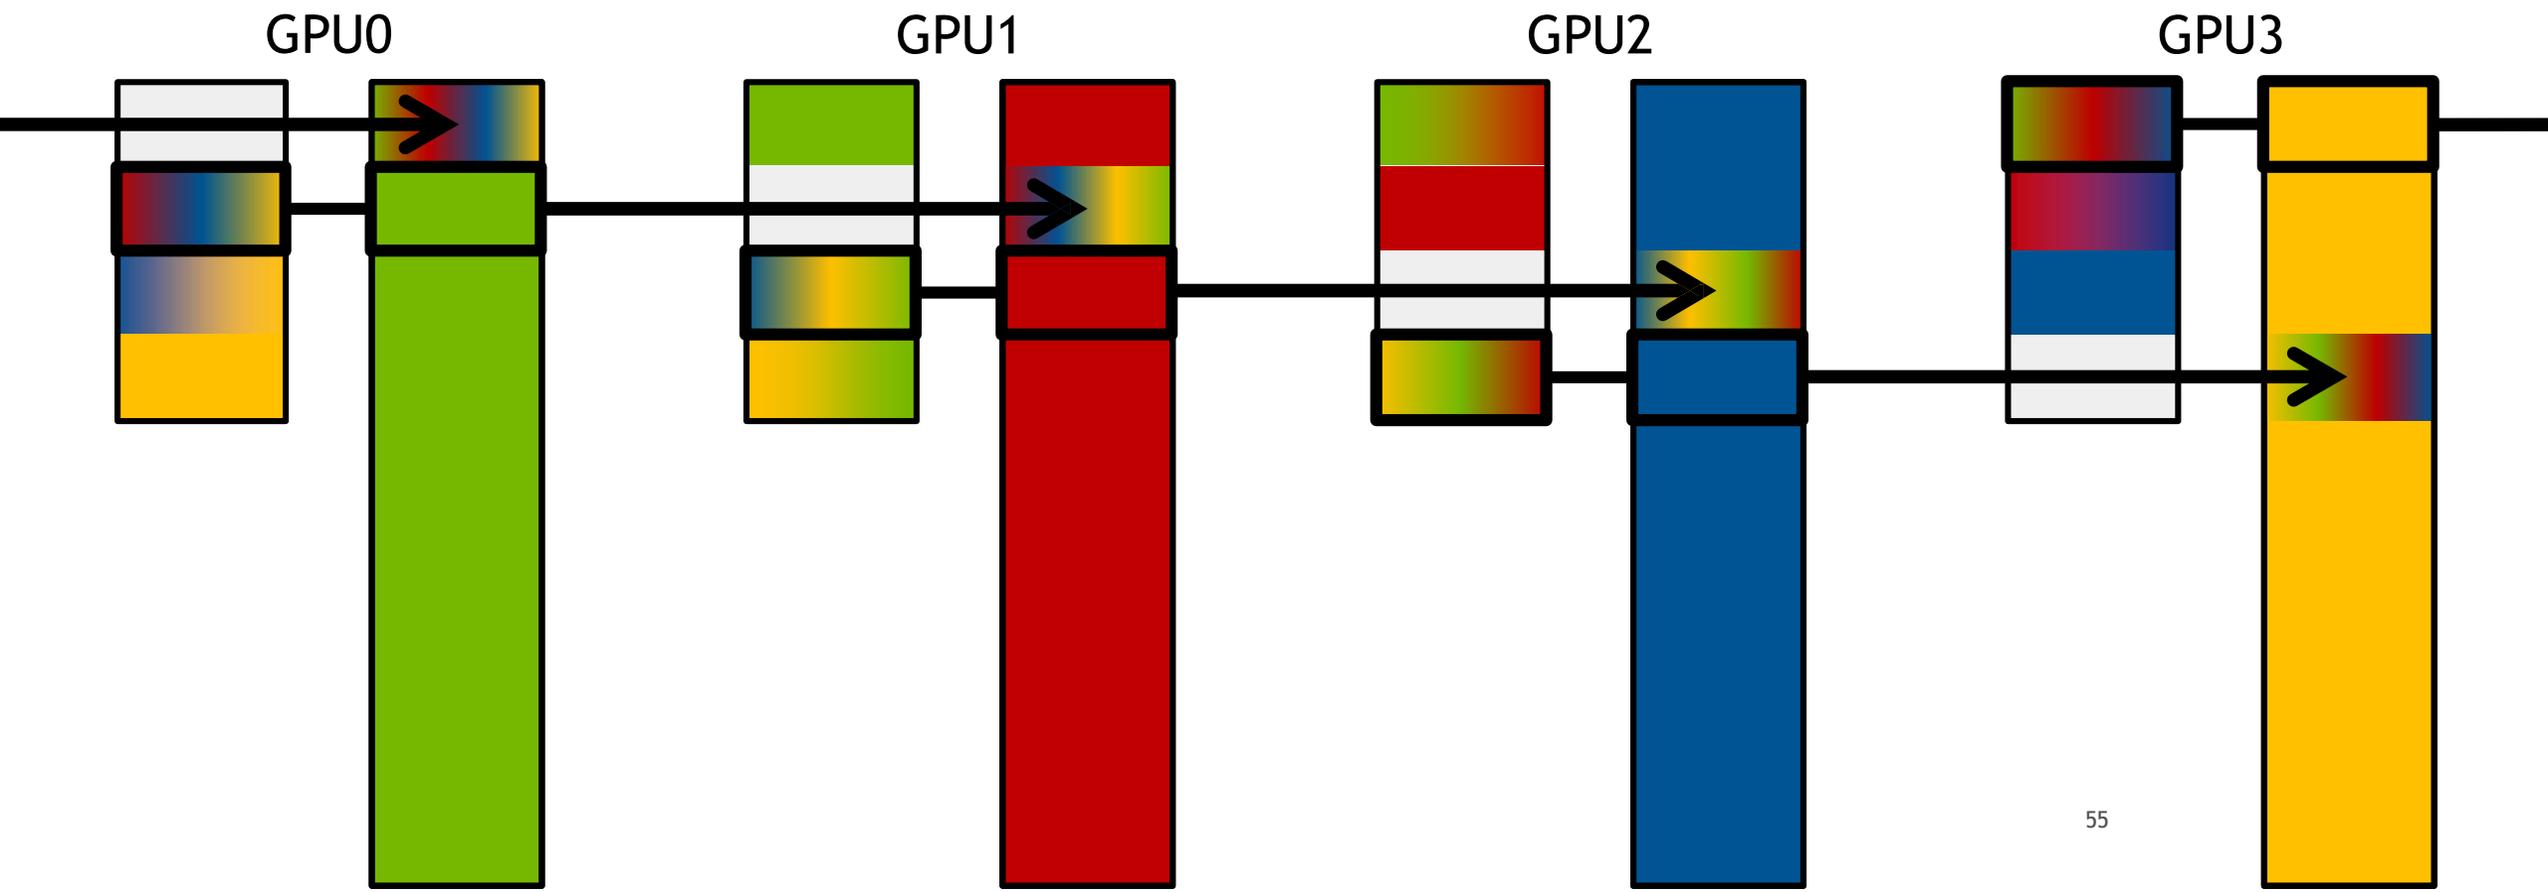
Chunk: 1  
Step: 3



# ALL-REDUCE

with unidirectional ring

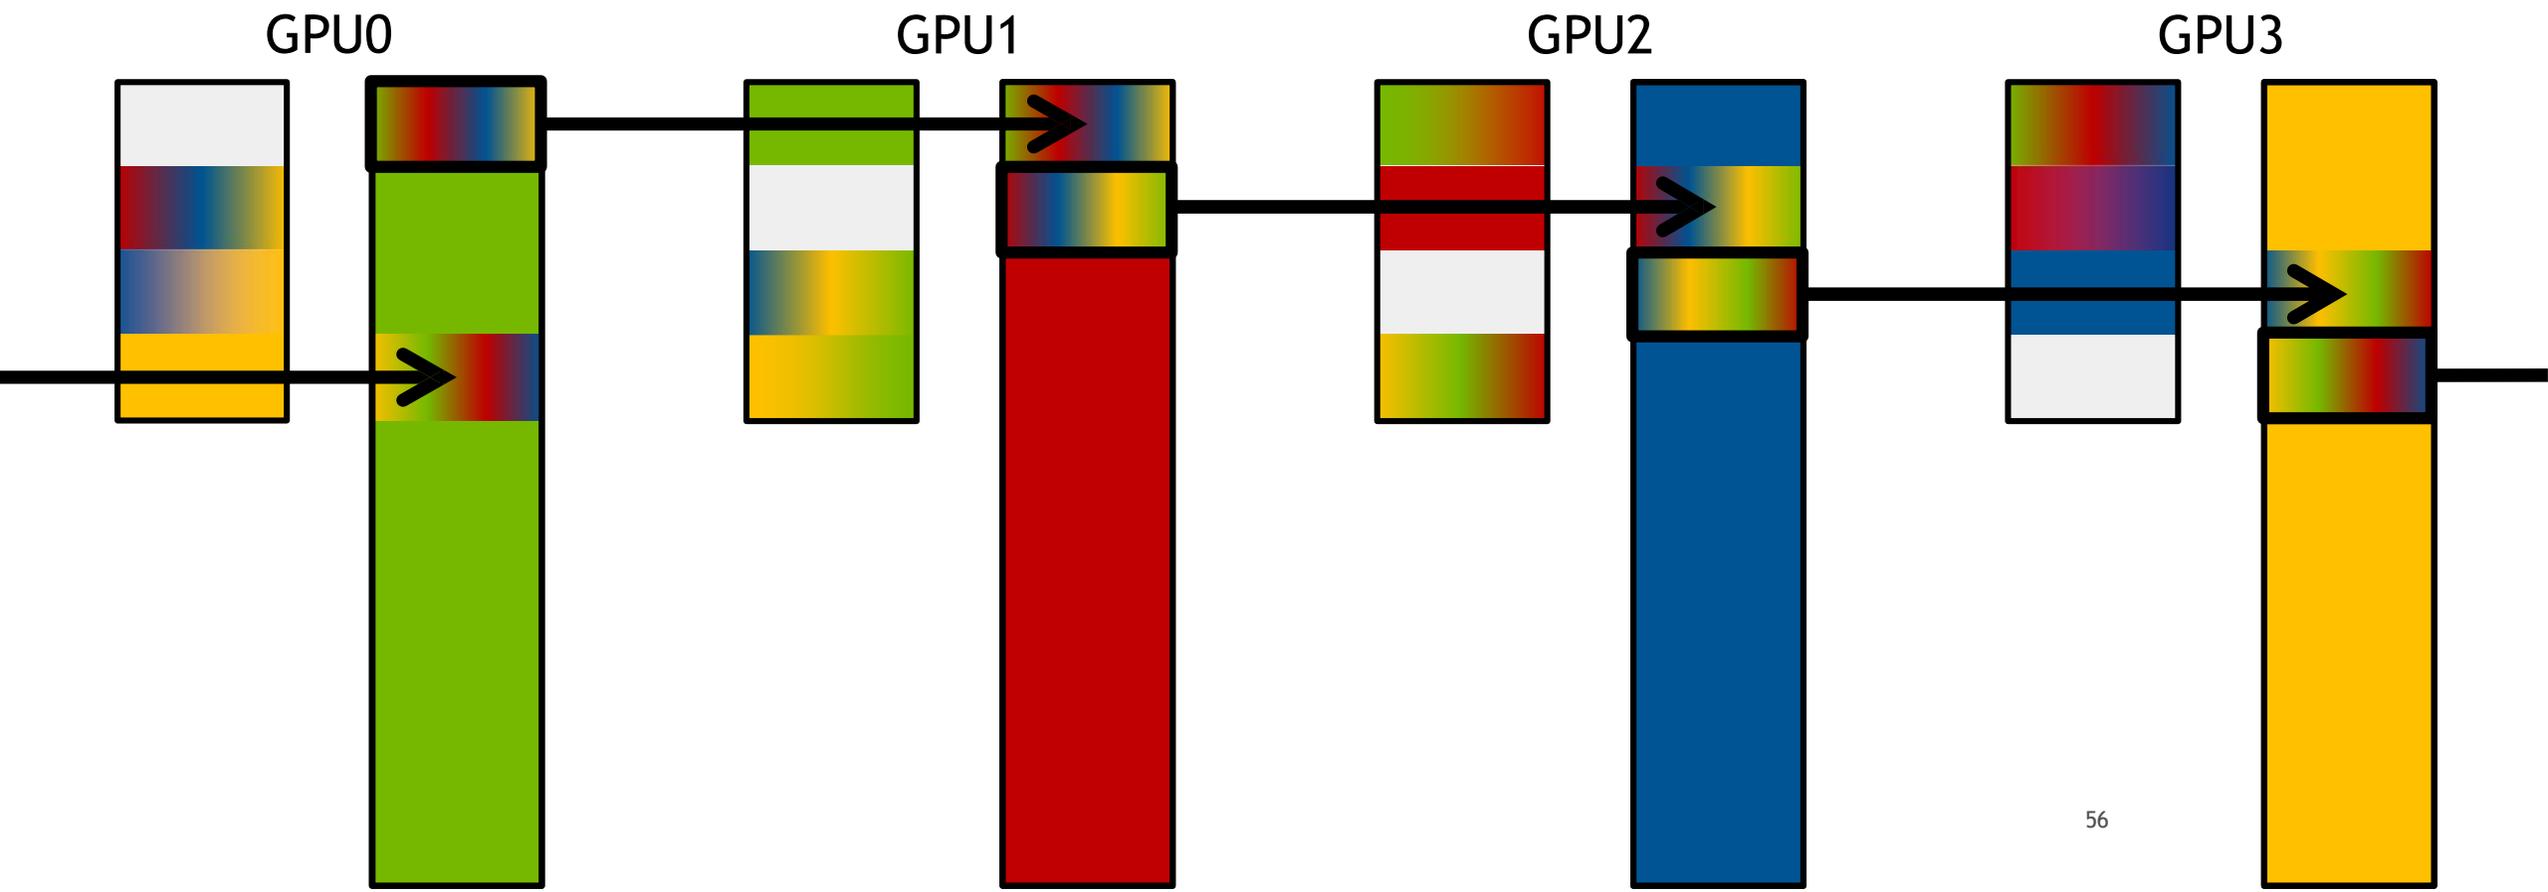
Chunk: 1  
Step: 4



# ALL-REDUCE

with unidirectional ring

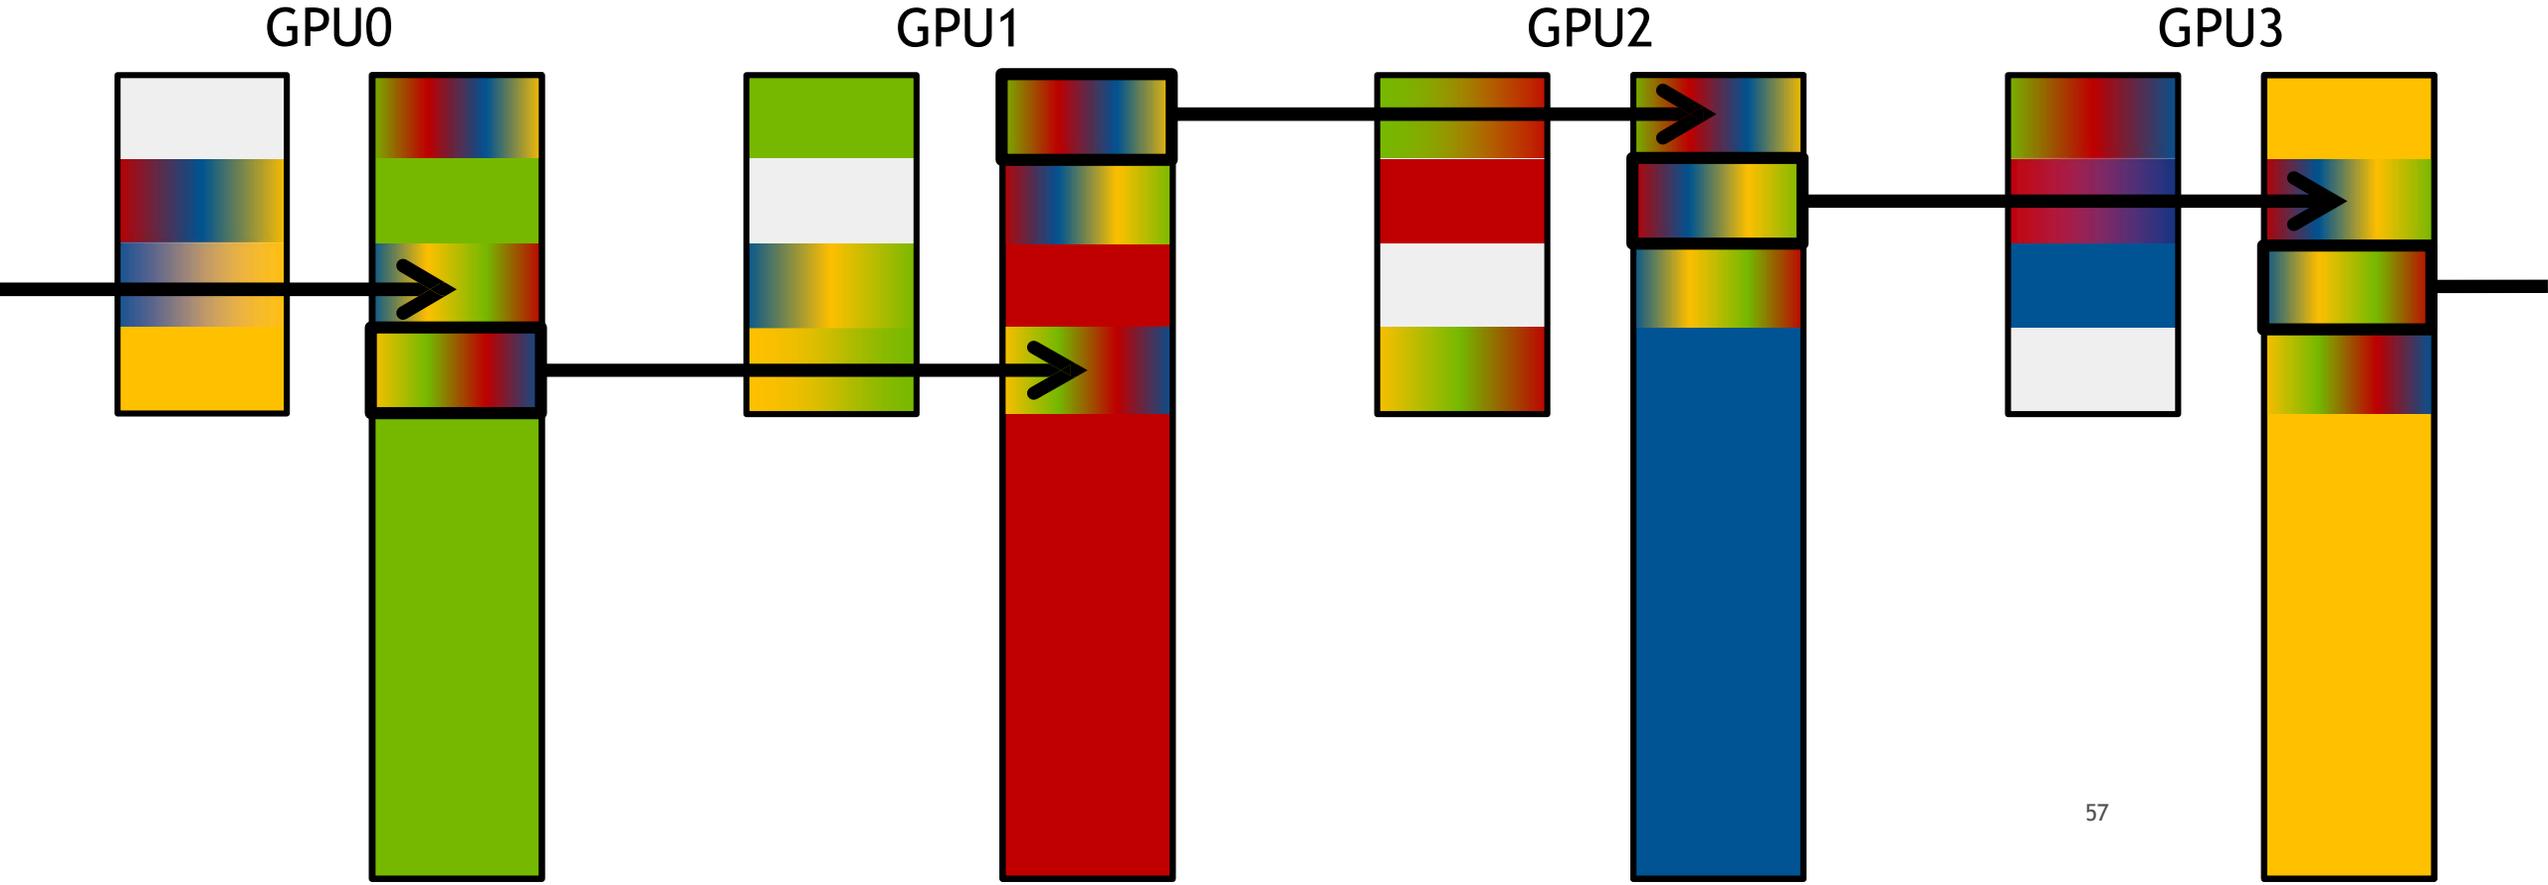
Chunk: 1  
Step: 5



# ALL-REDUCE

with unidirectional ring

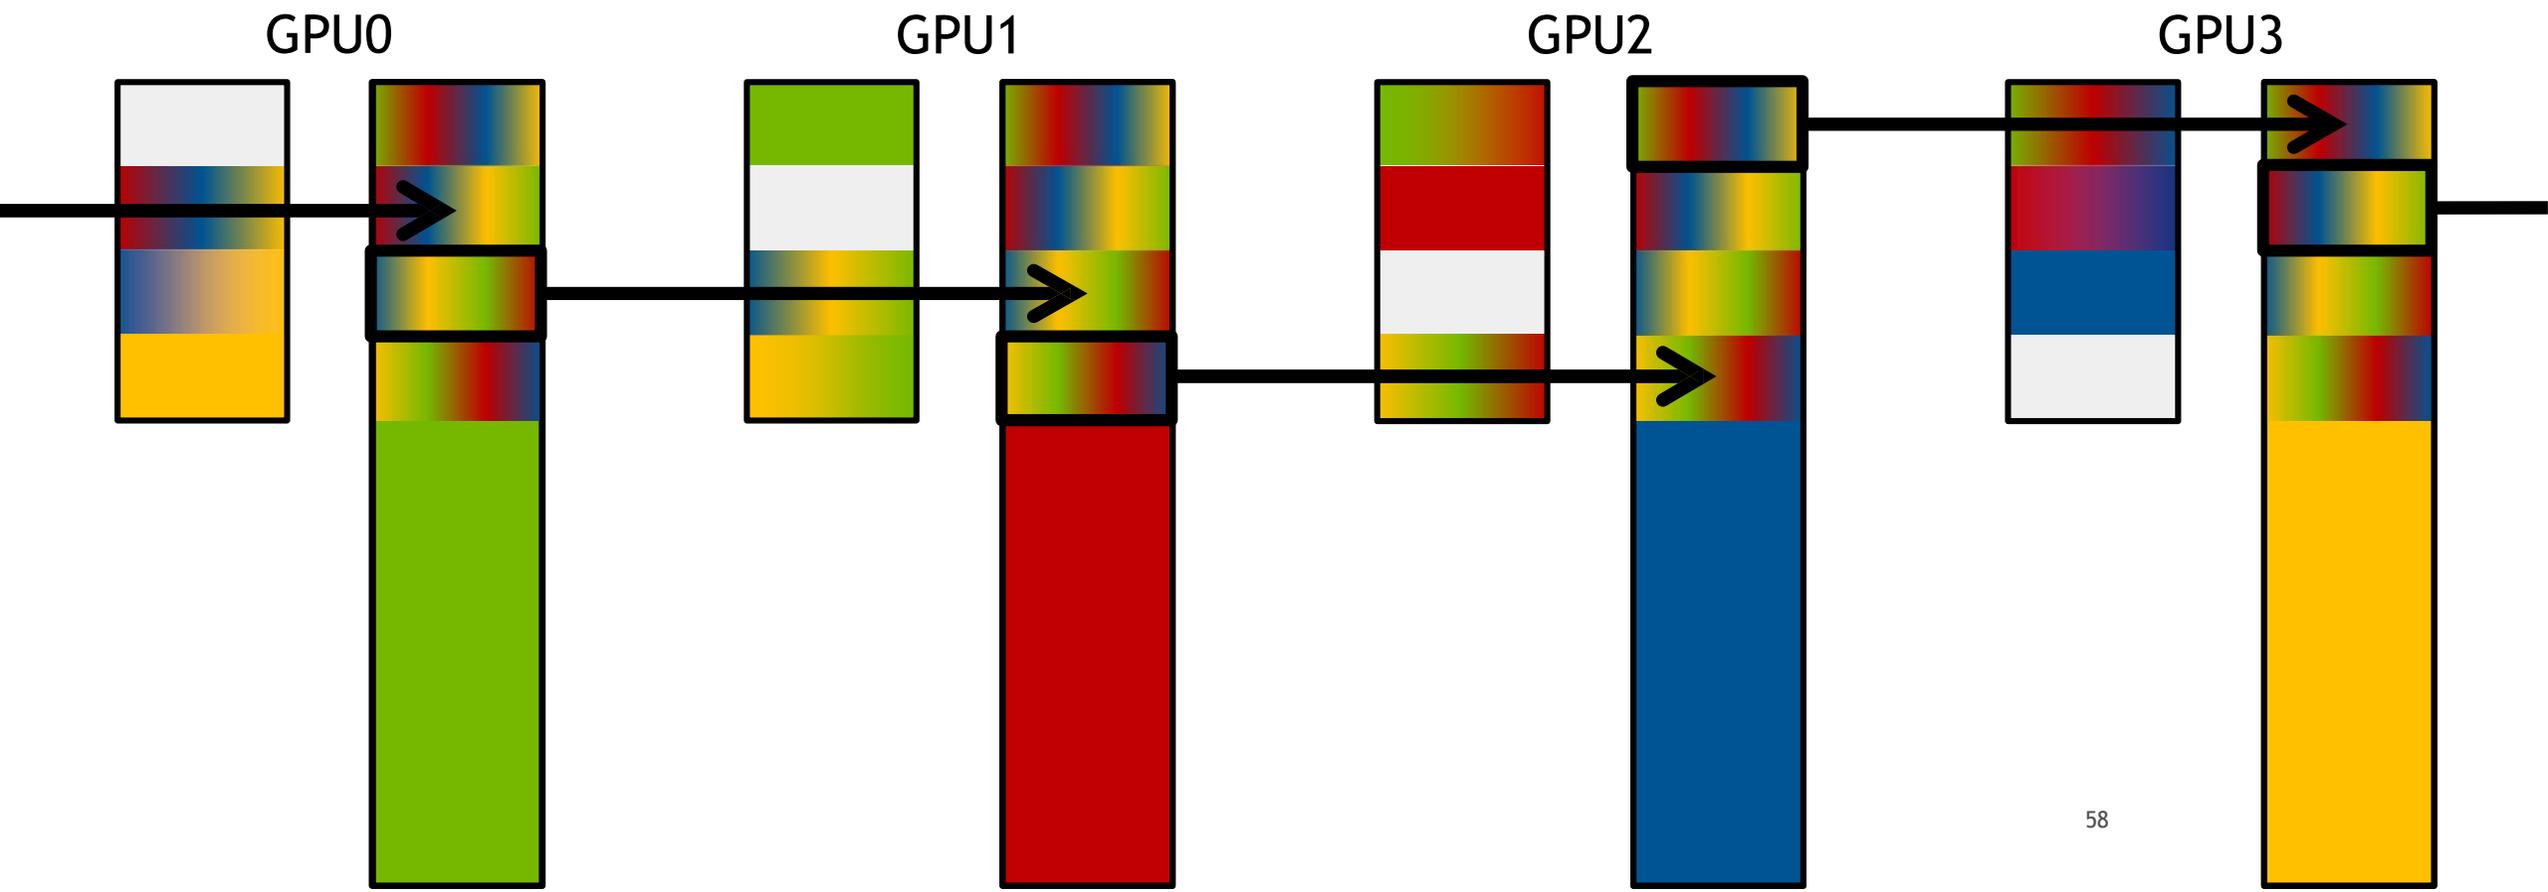
Chunk: 1  
Step: 6



# ALL-REDUCE

with unidirectional ring

Chunk: 1  
Step: 7

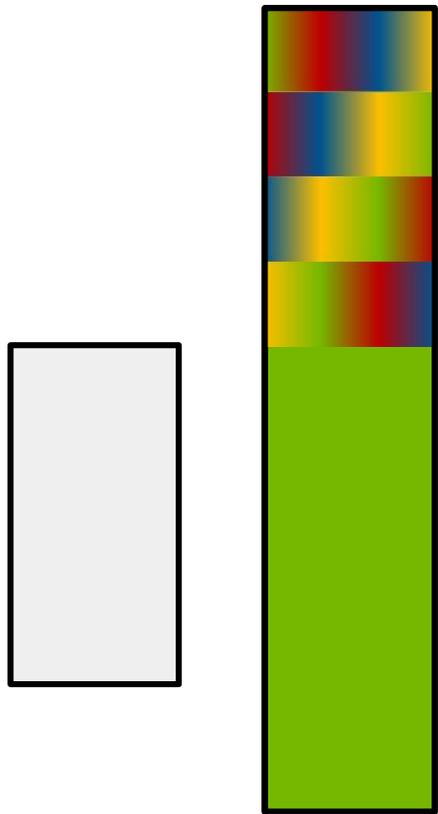


# ALL-REDUCE

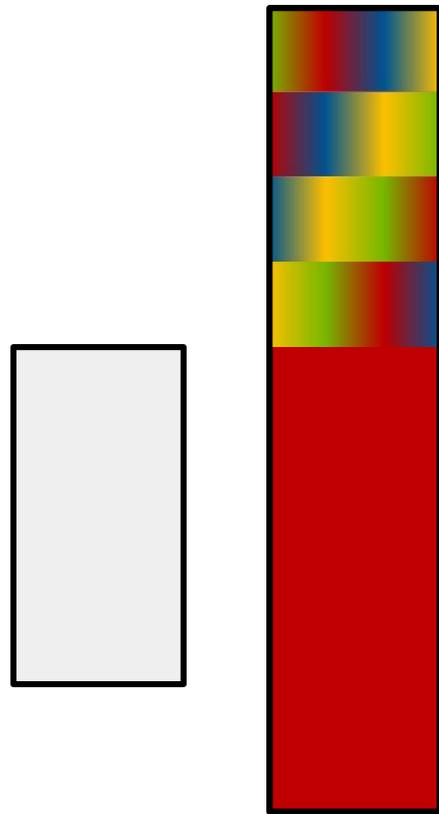
with unidirectional ring

Chunk: 2  
Step:

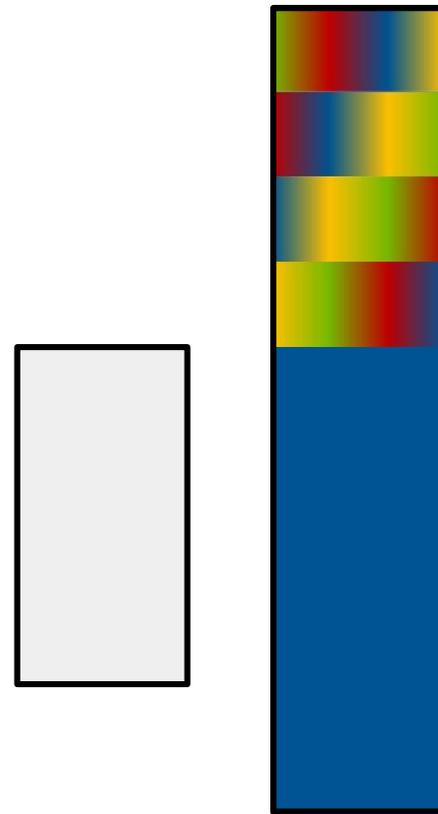
GPU0



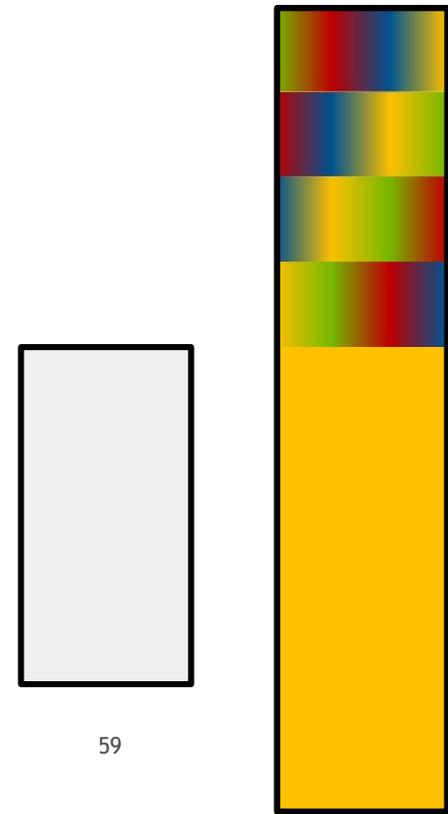
GPU1



GPU2



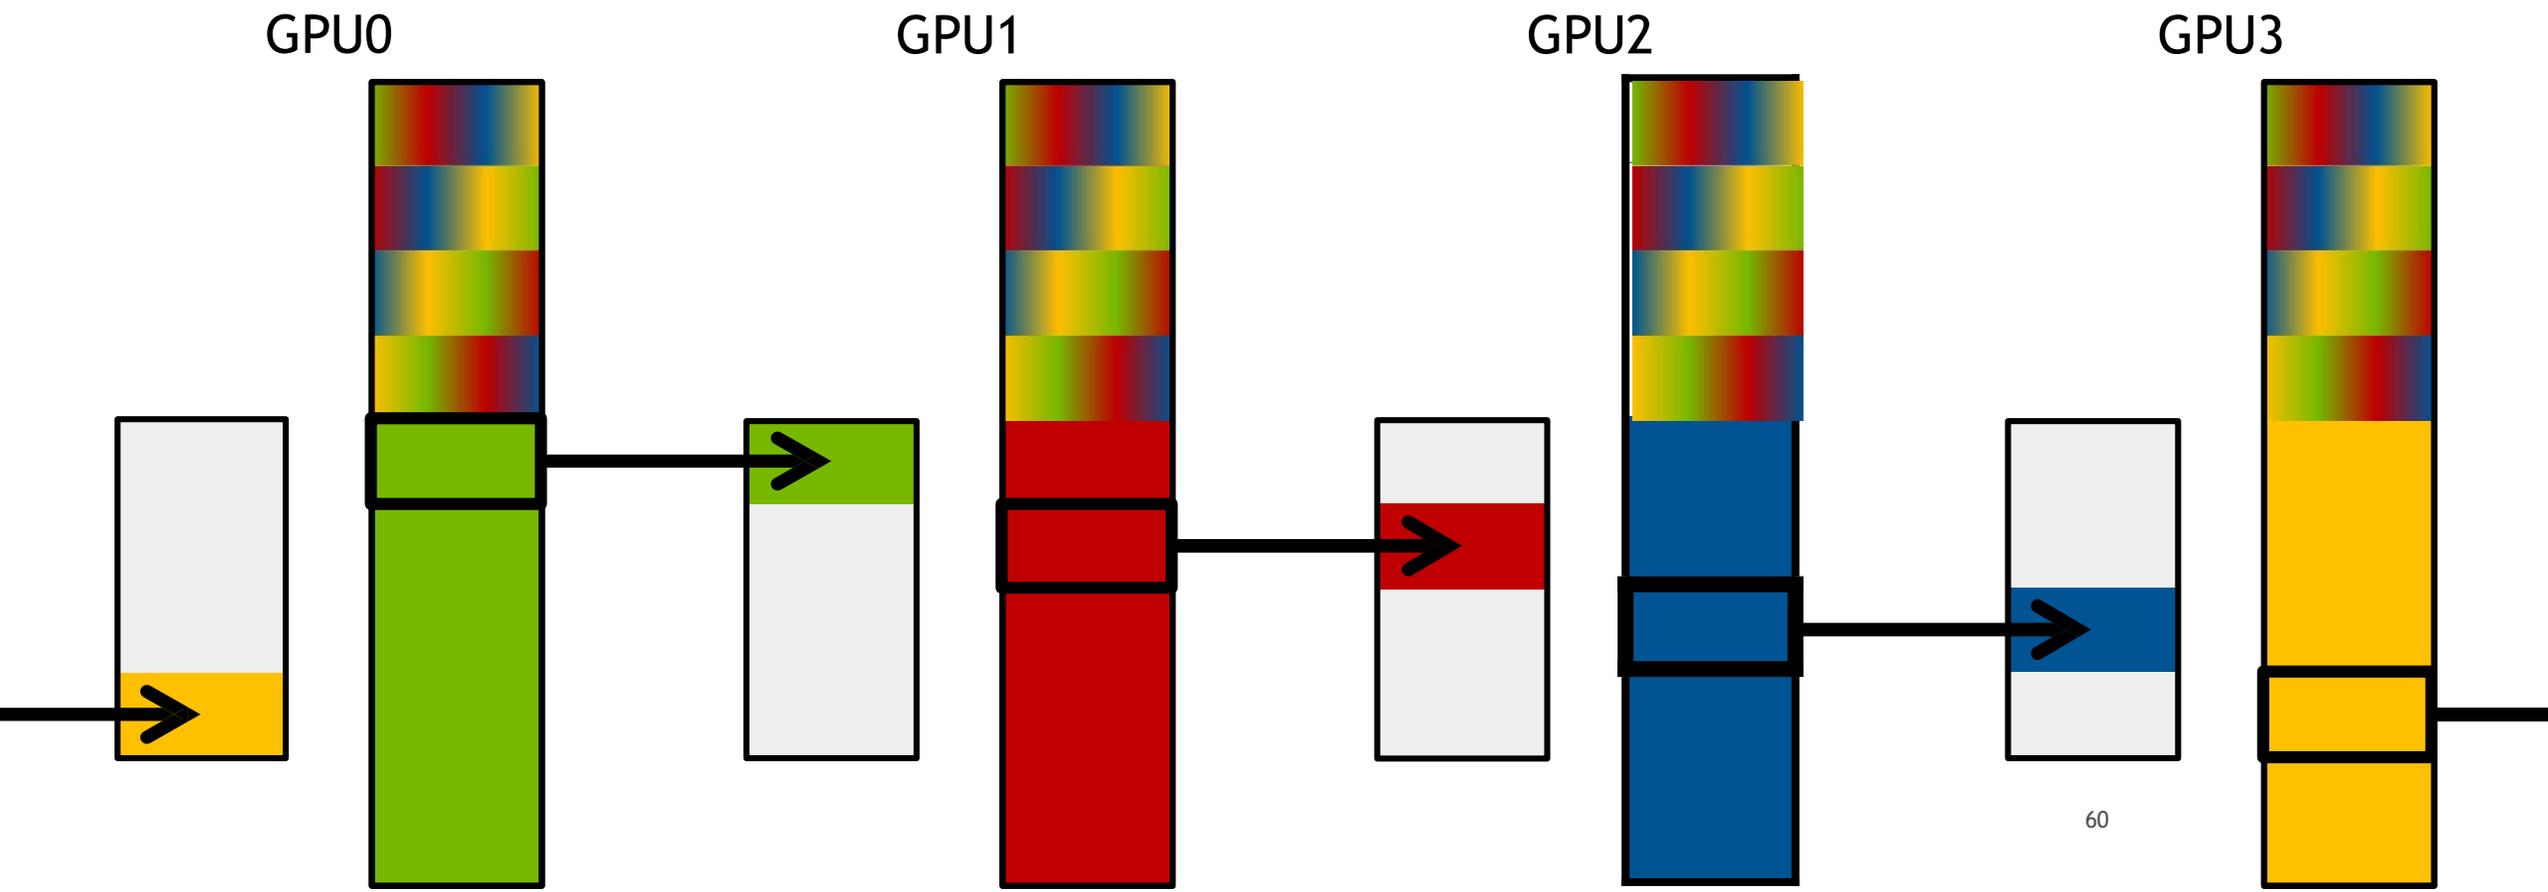
GPU3



# ALL-REDUCE

with unidirectional ring

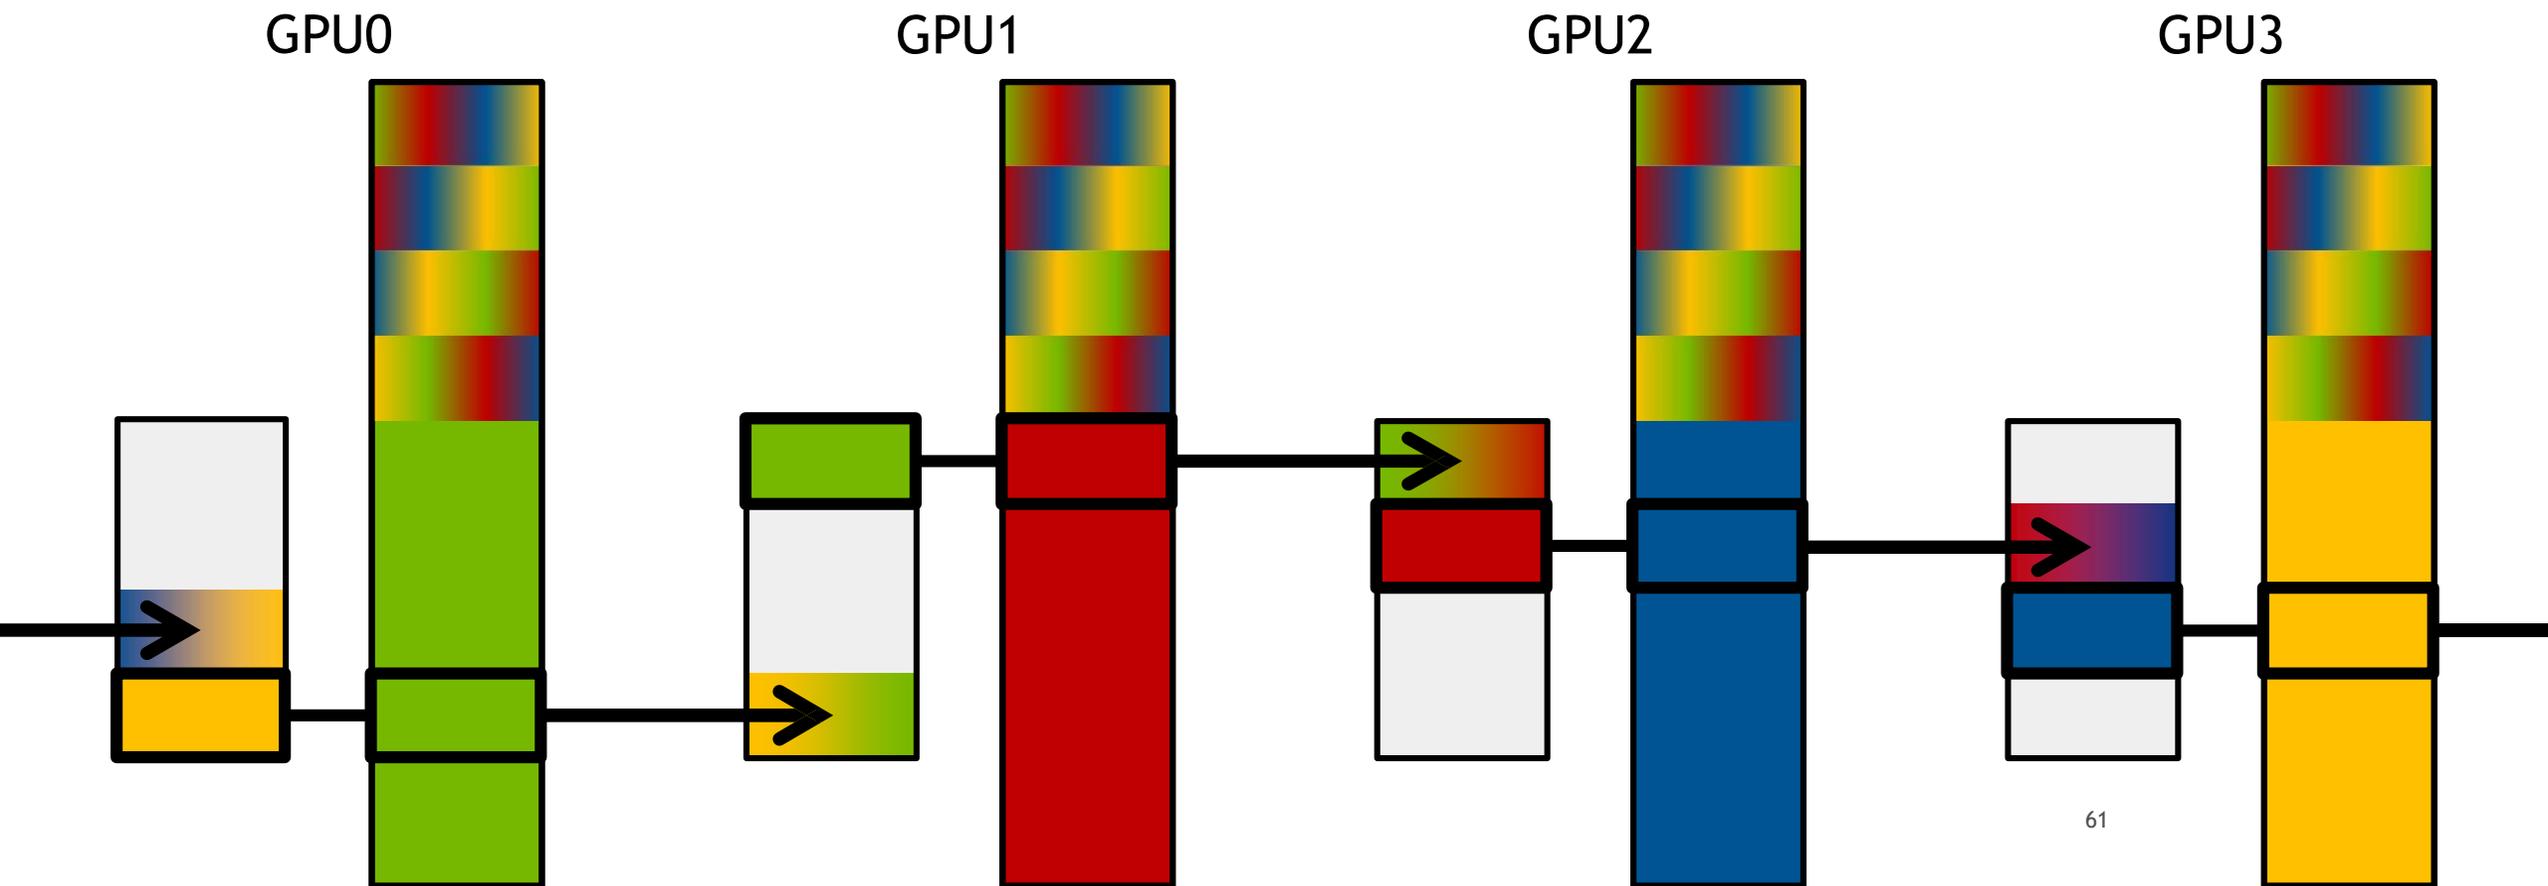
Chunk: 2  
Step: 1



# ALL-REDUCE

with unidirectional ring

Chunk: 2  
Step: 2

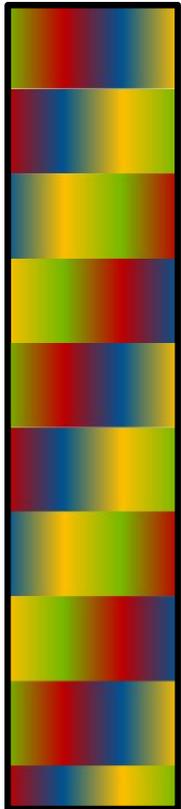


# ALL-REDUCE

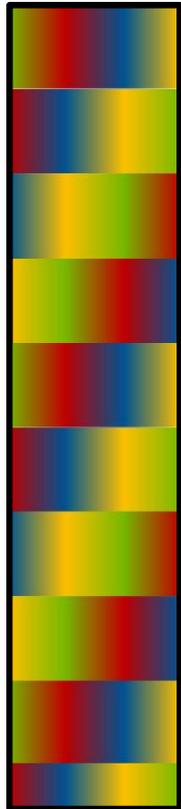
with unidirectional ring

done

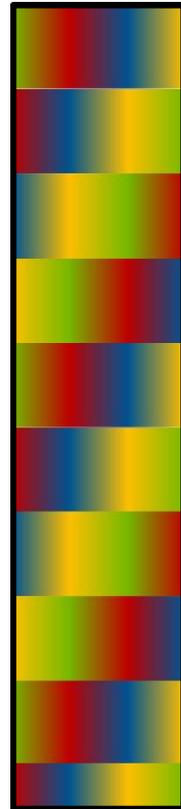
GPU0



GPU1



GPU2



GPU3

