



# EnvPipe: Performance-preserving DNN Training Framework for Saving Energy

Sangjin Choi, Inhoe Koo,  
Jeongseob Ahn, Myeongjae Jeon, Youngjin Kwon



# Need for Sustainable AI



- Training GPT-3 takes 14.8 days with 10,000 V100 GPUs consuming **1,287MWh**<sup>[1]</sup>
- Enough energy to power approx. **44,183 US households** for one day<sup>[2]</sup>

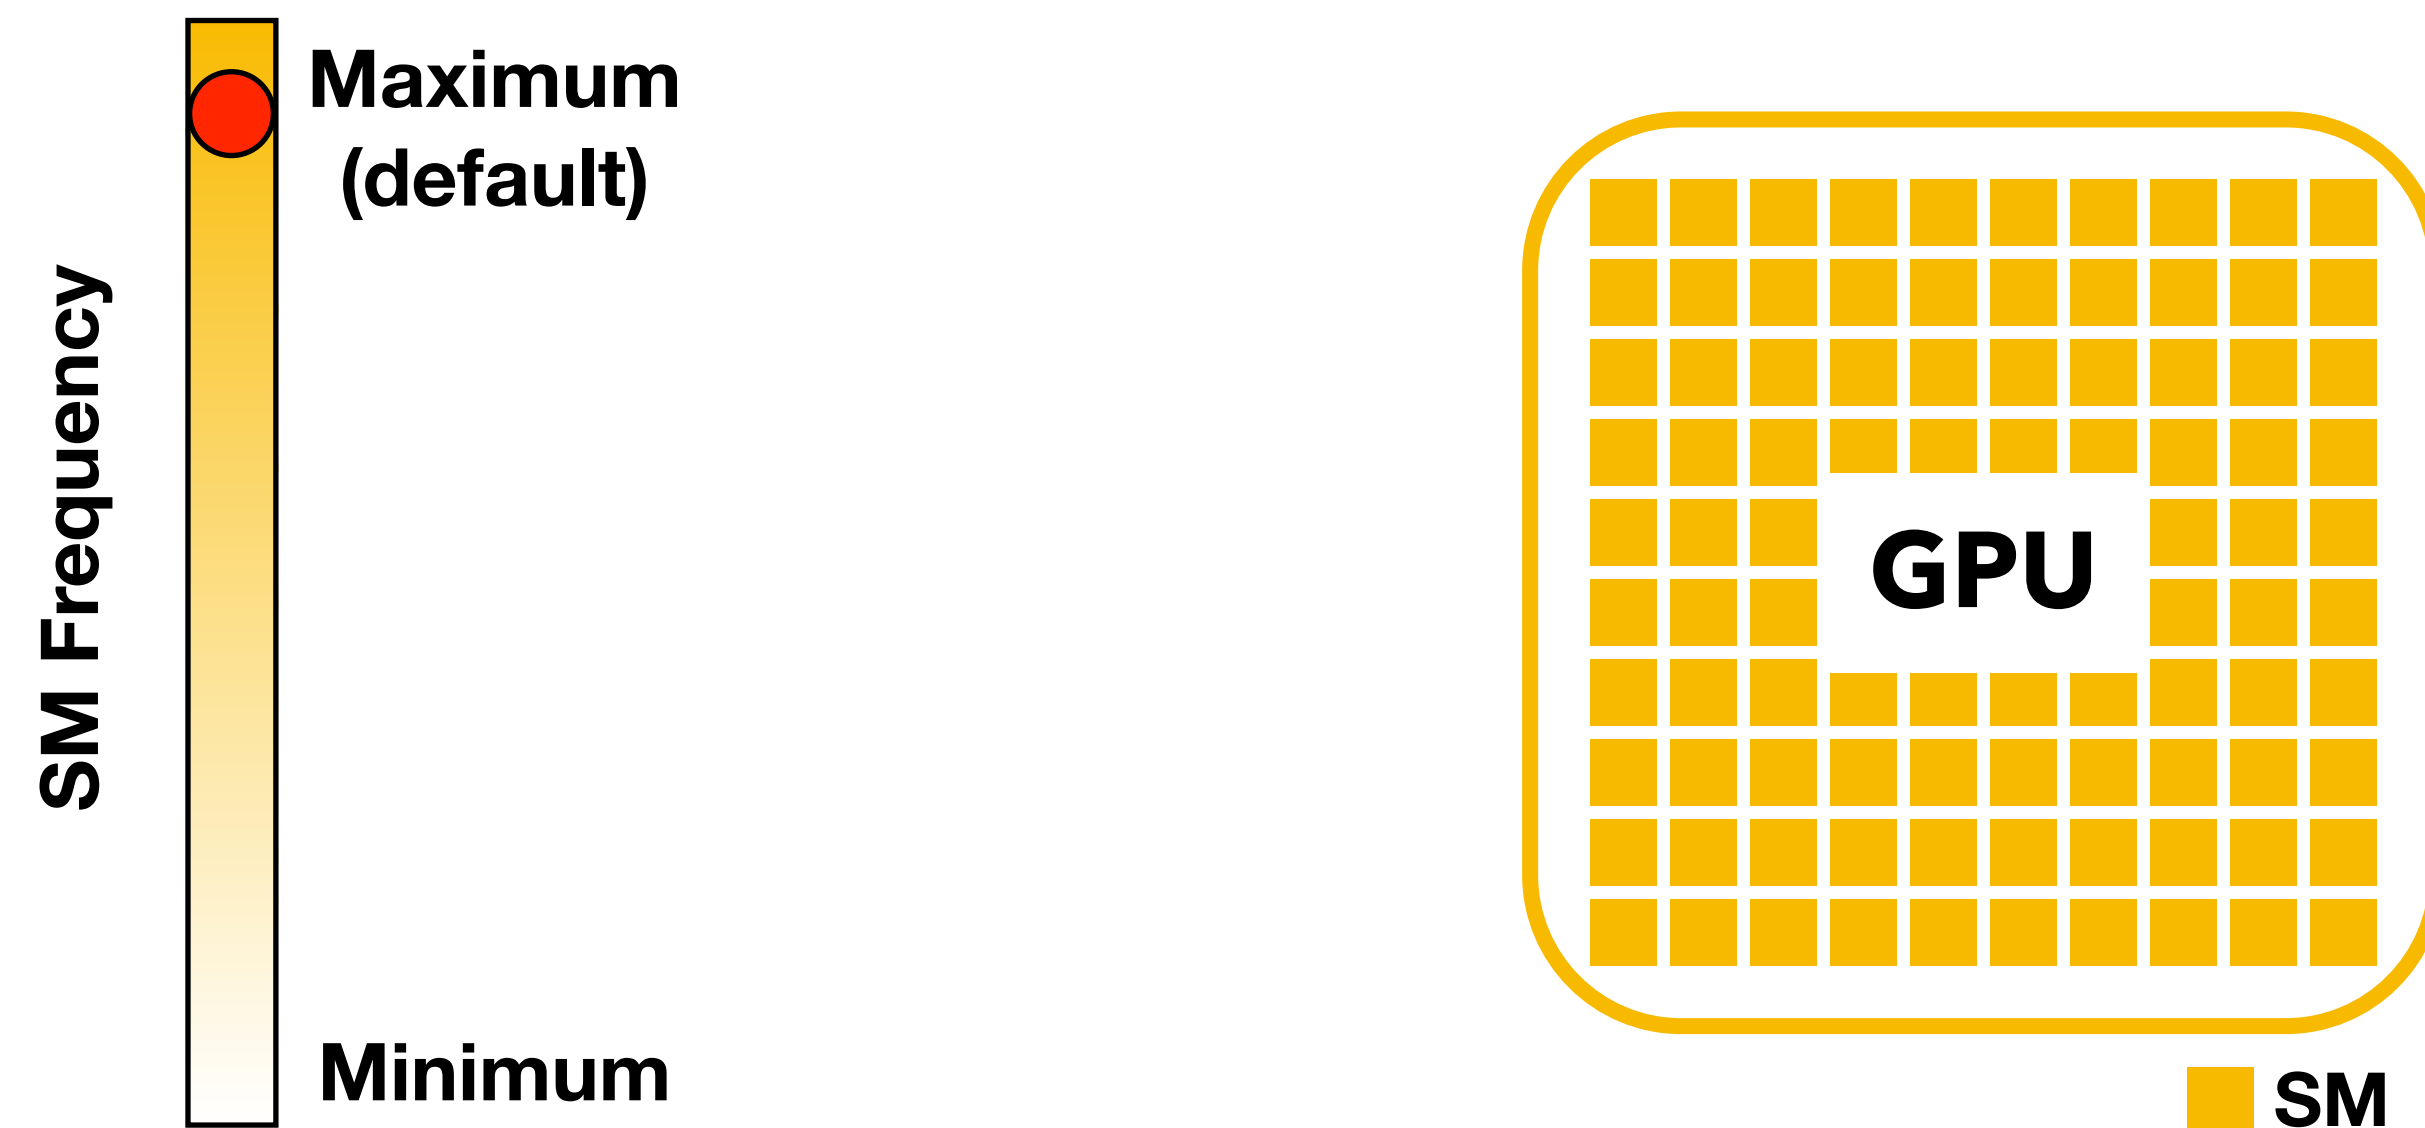
# Need for Sustainable AI



- Training GPT-3 takes 14.8 days with 10,000 V100 GPUs consuming **1,287MWh**<sup>[1]</sup>
- Enough energy to power approx. **44,183 US households** for one day<sup>[2]</sup>


# Prior Work to Save Energy in DNN Training

- **GPU DVFS (Dynamic Voltage and Frequency Scaling)**
  - Finding optimal streaming multiprocessor (SM) or memory clock frequency



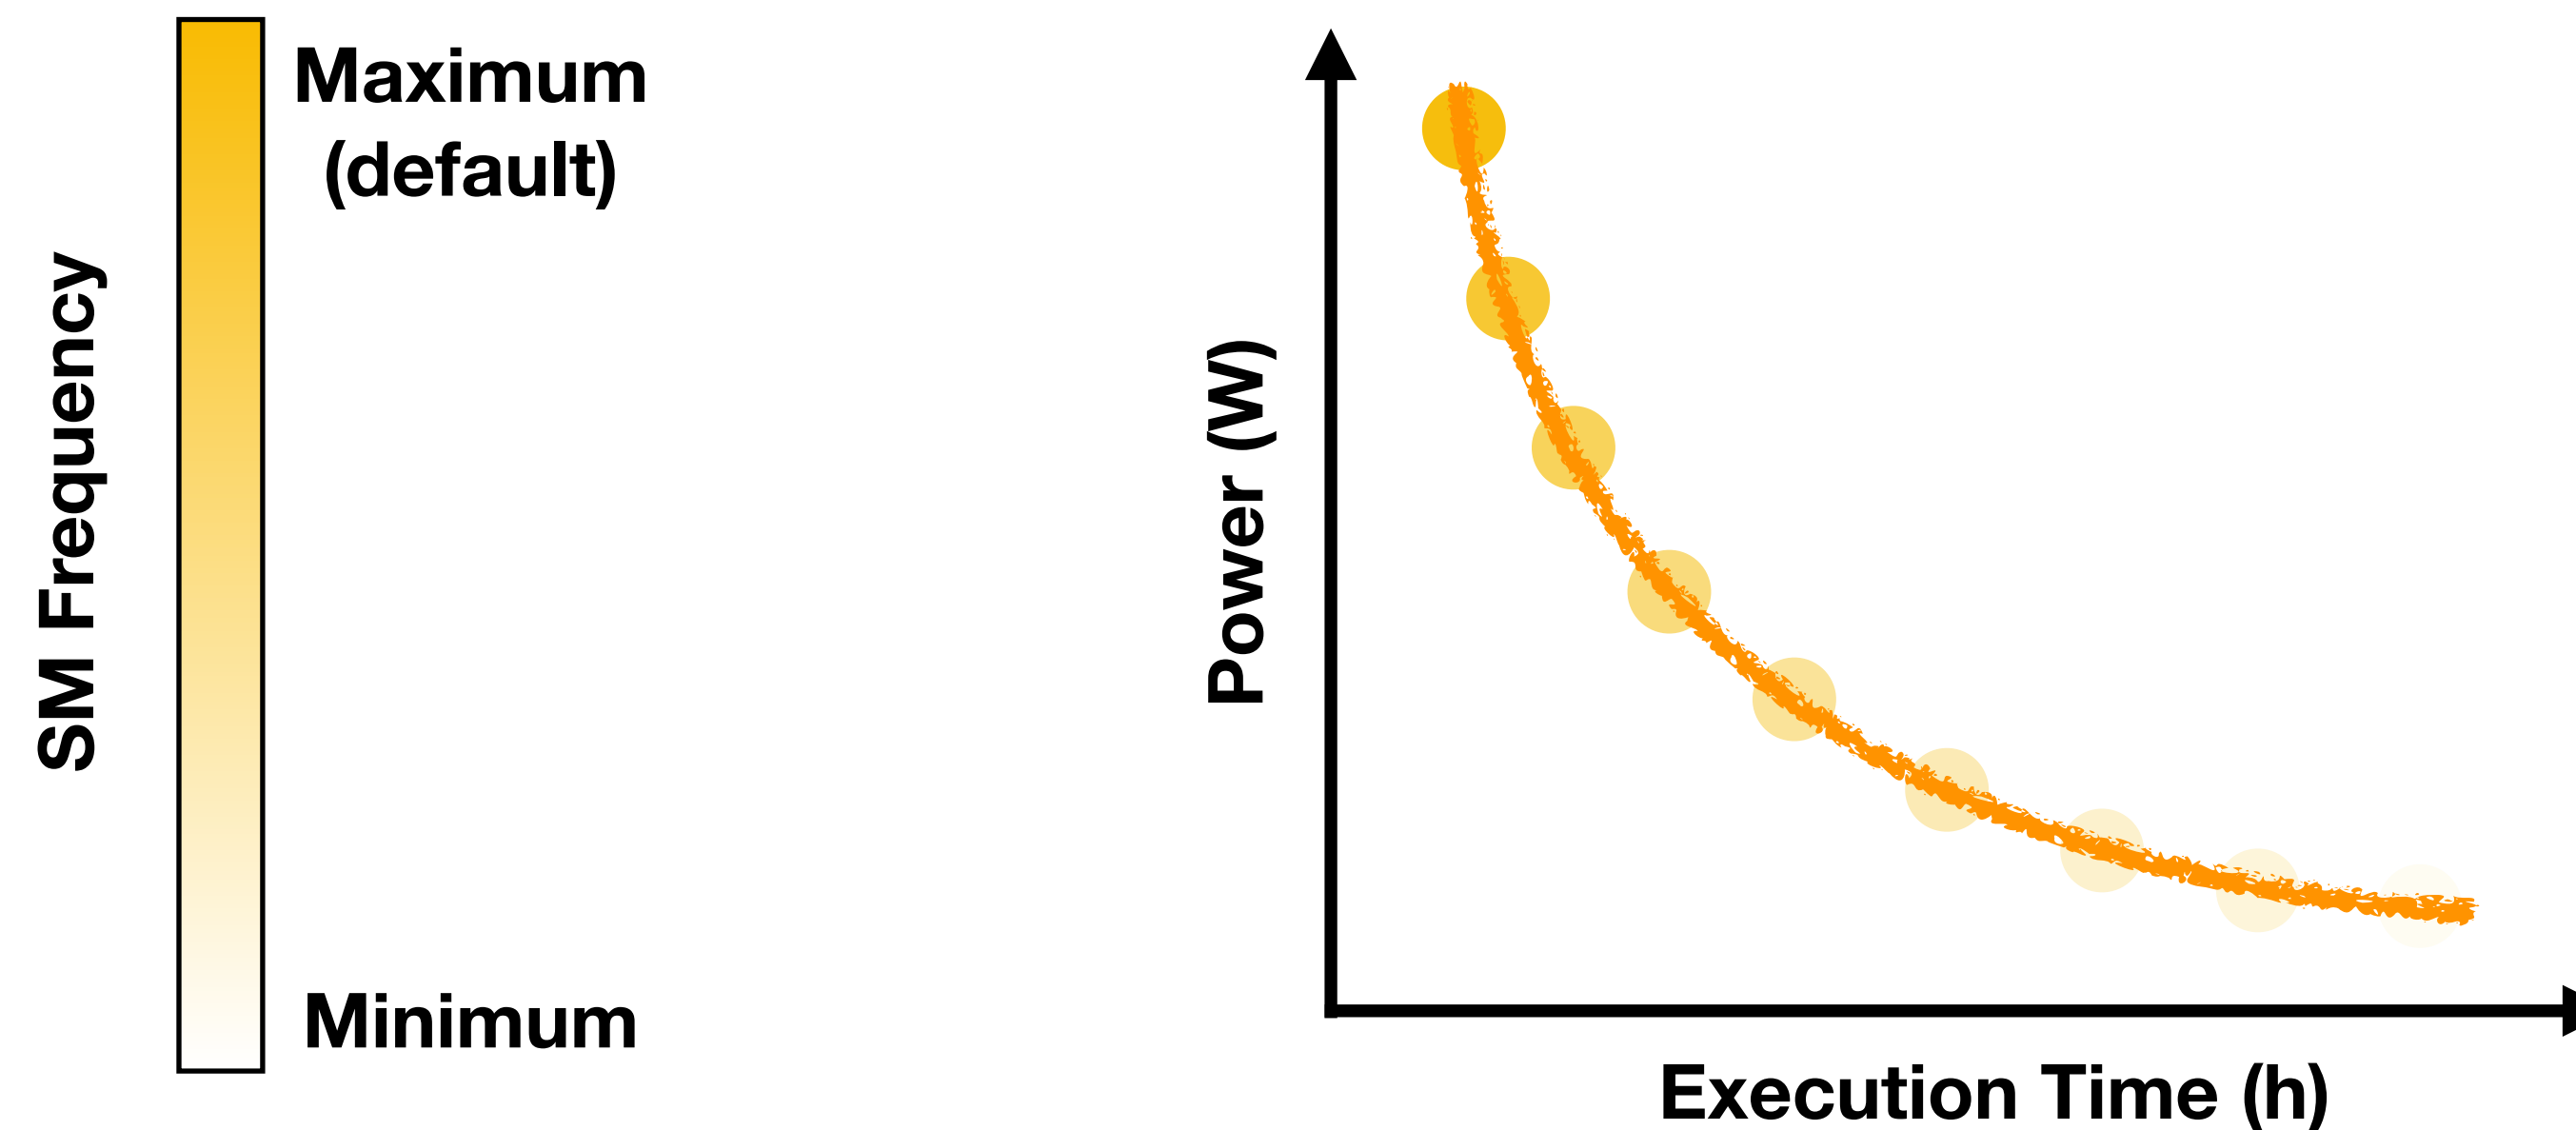
# Prior Work to Save Energy in DNN Training

- **GPU DVFS (Dynamic Voltage and Frequency Scaling)**
  - Finding optimal streaming multiprocessor (SM) or memory clock frequency



# Prior Work to Save Energy in DNN Training

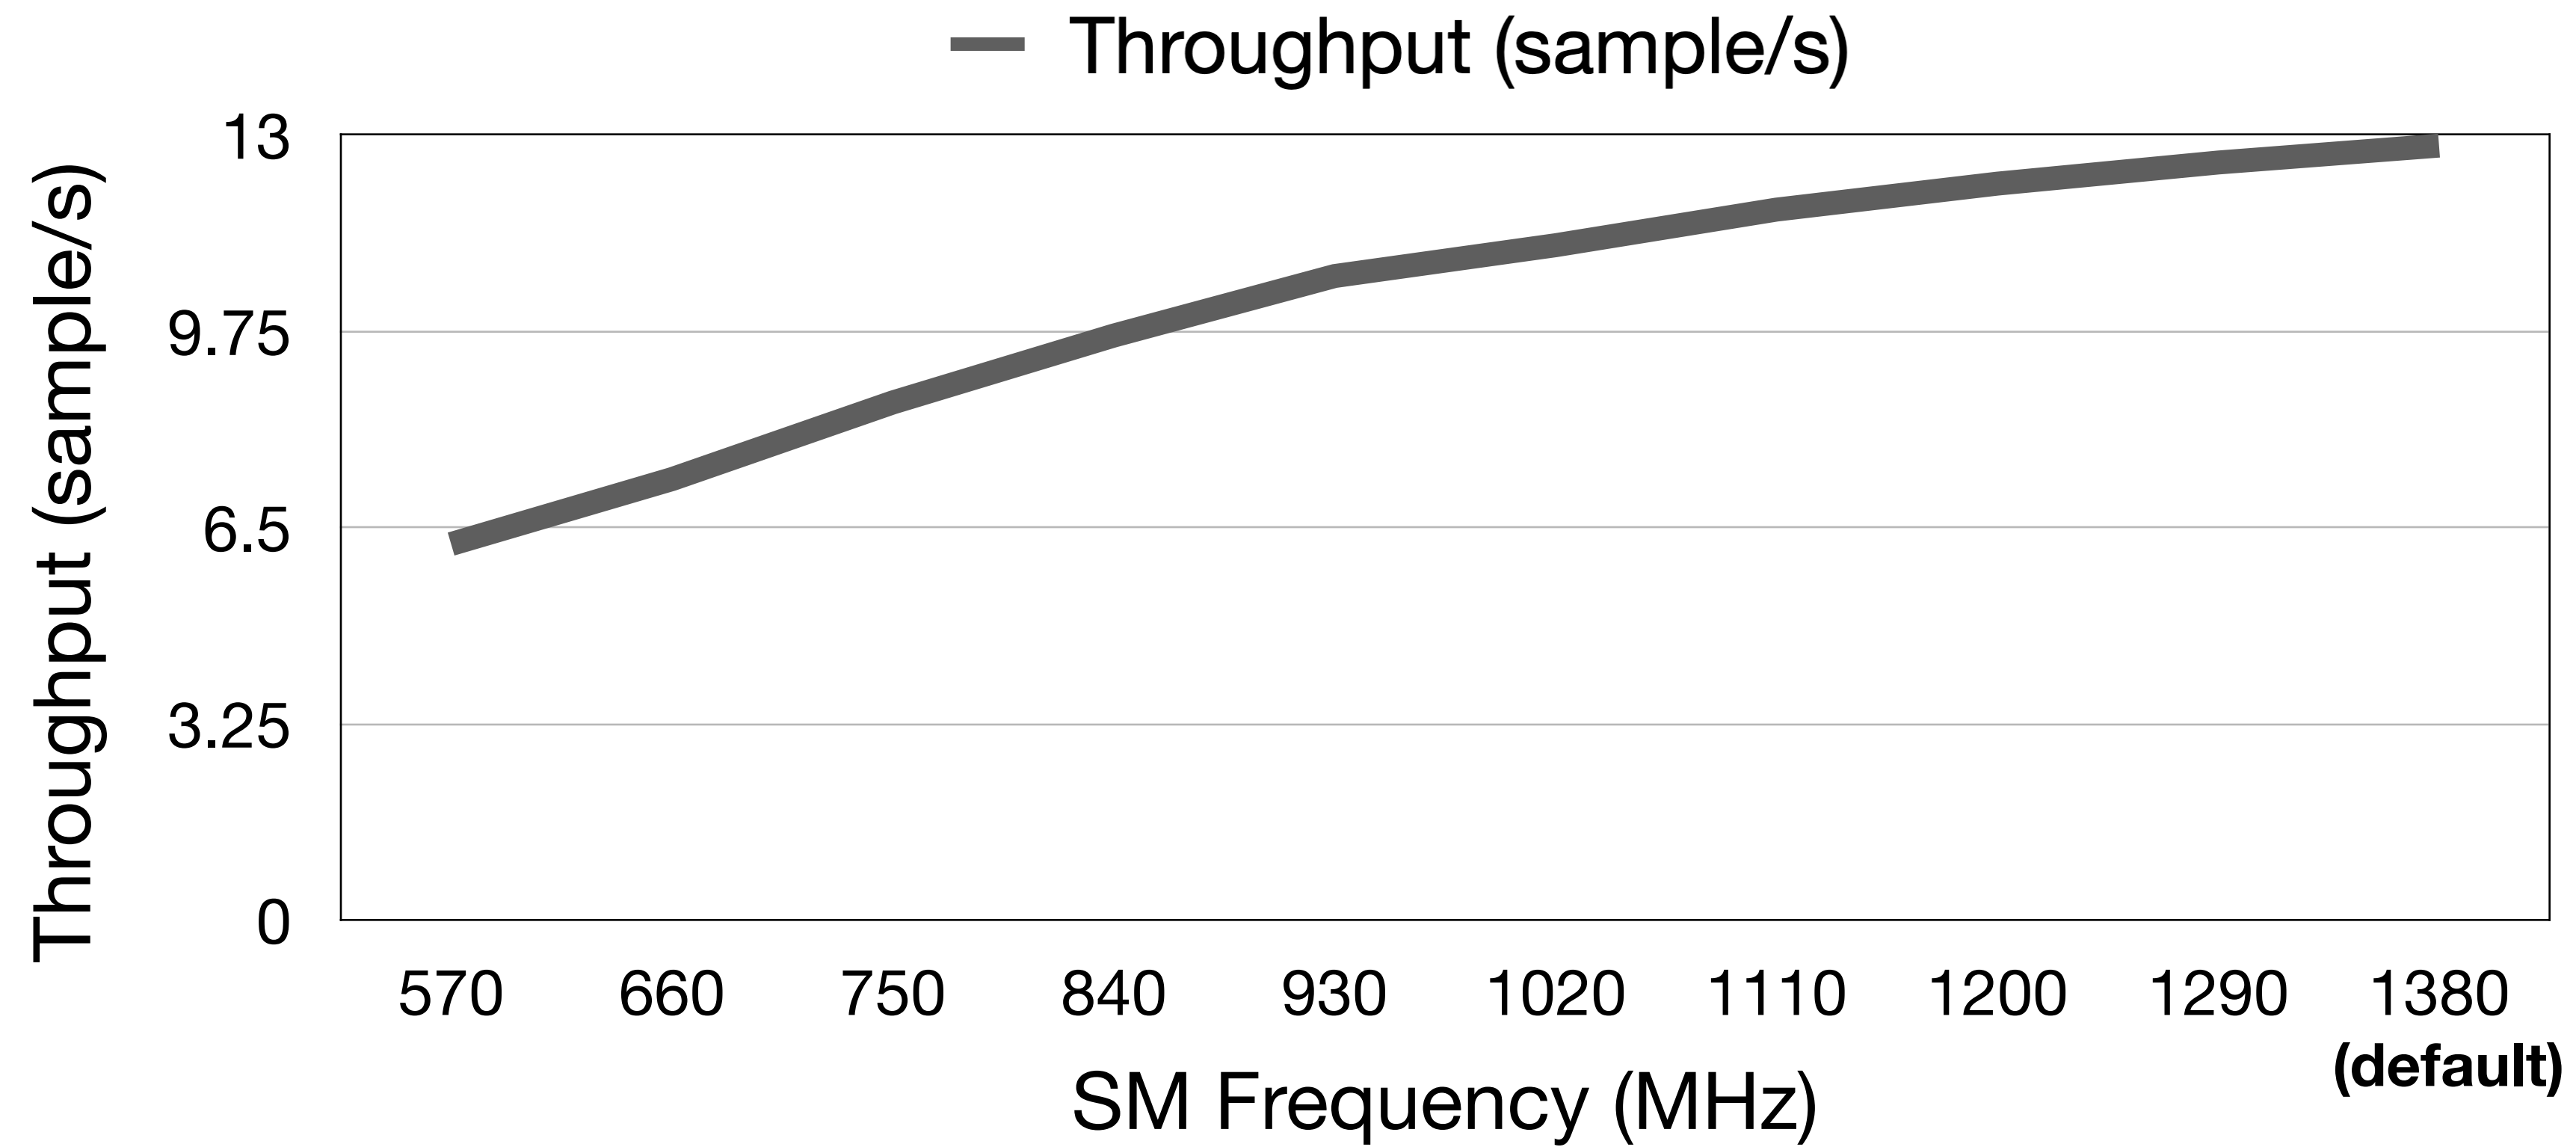
- **GPU DVFS (Dynamic Voltage and Frequency Scaling)**
  - Finding optimal streaming multiprocessor (SM) or memory clock frequency



**Reducing SM frequency saves energy but degrades performance**


# GPU DVFS: Energy Scaling Valley Trend

## Training Megatron-125M with V100



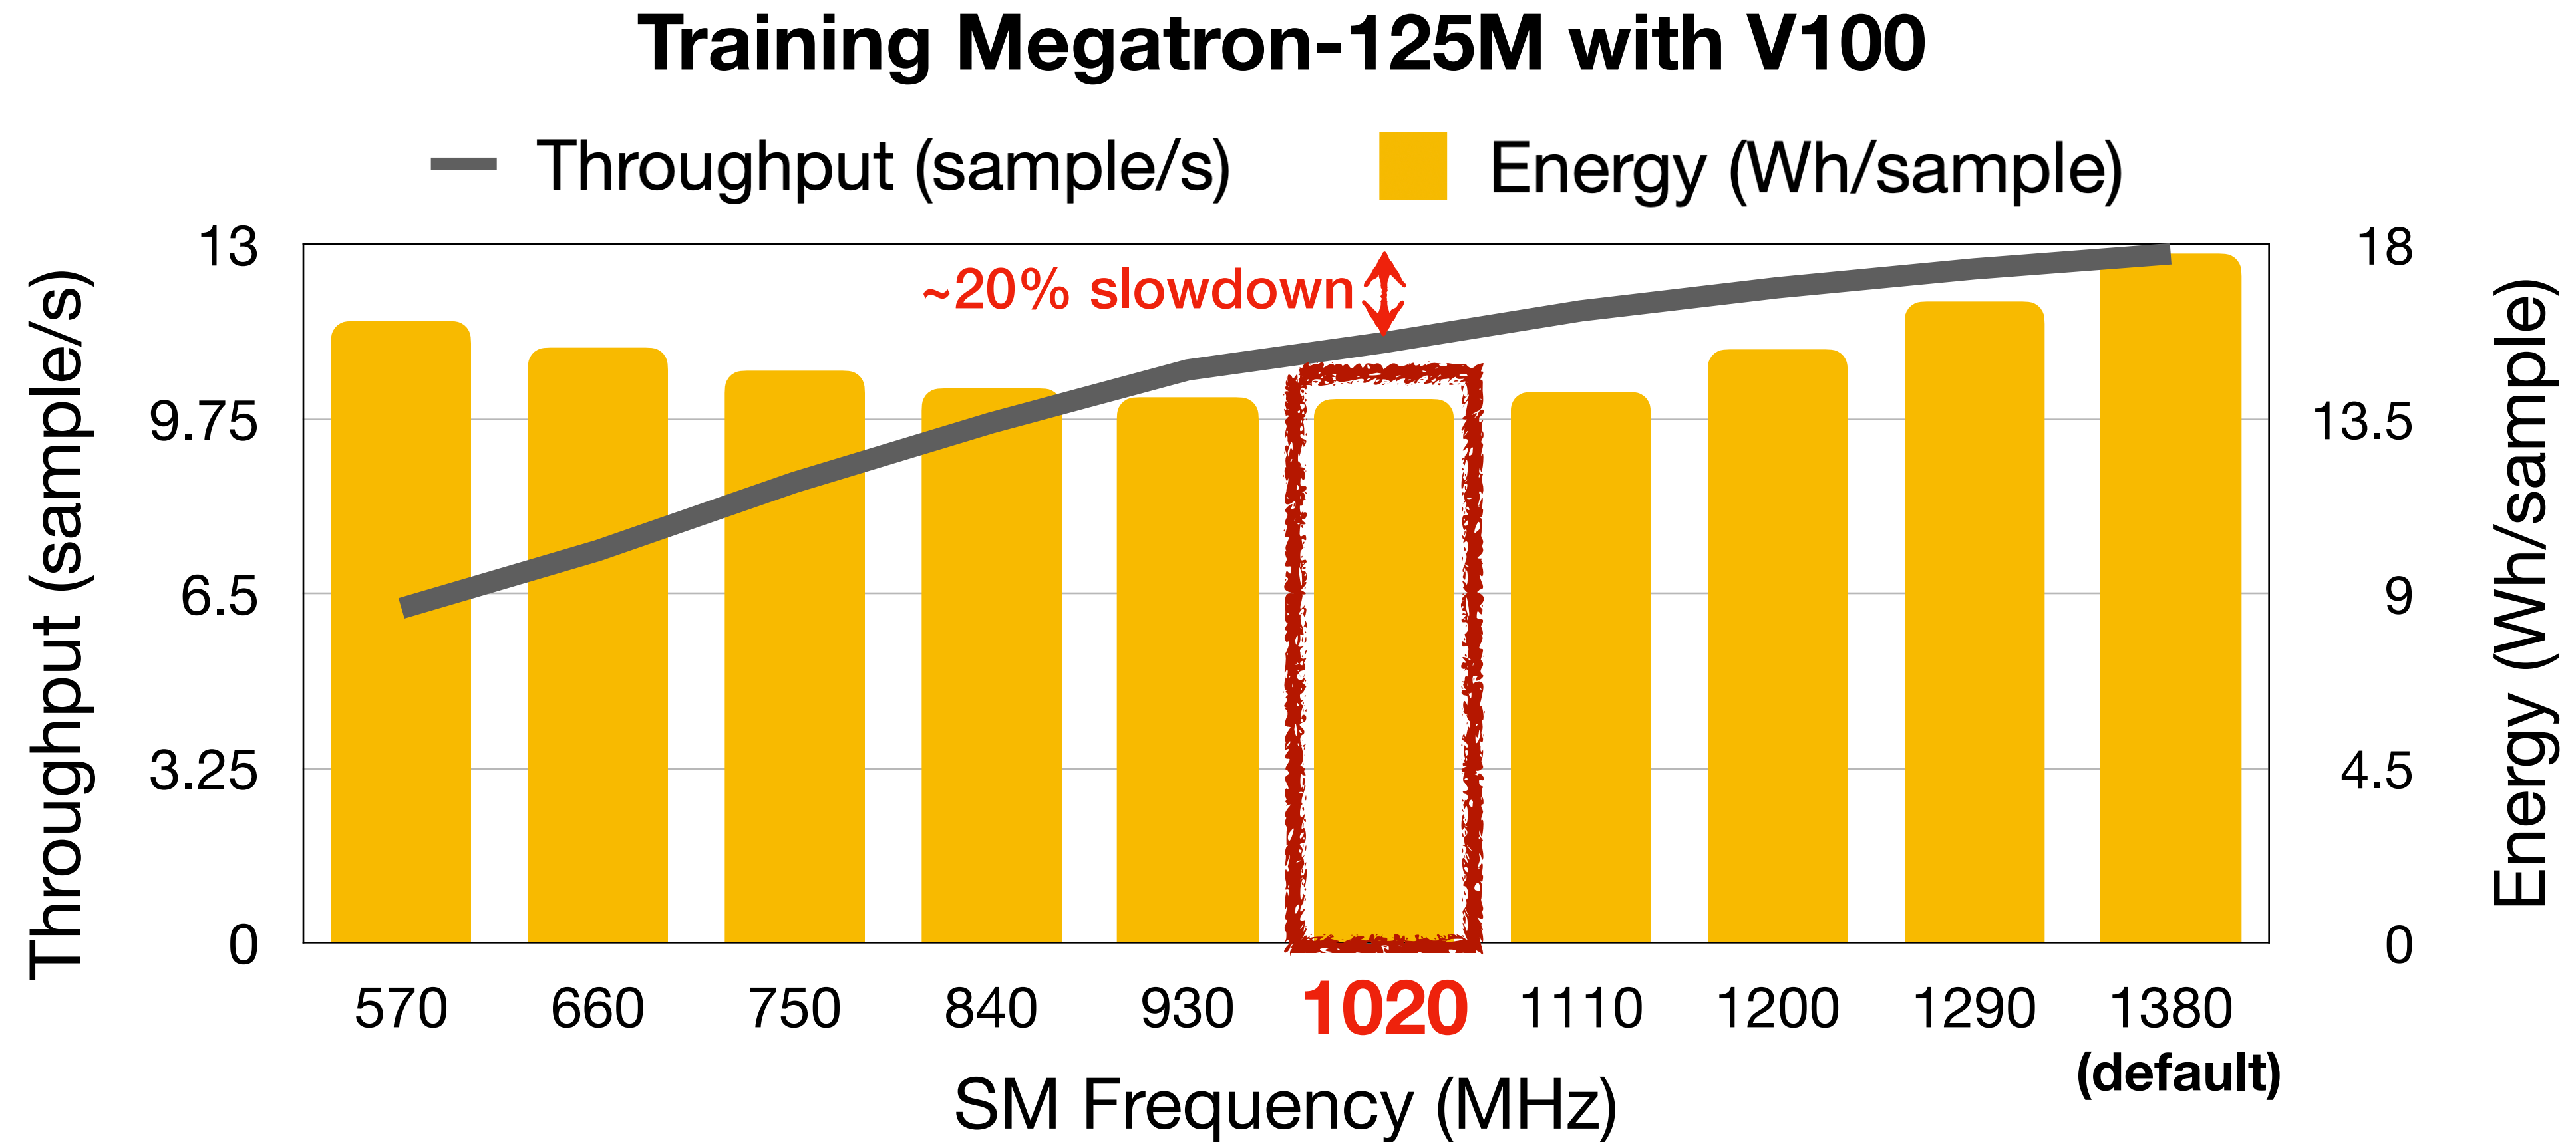


# GPU DVFS: Energy Scaling Valley Trend






# GPU DVFS: Energy Scaling Valley Trend



# GPU DVFS: Energy Scaling Valley Trend



Prior work focused on navigating the energy-performance tradeoff

# Problem: Hard to Sacrifice Perf. or Accuracy in Training

- **Performance:** how much slowdown is acceptable for long-running jobs?
  - 10% slowdown for 1 month training is 3 days
- **Accuracy:** the training goal for ML practitioners

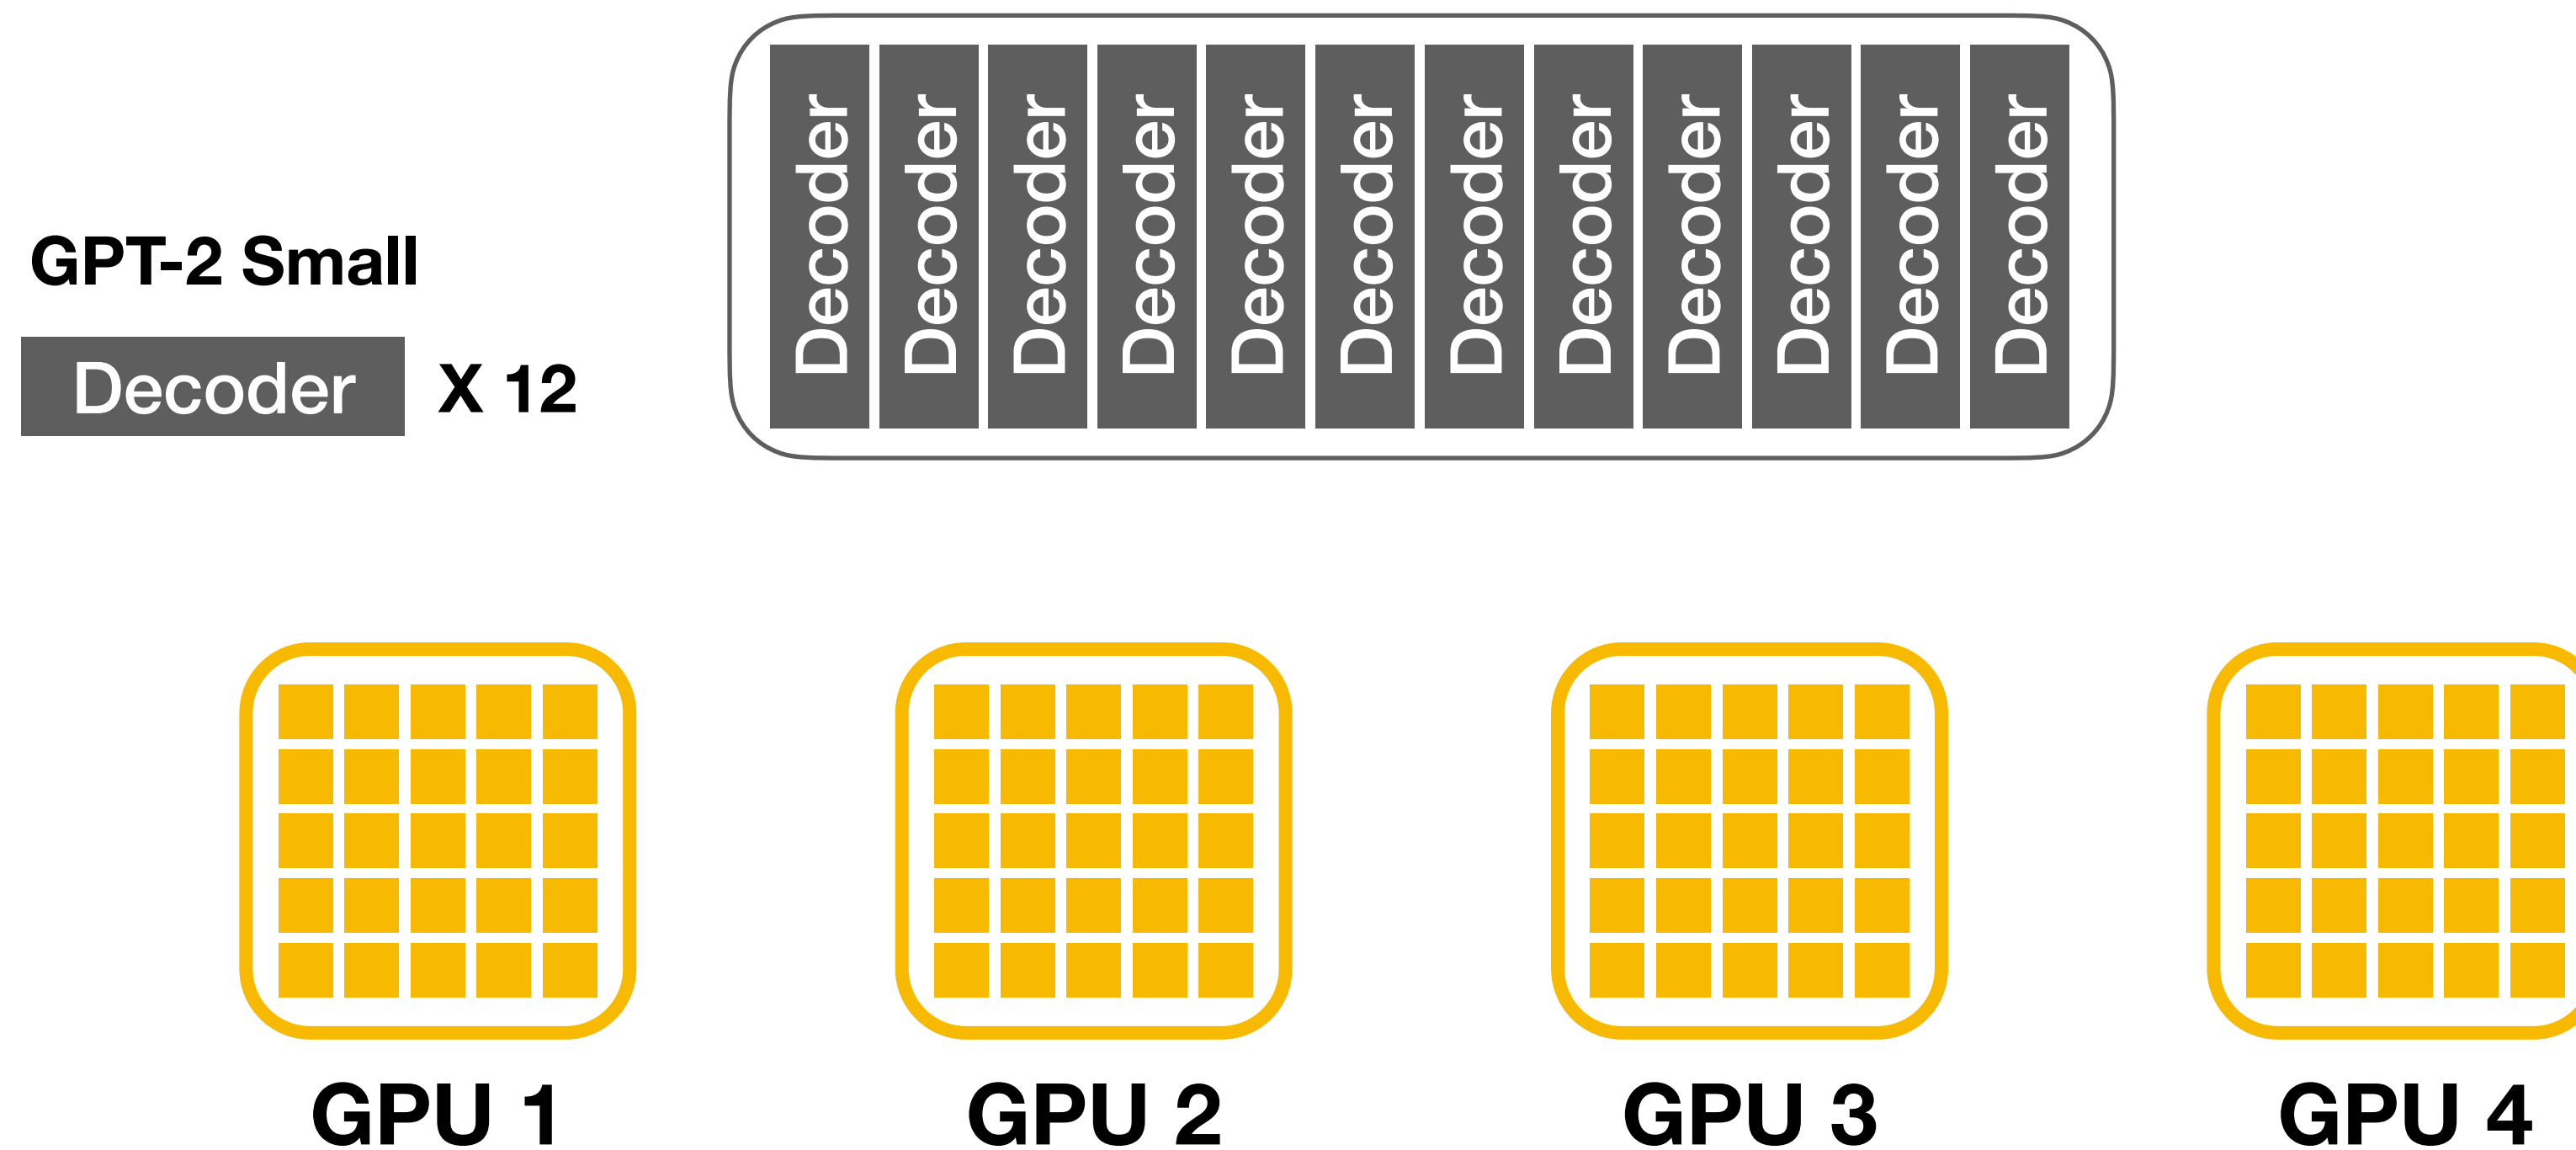
# Problem: Hard to Sacrifice Perf. or Accuracy in Training

- **Performance:** how much slowdown is acceptable for long-running jobs?
  - 10% slowdown for 1 month training is 3 days
- **Accuracy:** the training goal for ML practitioners

**Important to save energy without degrading any performance or accuracy**

# Opportunity: Bubbles in Pipeline Parallelism


- Model **vertically partitioned** to each GPU




# Opportunity: Bubbles in Pipeline Parallelism

- Model **vertically partitioned** to each GPU
- Splits input mini-batch into multiple micro-batches and pipelines execution


Mini-batch  
(Size 32)




GPU 1



GPU 2




GPU 3



GPU 4

# Opportunity: Bubbles in Pipeline Parallelism

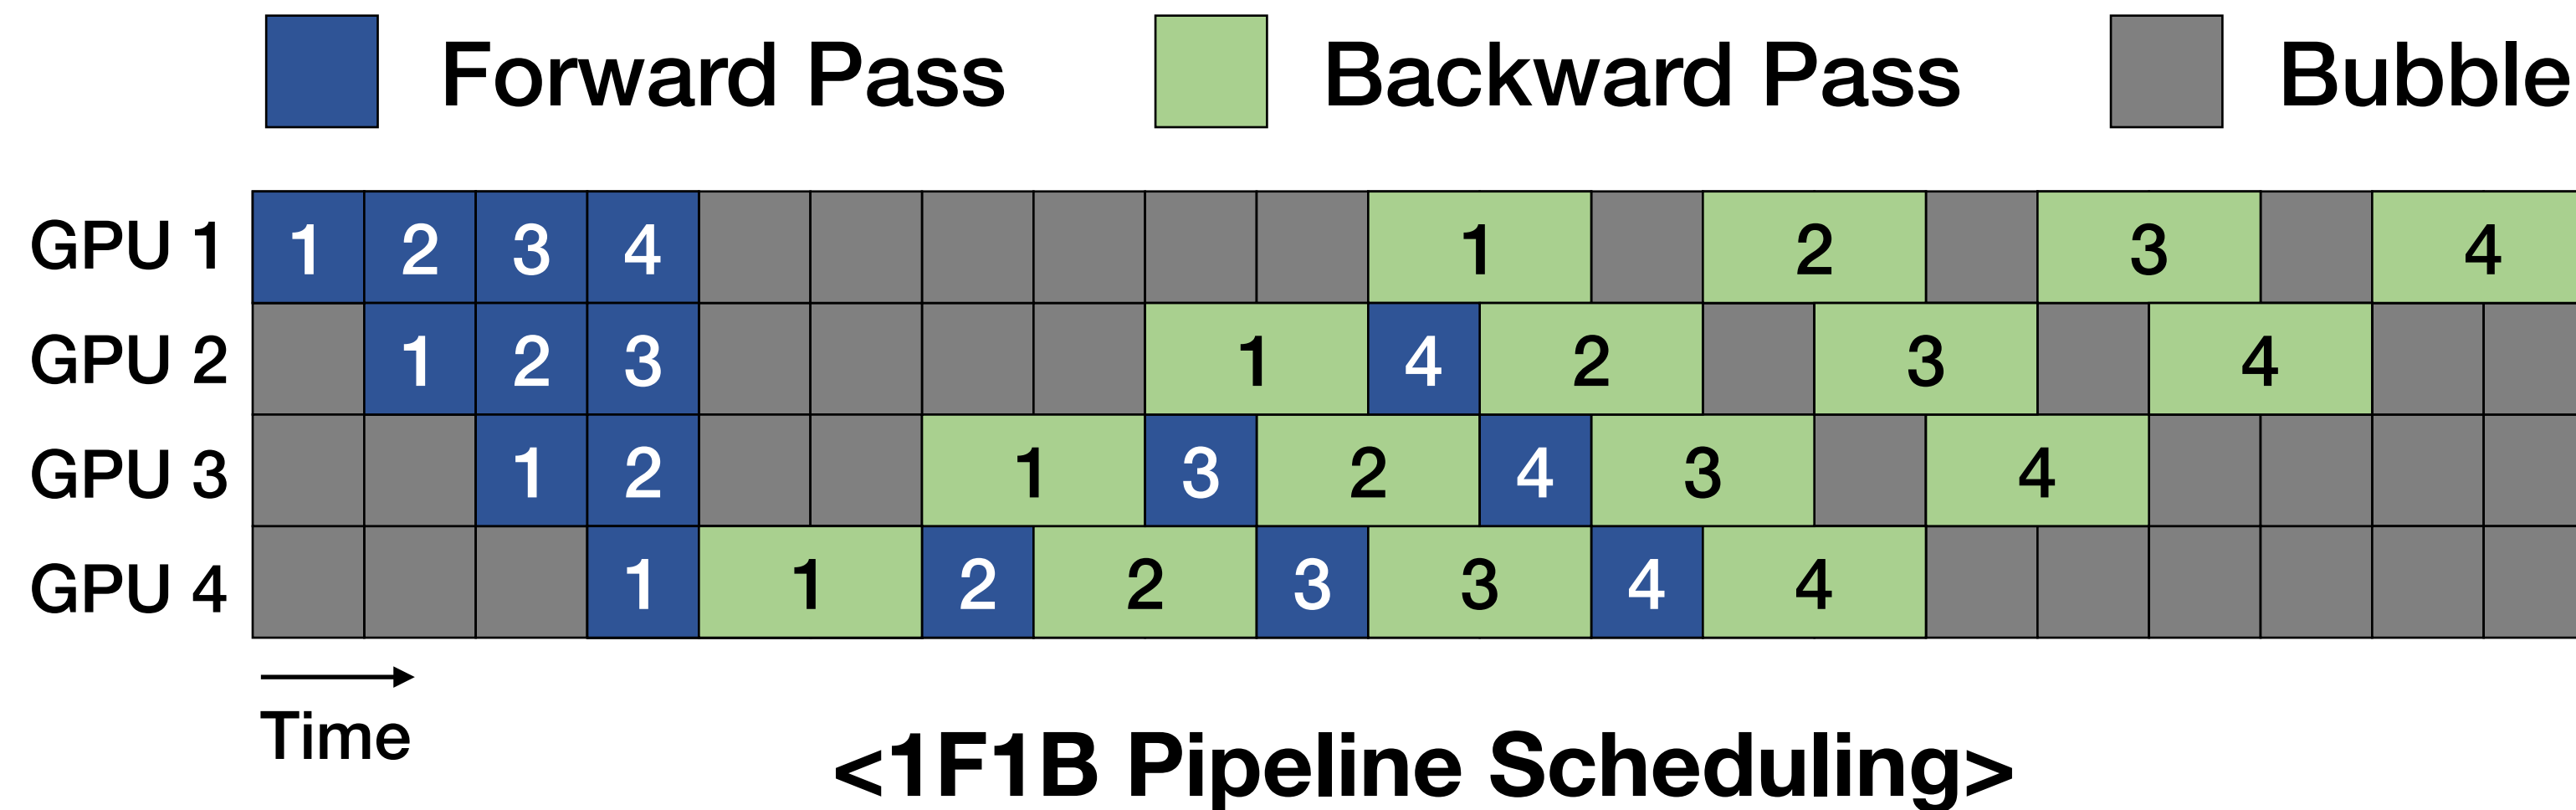
- Model **vertically partitioned** to each GPU
- Splits input mini-batch into multiple micro-batches and pipelines execution





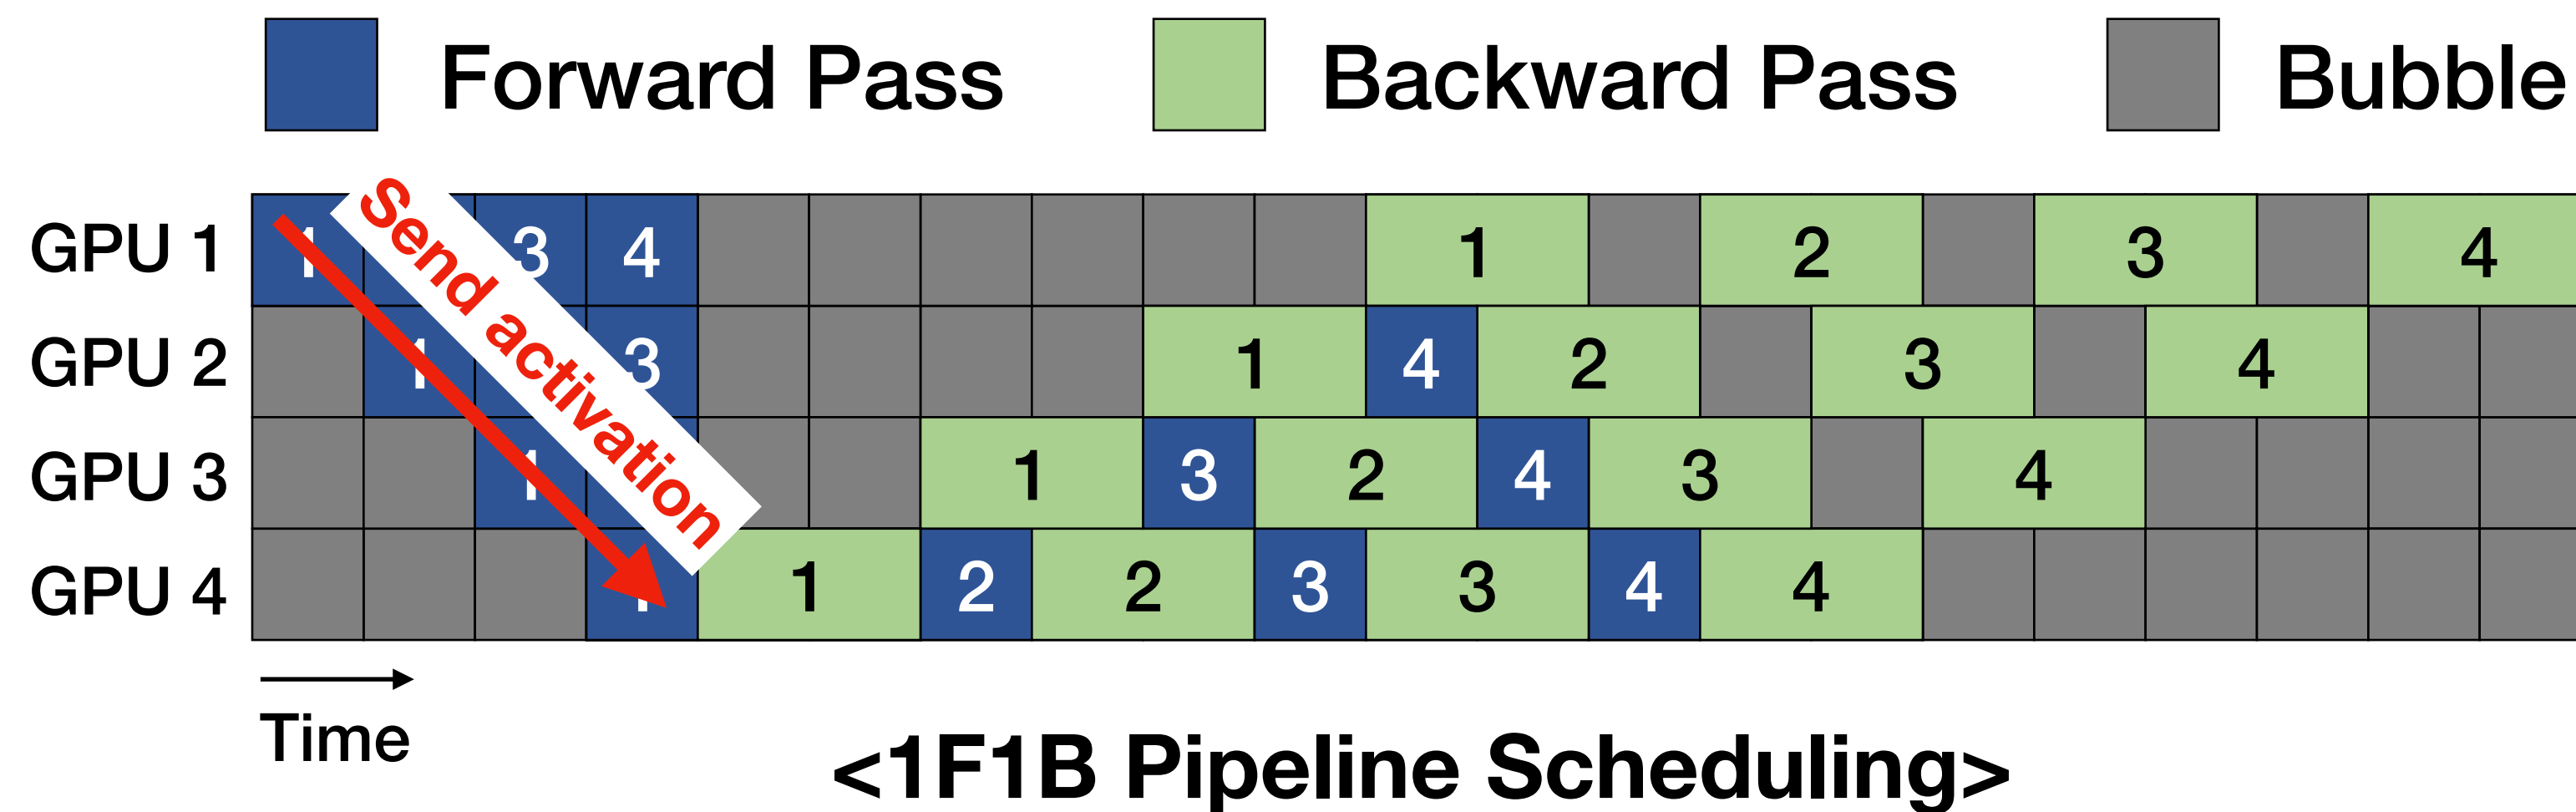
# Opportunity: Bubbles in Pipeline Parallelism

- Model **vertically partitioned** to each GPU
- Splits input mini-batch into multiple micro-batches and pipelines execution



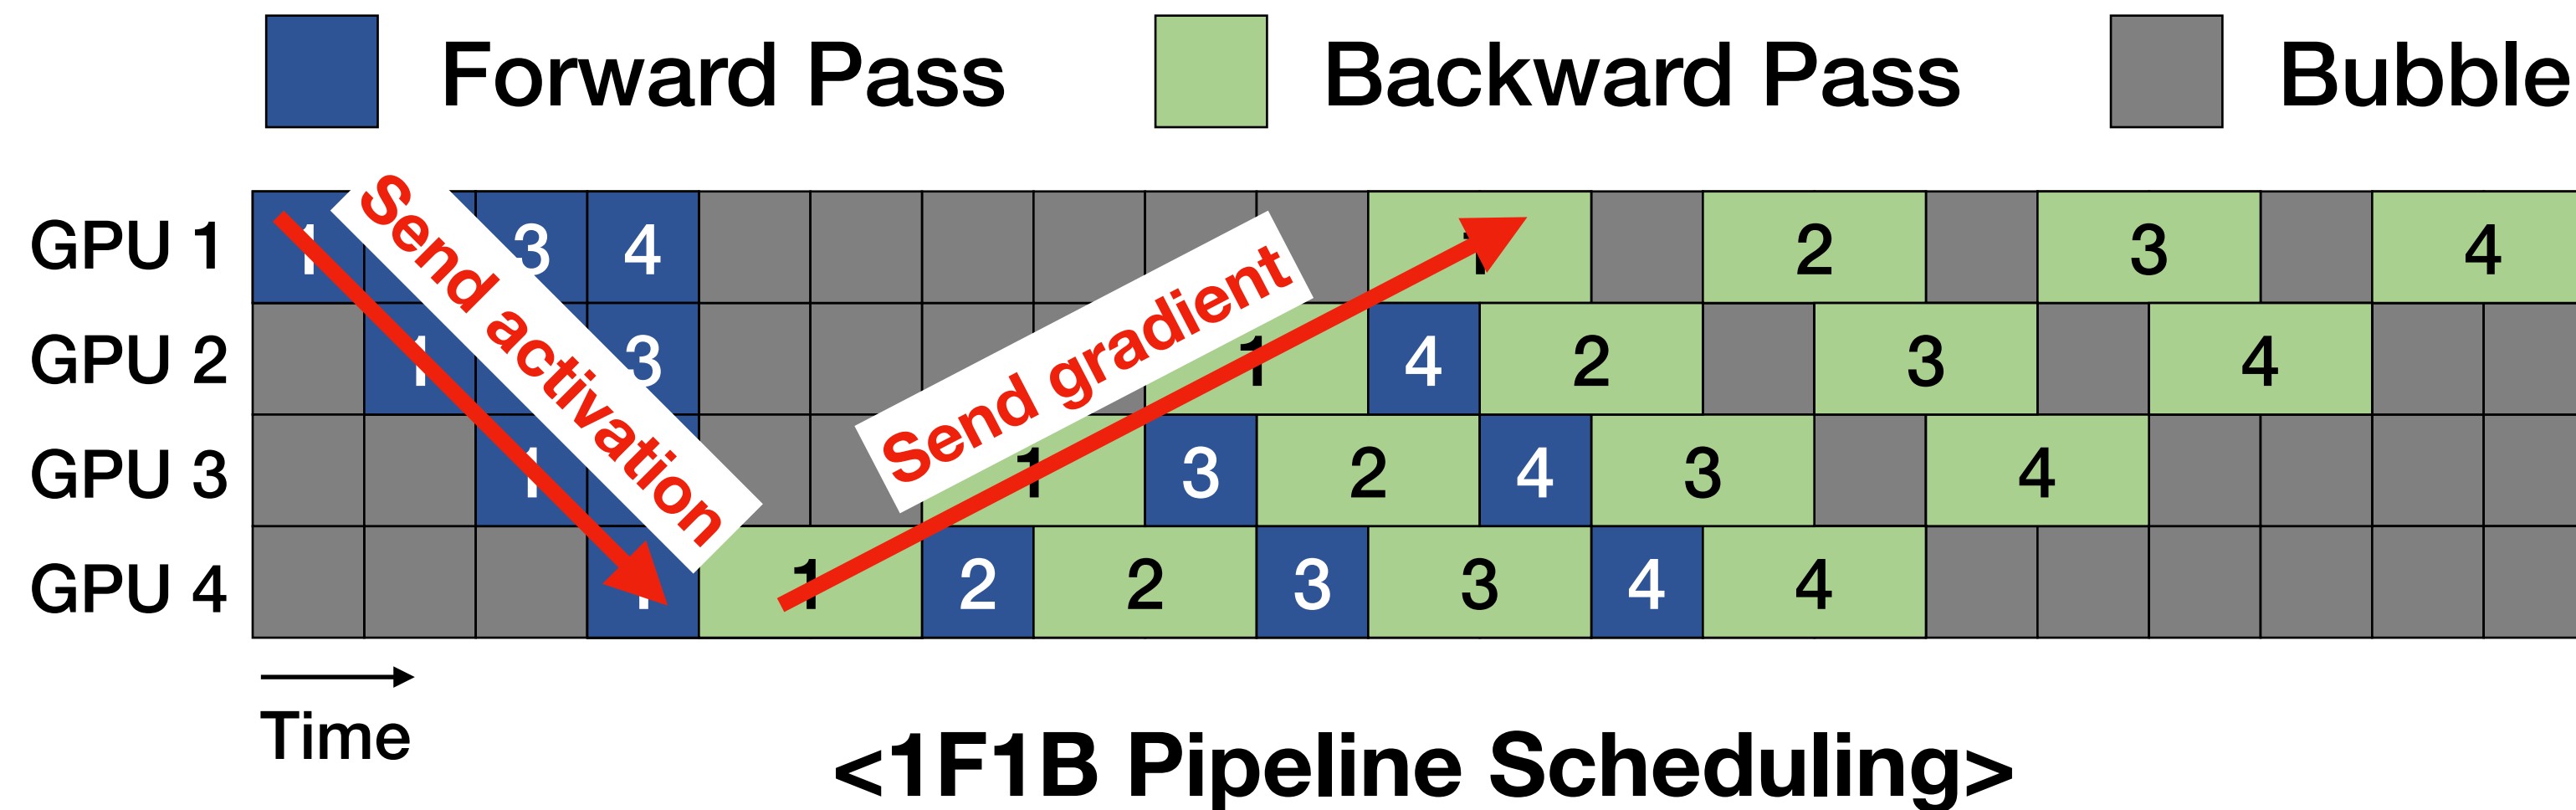
# Opportunity: Bubbles in Pipeline Parallelism

- Model **vertically partitioned** to each GPU
- Splits input mini-batch into multiple micro-batches and pipelines execution



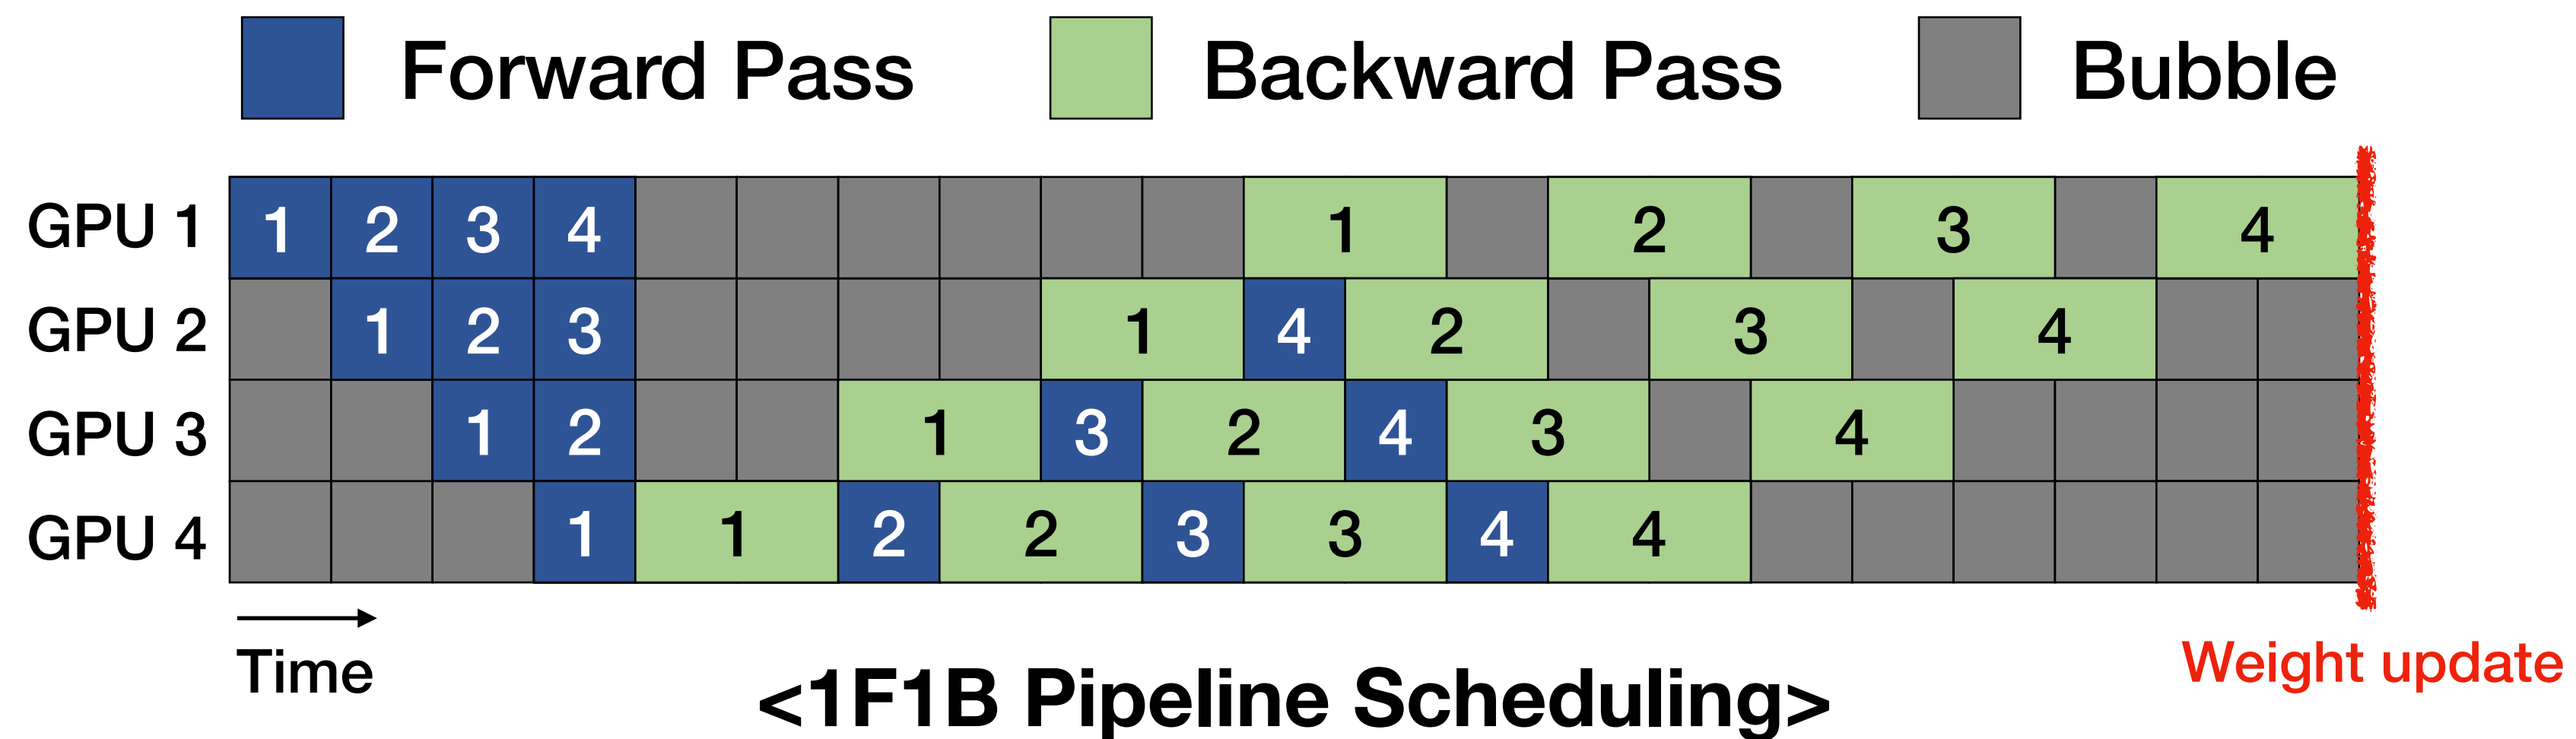
# Opportunity: Bubbles in Pipeline Parallelism

- Model **vertically partitioned** to each GPU
- Splits input mini-batch into multiple micro-batches and pipelines execution



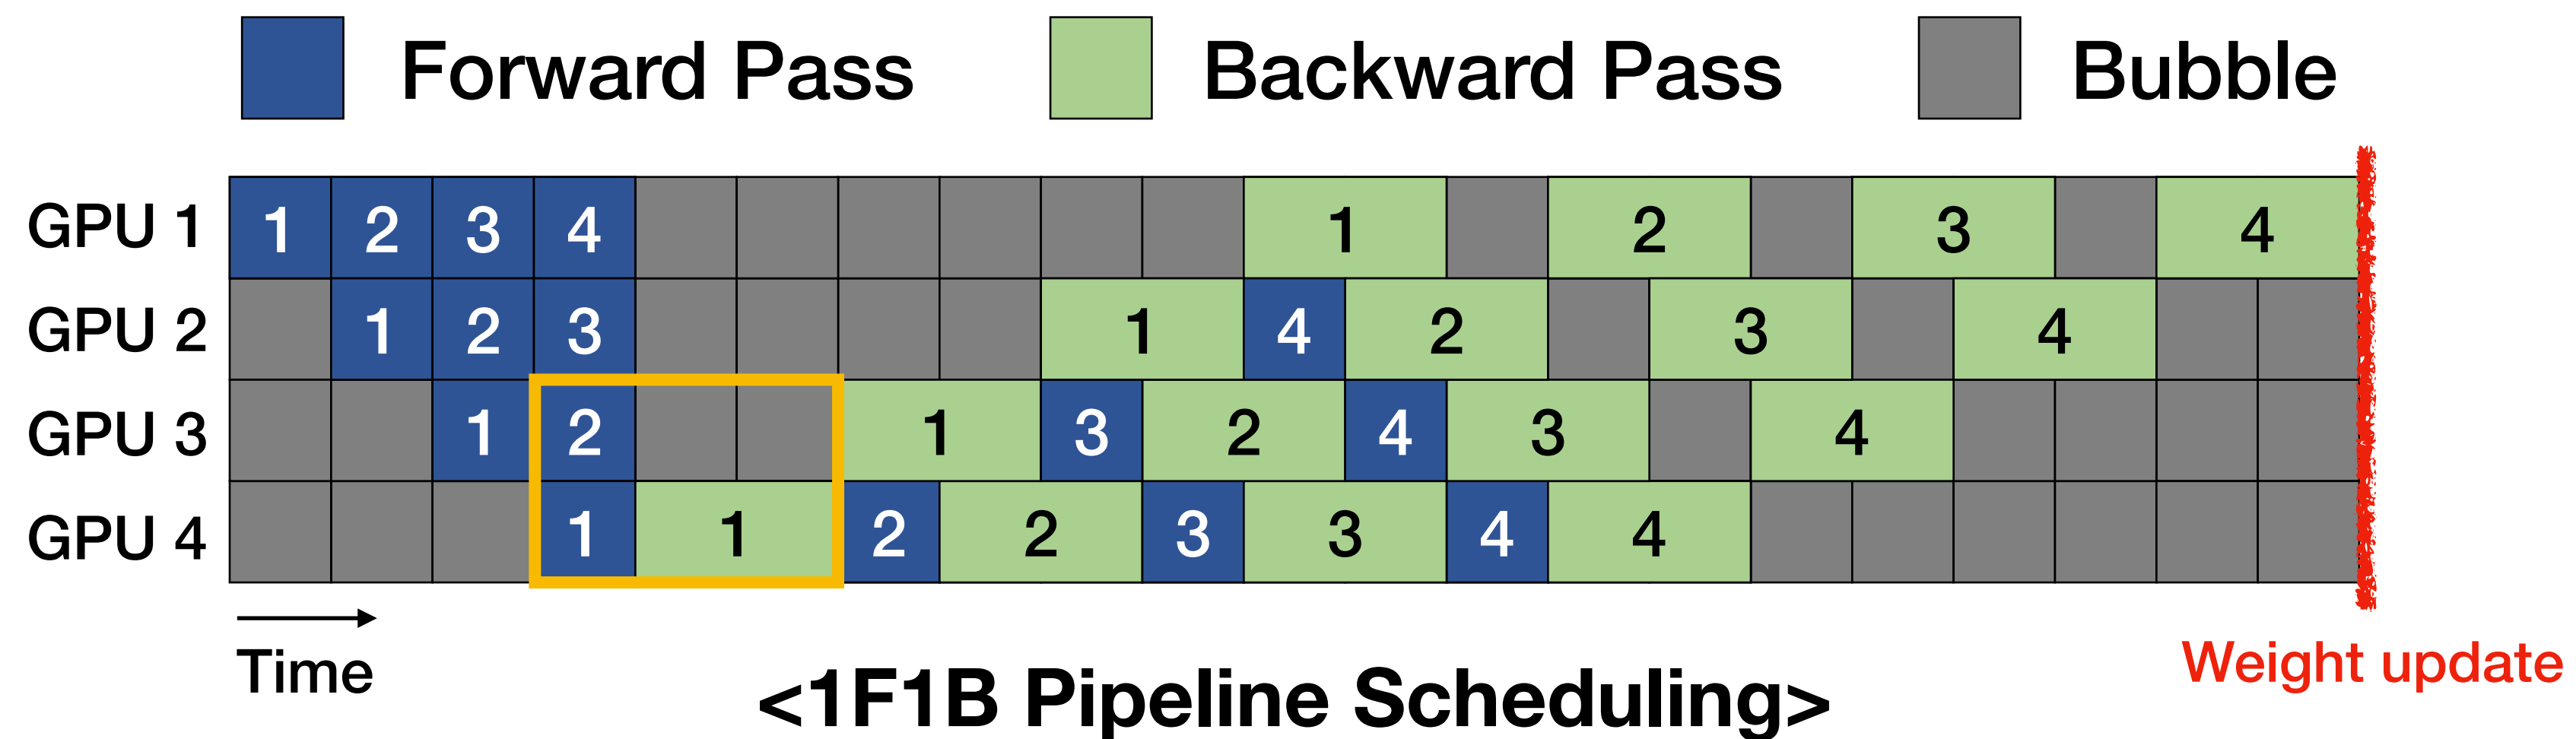
# Opportunity: Bubbles in Pipeline Parallelism

- Model **vertically partitioned** to each GPU
- Splits input mini-batch into multiple micro-batches and pipelines execution
- Weight update introduces **inevitable pipeline bubbles**




# Opportunity: Bubbles in Pipeline Parallelism

- Model **vertically partitioned** to each GPU
- Splits input mini-batch into multiple micro-batches and pipelines execution
- Weight update introduces **inevitable pipeline bubbles**




# Opportunity: Bubbles in Pipeline Parallelism



■ Forward Pass    ■ Backward Pass    ■ Bubble



<1F1B Pipeline Scheduling>



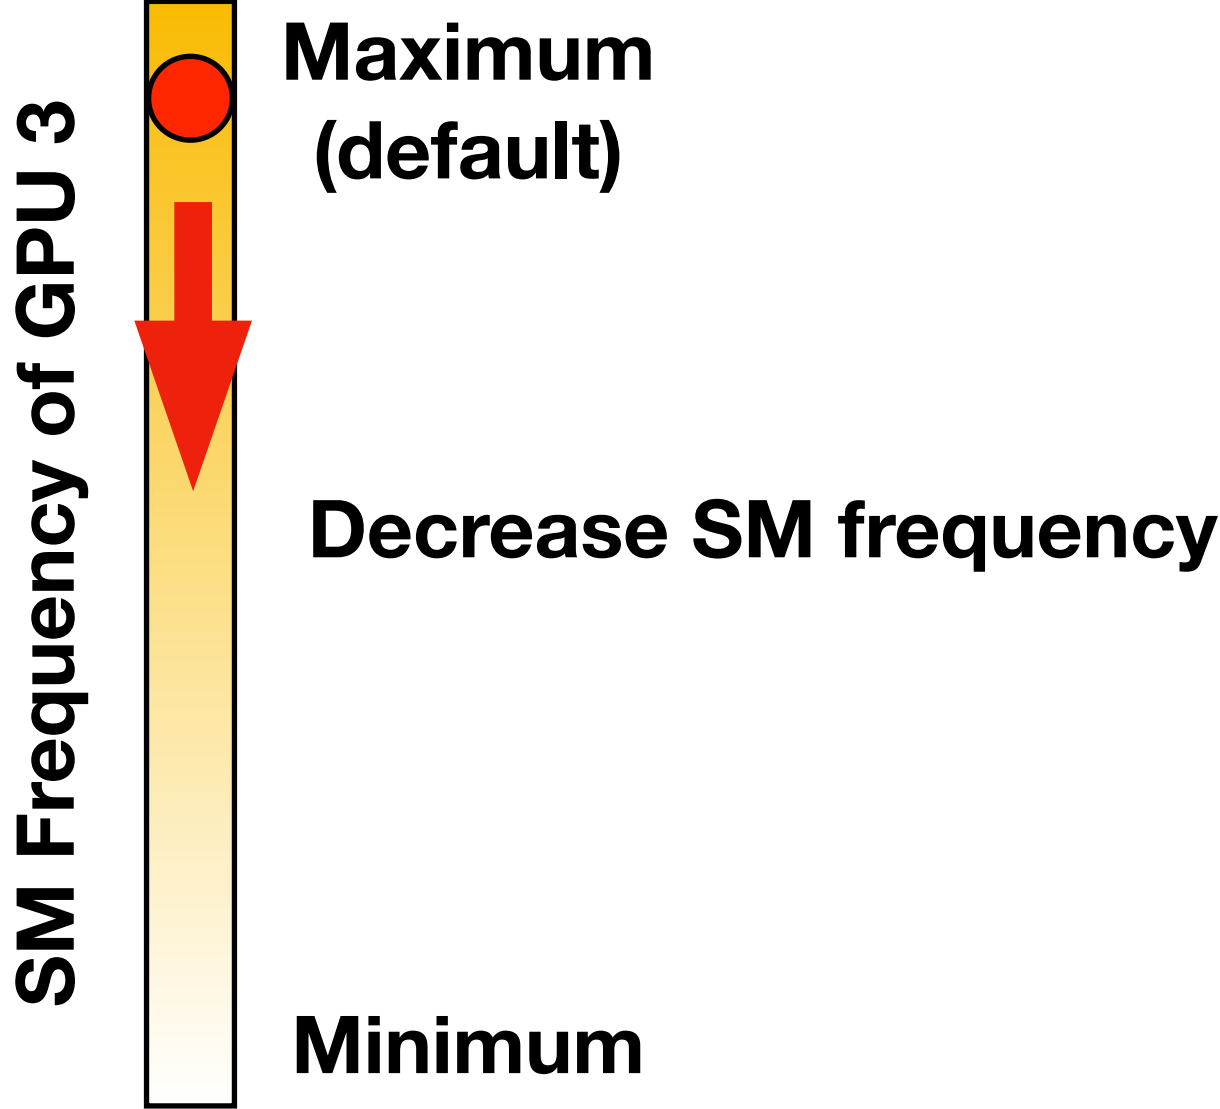
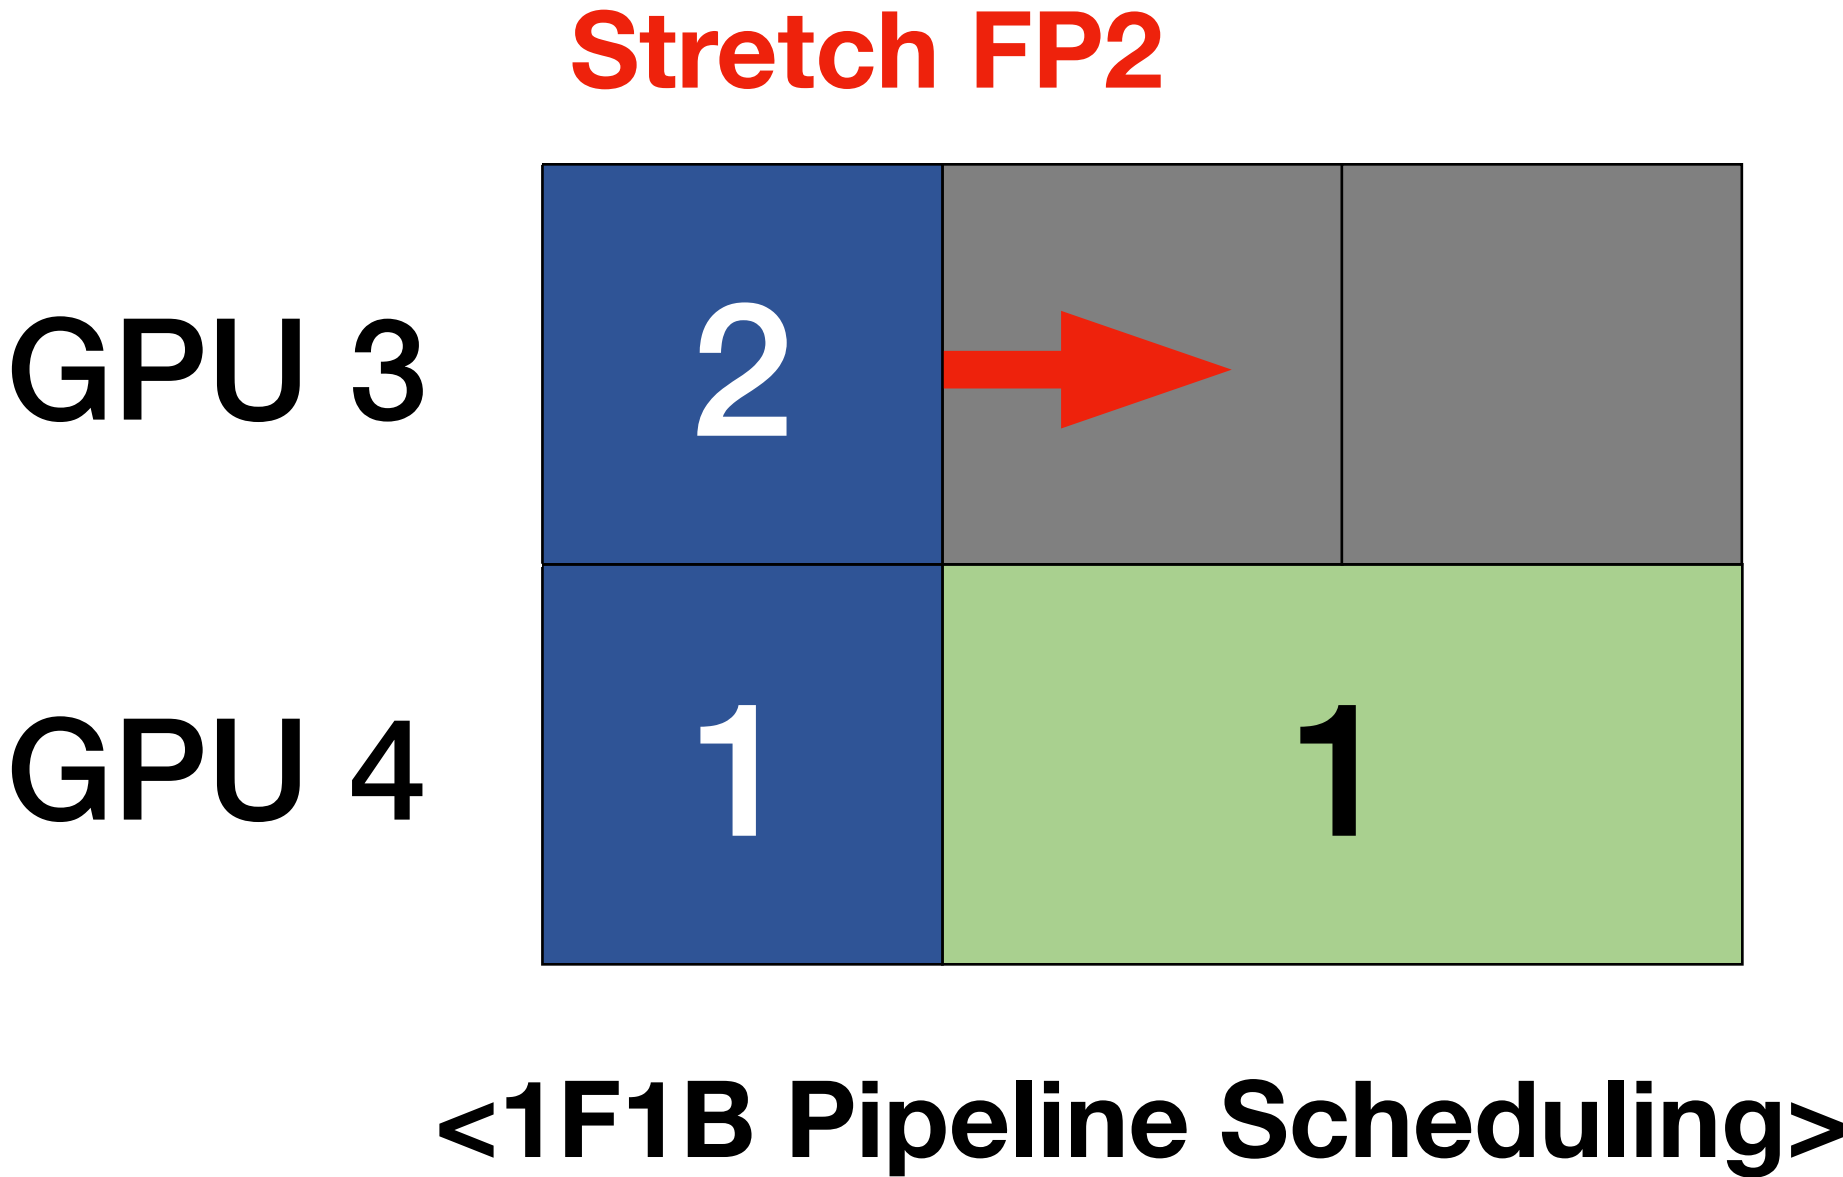
# Opportunity: Bubbles in Pipeline Parallelism





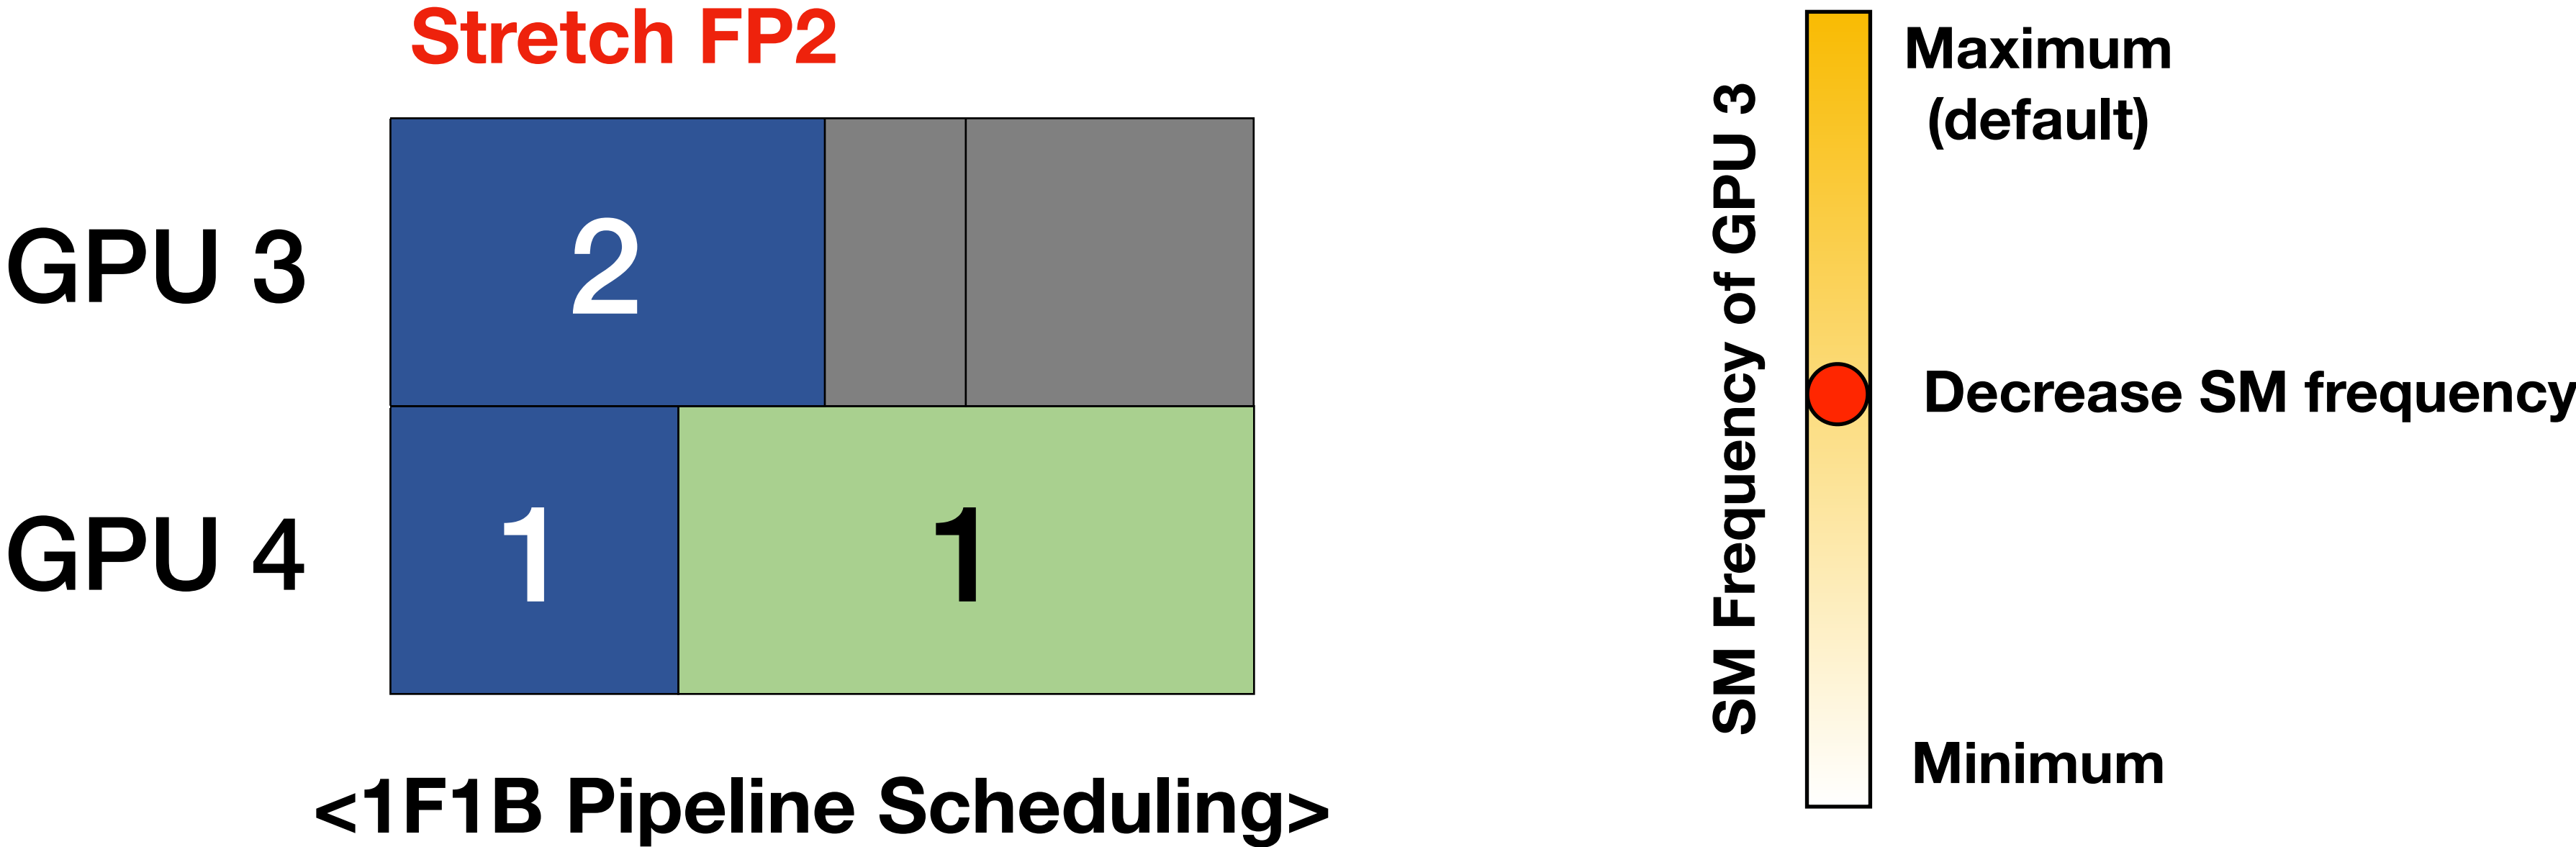
# Opportunity: Bubbles in Pipeline Parallelism

■ Forward Pass    ■ Backward Pass    ■ Bubble

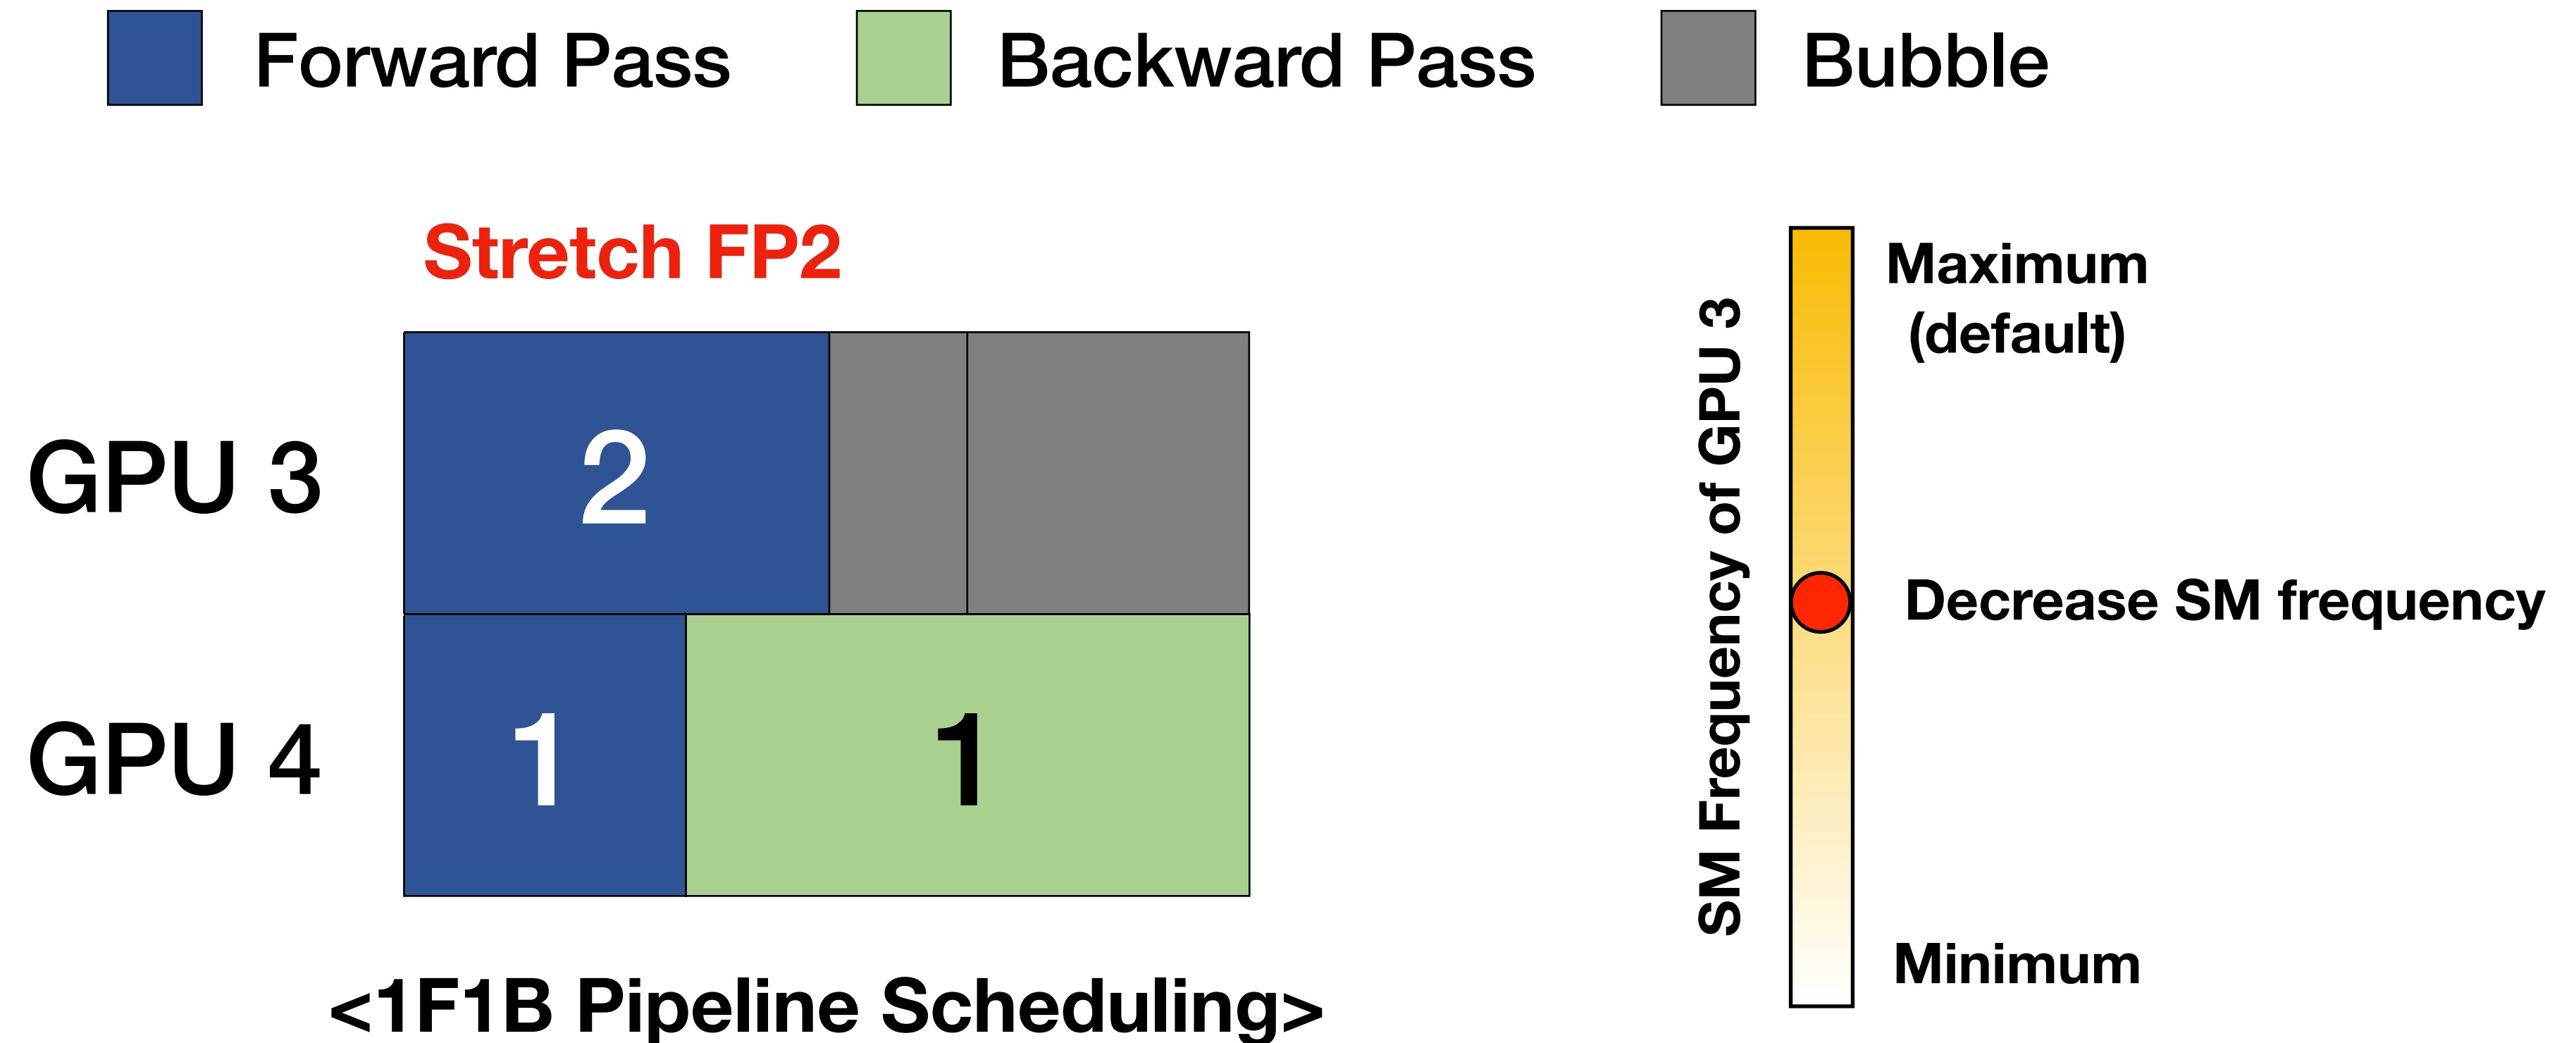


# Opportunity: Bubbles in Pipeline Parallelism

■ Forward Pass    ■ Backward Pass    ■ Bubble



# Opportunity: Bubbles in Pipeline Parallelism



Leverage pipeline bubbles to save energy without performance degradation

# Challenge 1: Not All Bubbles are Usable

■ Forward Pass    ■ Backward Pass    ■ Bubble

GPU 1	1	2	3	4									1	2	3	4	
GPU 2		1	2	3	4							1	2	3	4		
GPU 3			1	2	3	4				1	2	3	4				
GPU 4				1	2	3	4	1	2	3	4						

<GPipe Pipeline Scheduling>


# Challenge 1: Not All Bubbles are Usable

Forward Pass
  Backward Pass
  Bubble

GPU 1	1	2	3	4									1	2	3	4		
GPU 2		1	2	3	4								1	2	3	4		
GPU 3			1	2	3	4					1	2	3	4				
GPU 4				1	2	3	4	1	2	3	4							

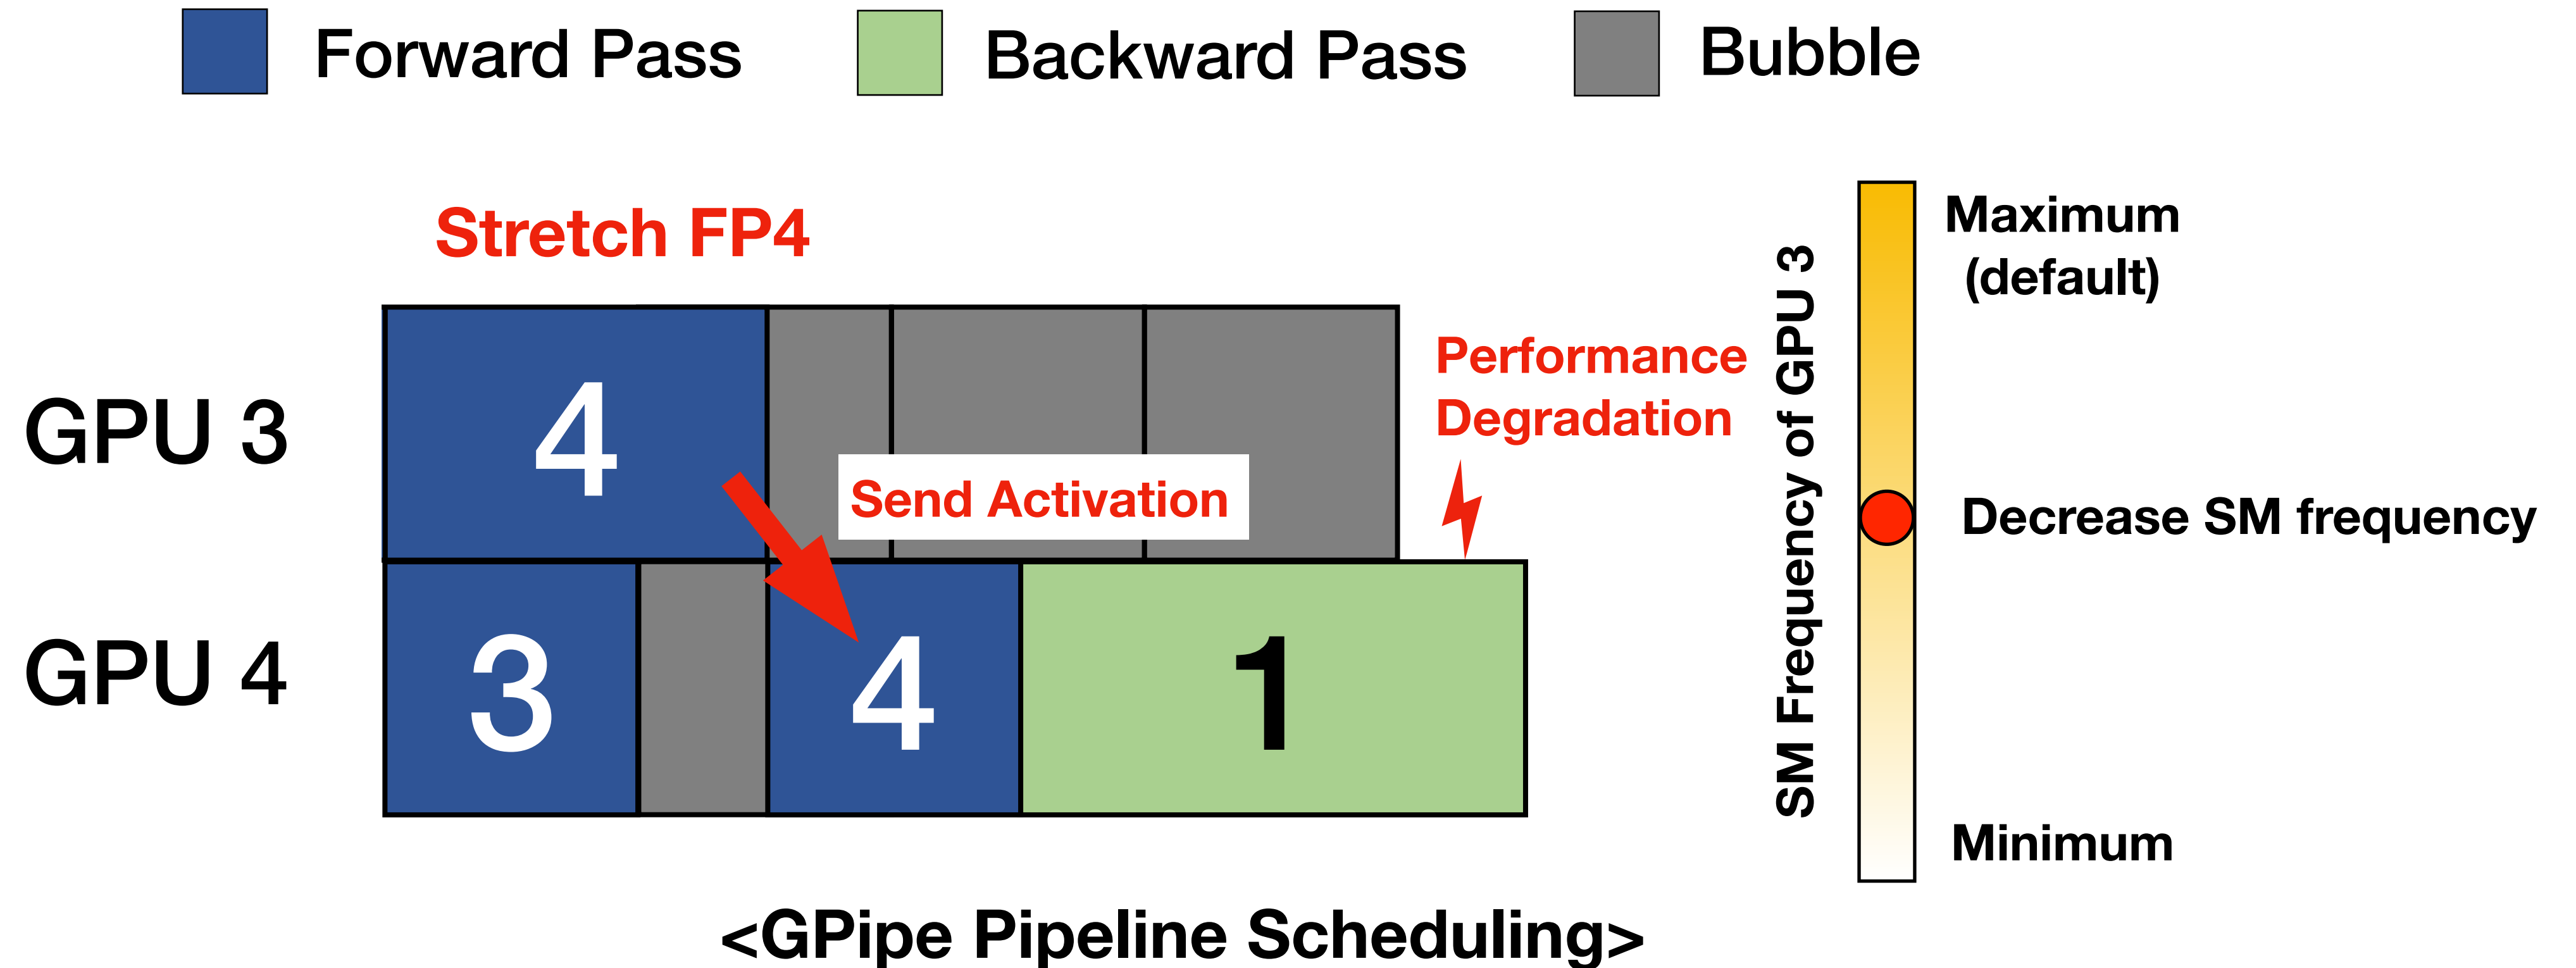
<GPipe Pipeline Scheduling>

# Challenge 1: Not All Bubbles are Usable



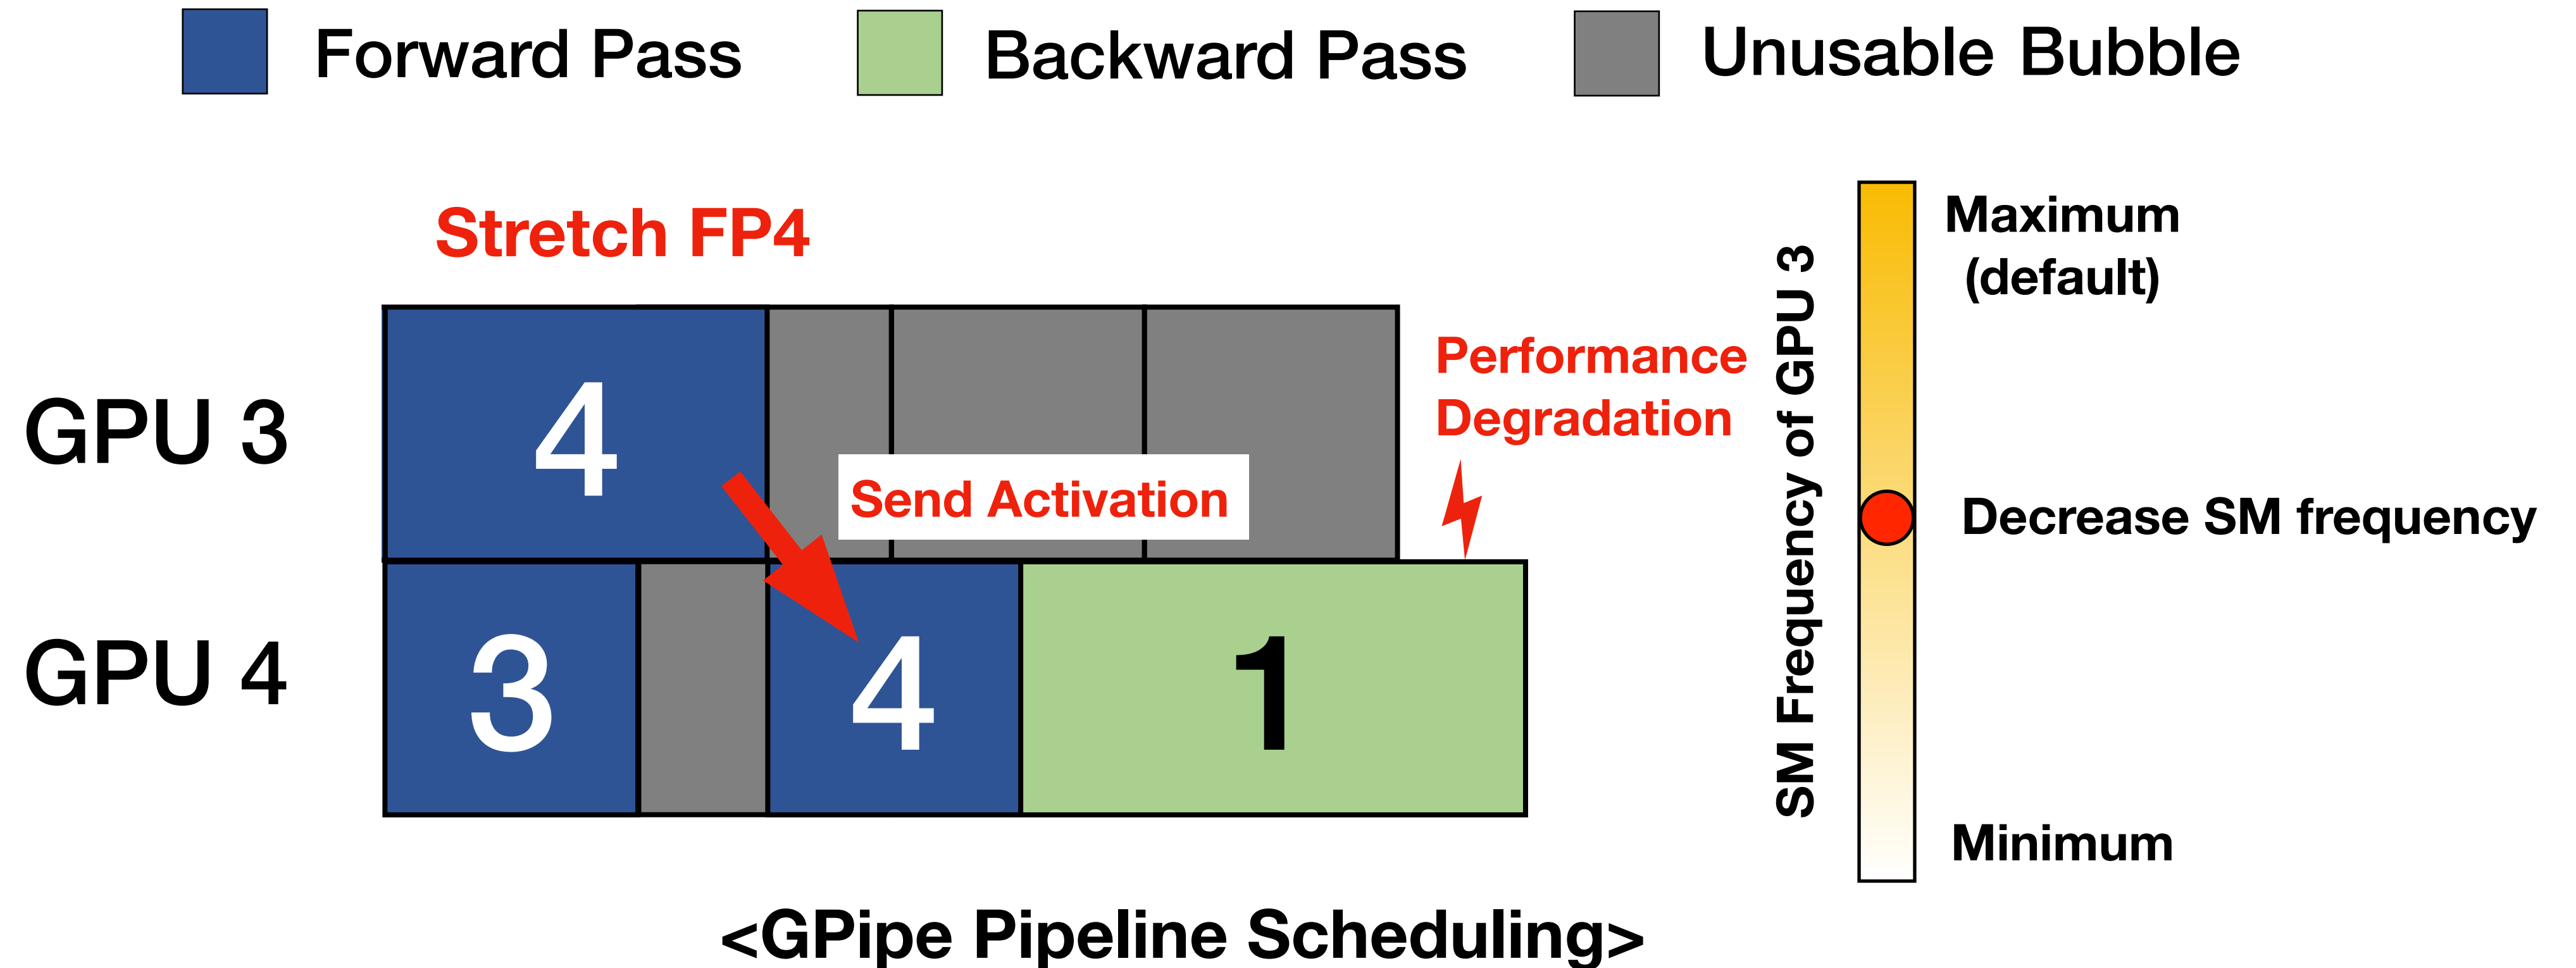
<GPipe Pipeline Scheduling>

# Challenge 1: Not All Bubbles are Usable







# Challenge 1: Not All Bubbles are Usable





**Need to identify usable and unusable bubbles**

# Challenge 2: Not Enough Usable Bubbles



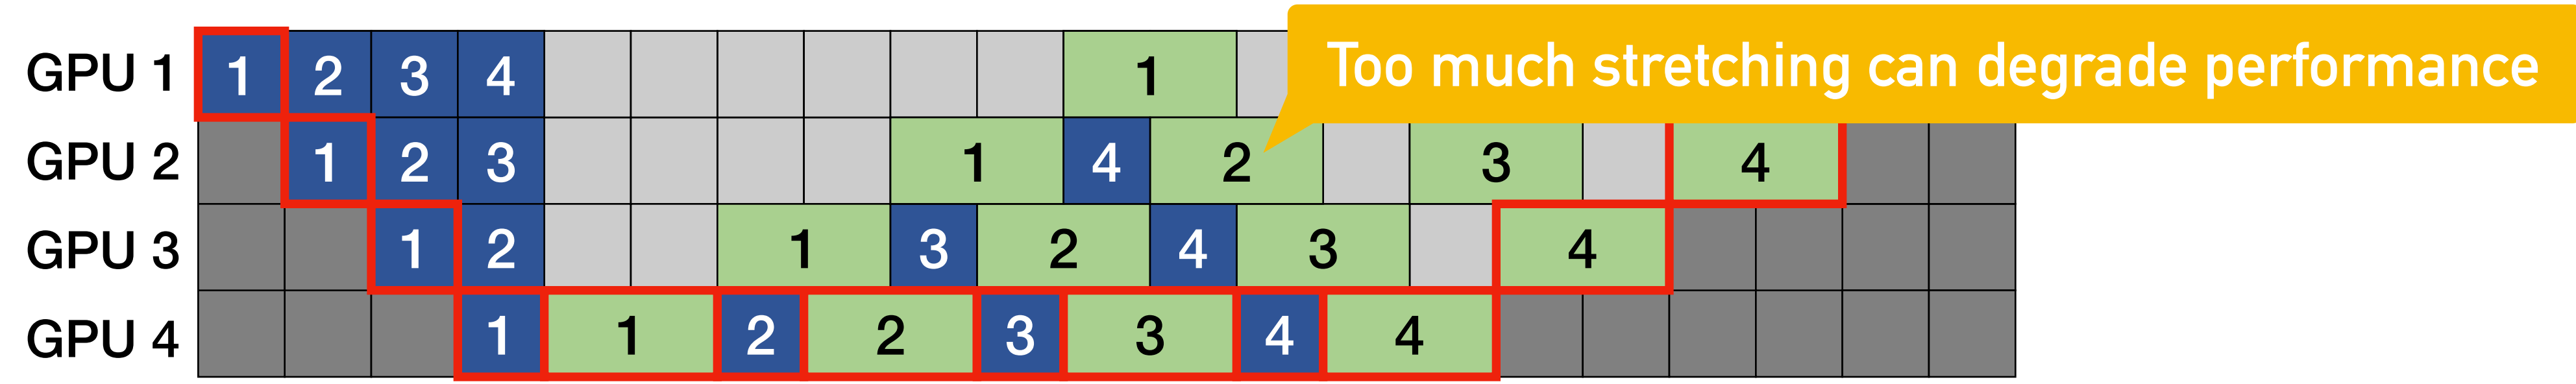

<1F1B Pipeline Scheduling>

# Challenge 2: Not Enough Usable Bubbles



<1F1B Pipeline Scheduling>

# Challenge 2: Not Enough Usable Bubbles



<1F1B Pipeline Scheduling>

**Leverage bubbles just the right amount to avoid performance degradation**

# Goals of EnvPipe

Envelope + Pipeline Parallelism

# Goals of EnvPipe

Envelope + Pipeline Parallelism

**Maximize Energy Saving**

Leverage as many pipeline bubbles as possible

# Goals of EnvPipe

Envelope + Pipeline Parallelism

**Maximize Energy Saving**

Leverage as many pipeline bubbles as possible

**No Accuracy Degradation**

Not modify user-provided hyperparameters



# Goals of EnvPipe

Envelope + Pipeline Parallelism

**Maximize Energy Saving**

Leverage as many pipeline bubbles as possible

**No Accuracy Degradation**

Not modify user-provided hyperparameters


**Minimize Perf. Degradation**

Keep performance degradation to near zero

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

# EnvPipe

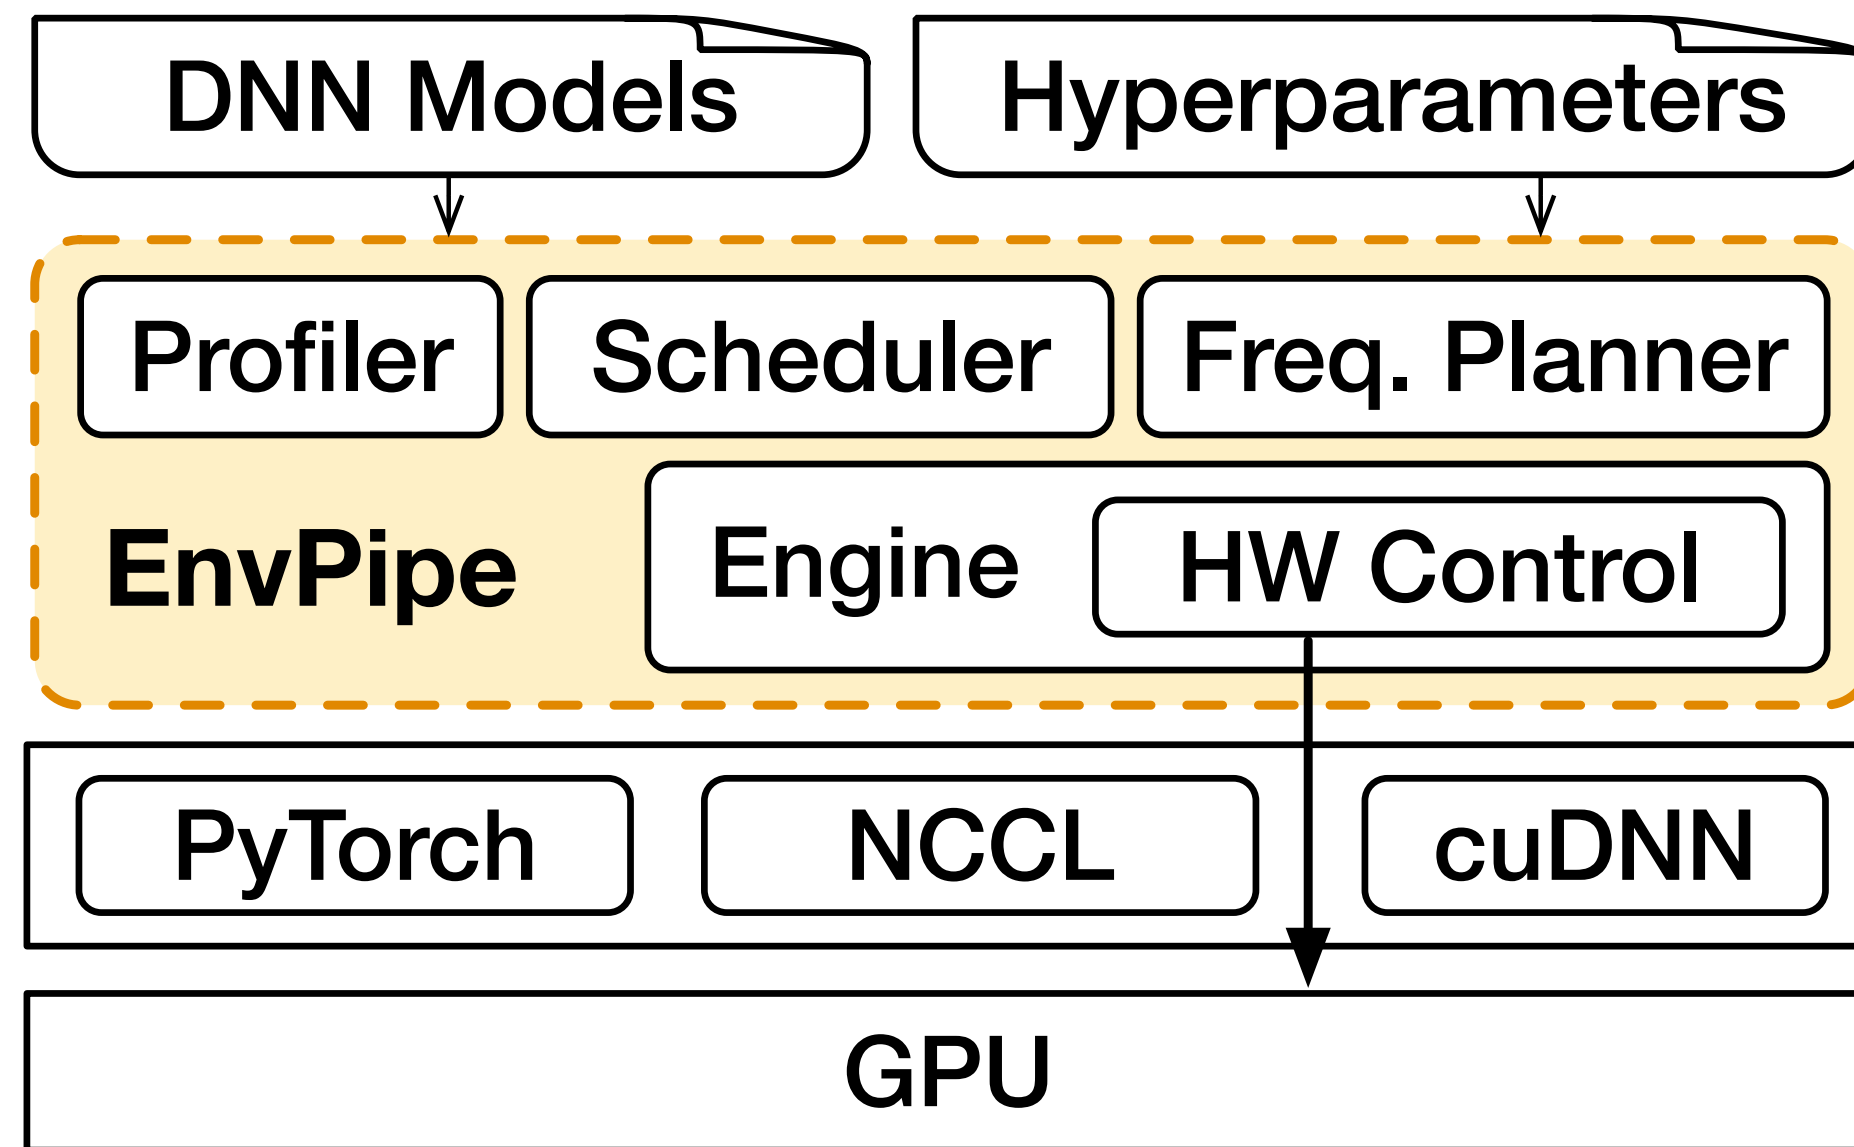
Performance-preserving DNN Training Framework for Saving Energy



Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

# EnvPipe

Performance-preserving DNN Training Framework for Saving Energy



1. Profiler

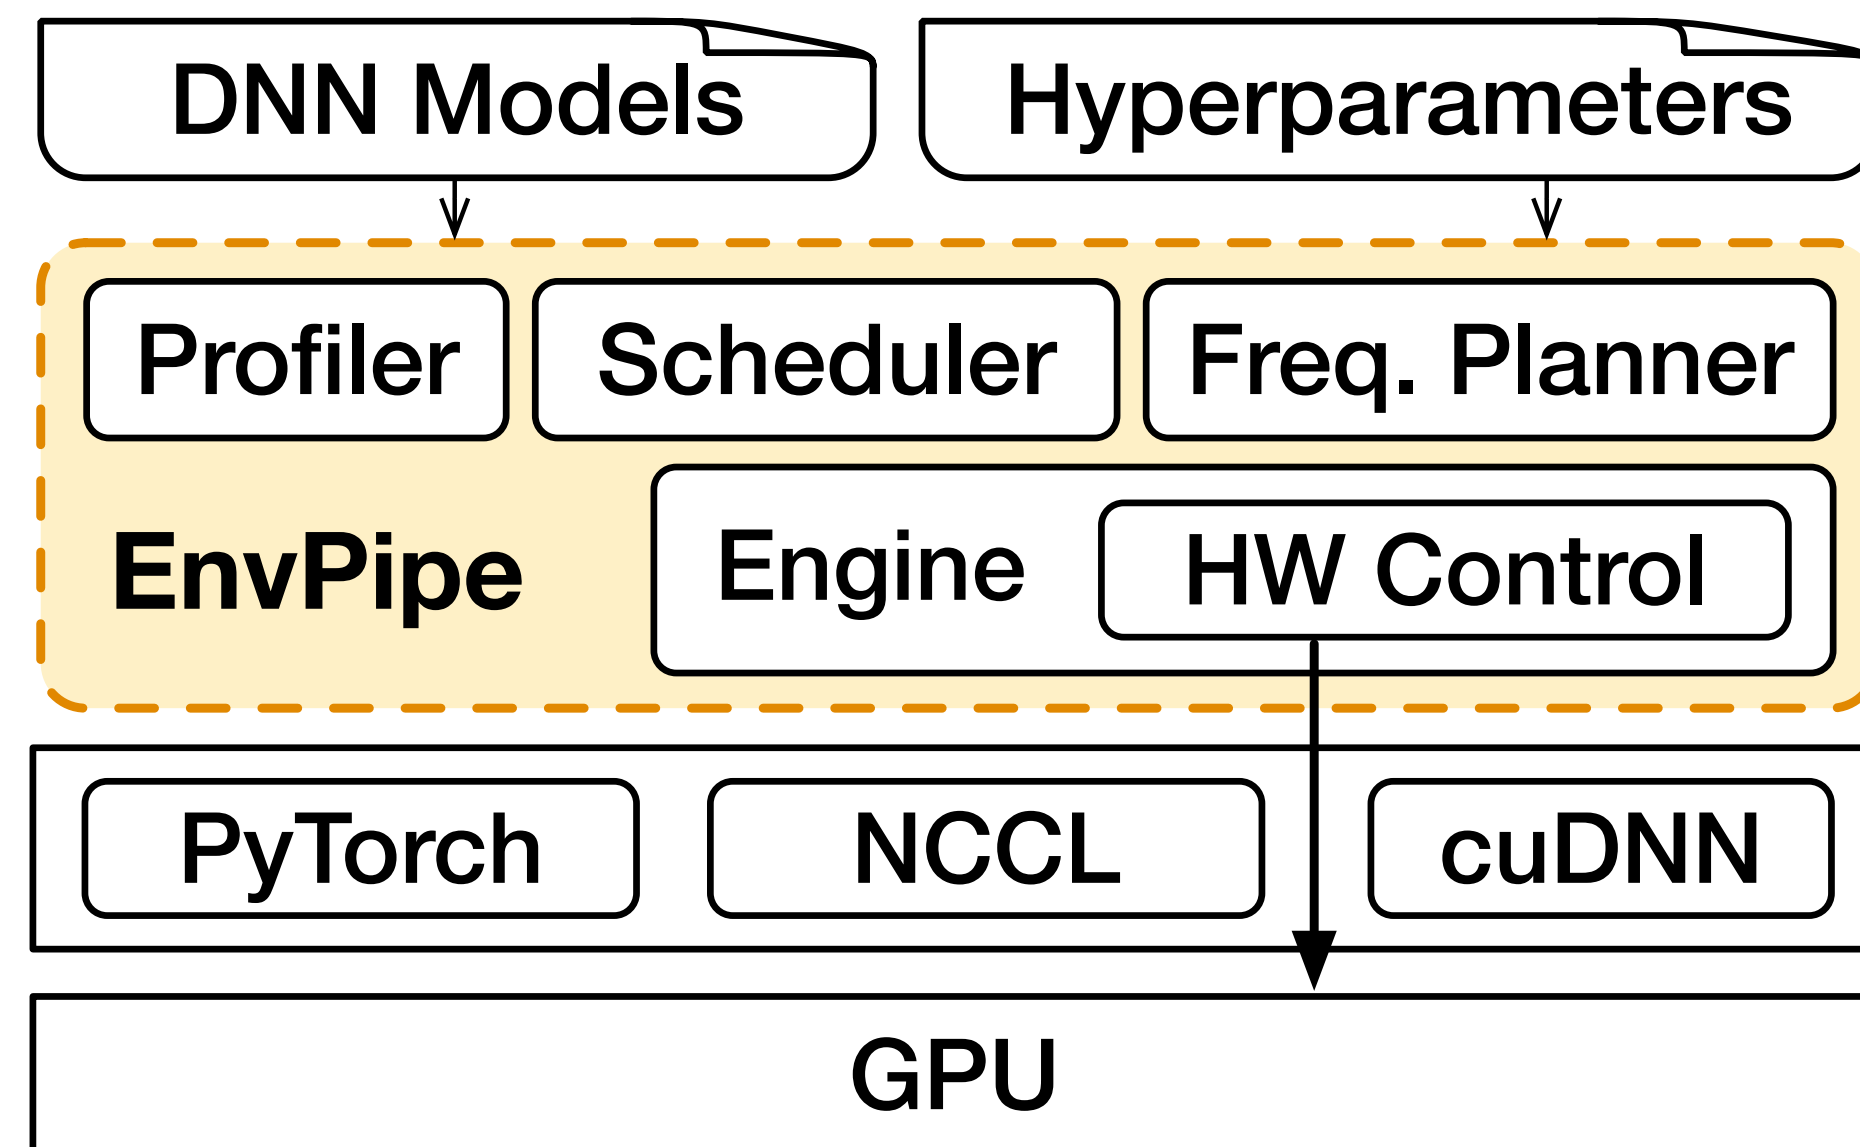
2. Scheduler

3. Freq. Planner

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

# EnvPipe

Performance-preserving DNN Training Framework for Saving Energy



## 1. Profiler

Profiles performance-energy trend for each stage

## 2. Scheduler

Schedules pipeline units to maximize energy saving

## 3. Freq. Planner

Reconfigures to minimize performance degradation


# 1. Profiler

2. Scheduler

3. Freq. Planner

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

- Energy consumption shows **valley trend** according to SM frequency
  - Depends on hardware, batch size, and partitioned model
- Online profiler profiles **each pipeline stage** to obtain the performance-energy trend




# 1. Profiler

2. Scheduler

3. Freq. Planner

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

- Energy consumption shows **valley trend** according to SM frequency
  - Depends on hardware, batch size, and partitioned model
- Online profiler profiles **each pipeline stage** to obtain the performance-energy trend



- How to schedule pipeline units determines the amount of usable bubbles

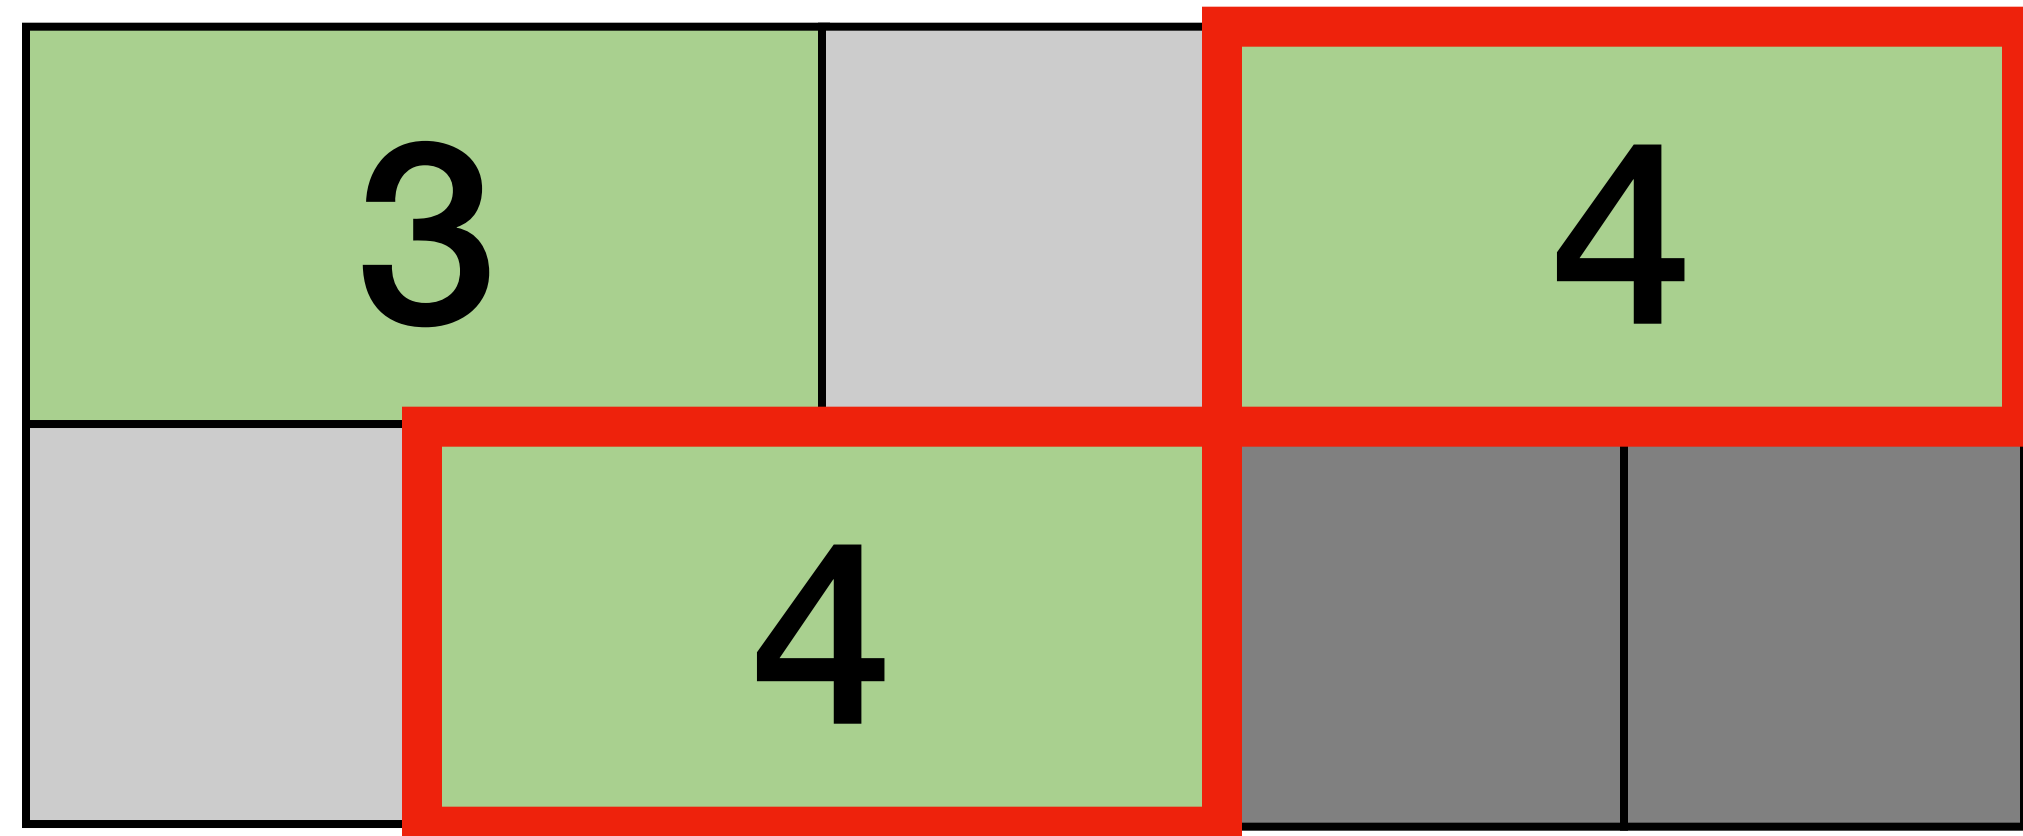

## Design questions

- 1) How to effectively identify usable bubbles?
- 2) How to maximize utilization of usable bubbles for energy saving?







1) How to effectively *identify* usable bubbles?



1) How to effectively *identify* usable bubbles?





1. Profiler

# 2. Scheduler



3. Freq. Planner

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation


1) How to effectively *identify* usable bubbles?




1) How to effectively *identify* usable bubbles?



1) How to effectively *identify* usable bubbles?



Bubbles after non-performance-critical path are usable



Bubbles after critical path are unusable











# 2. Scheduler

Maximize Energy Saving  
 No Accuracy Degradation  
 Minimize Perf. Degradation



2) How to *maximize utilization* of usable bubbles?



# 2. Scheduler

Maximize Energy Saving  
 No Accuracy Degradation  
 Minimize Perf. Degradation

2) How to *maximize utilization* of usable bubbles?


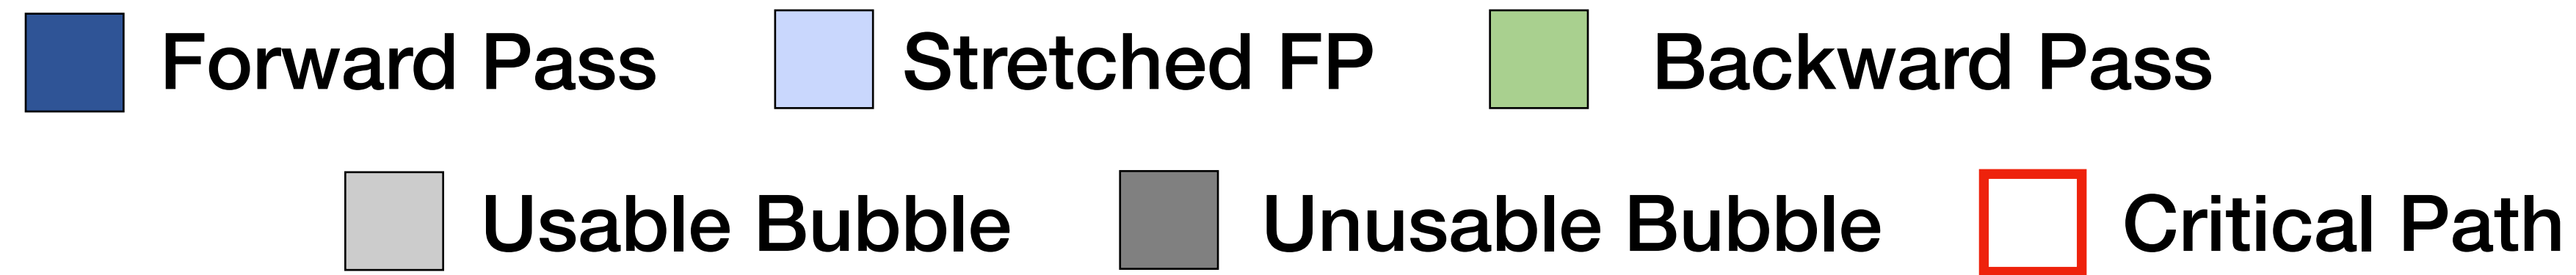


Far dependency

# 2. Scheduler

Maximize Energy Saving  
 No Accuracy Degradation  
 Minimize Perf. Degradation

2) How to *maximize utilization* of usable bubbles?





Far dependency

# 2. Scheduler

Maximize Energy Saving  
 No Accuracy Degradation  
 Minimize Perf. Degradation



2) How to *maximize utilization* of usable bubbles?



# 2. Scheduler

Maximize Energy Saving  
 No Accuracy Degradation  
 Minimize Perf. Degradation

2) How to *maximize utilization* of usable bubbles?









# 2. Scheduler

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

2) How to *maximize utilization* of usable bubbles?



2) How to *maximize utilization* of usable bubbles?



**Change immediate dependency to far dependency**






# 2. Scheduler


Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

2) How to *maximize utilization* of usable bubbles?



Too much wasted bubbles at the front

Lots of backward computation but not enough usable bubbles to utilize

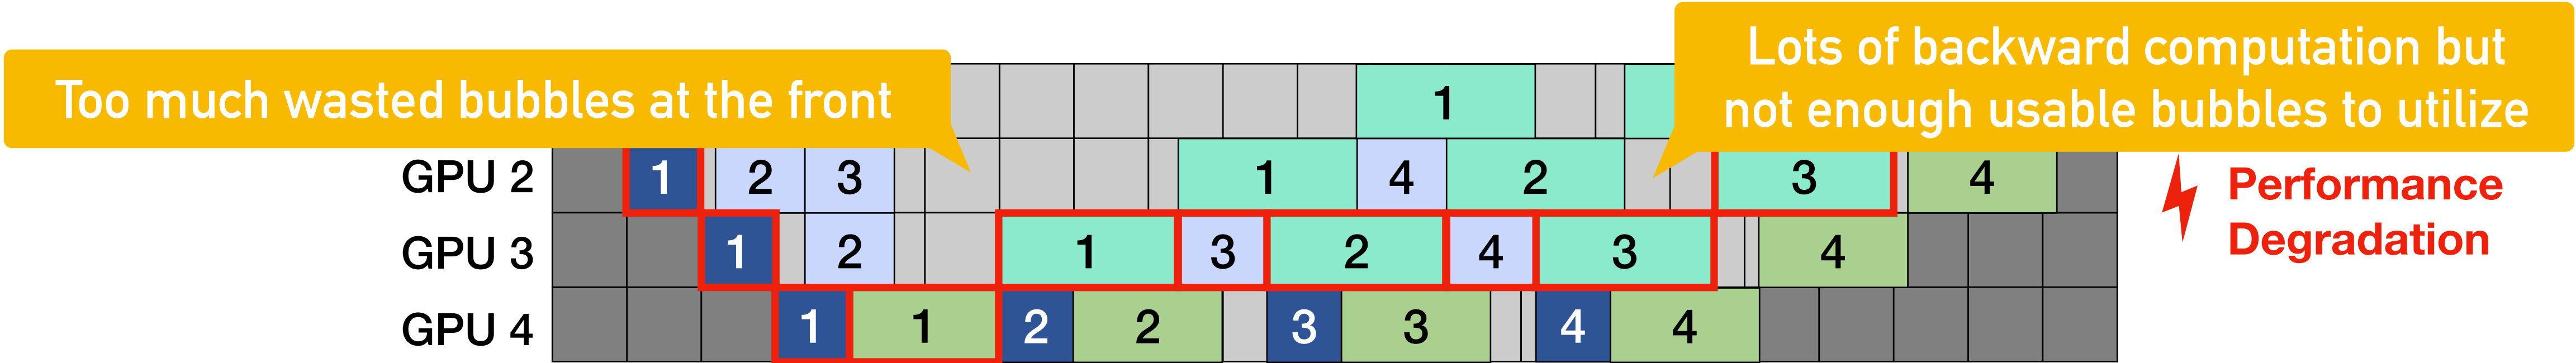



⚡ Performance Degradation

# 2. Scheduler

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

2) How to *maximize utilization* of usable bubbles?



**Distribute usable bubbles to maximize utilization**


# 2. Scheduler

Maximize Energy Saving  
 No Accuracy Degradation  
 Minimize Perf. Degradation

## Scheduler to maximize utilization of usable bubbles

Step 1) Initialize with 1F1B scheduling

► Among existing pipeline scheduling methods, 1F1B has the fewest units on the performance-critical path




### Scheduler to maximize utilization of usable bubbles

Step 1) Initialize with 1F1B scheduling

- Among existing pipeline scheduling methods, 1F1B has the fewest units on the performance-critical path

### Step 2) Reschedule FP units

- Change immediate dependency and distribute usable bubbles to maximize bubble utilization




### Scheduler to maximize utilization of usable bubbles

Step 1) Initialize with 1F1B scheduling

- Among existing pipeline scheduling methods, 1F1B has the fewest units on the performance-critical path

### Step 2) Reschedule FP units

- Change immediate dependency and distribute usable bubbles to maximize bubble utilization
- Keep data dependency (send/recv activation/gradient)



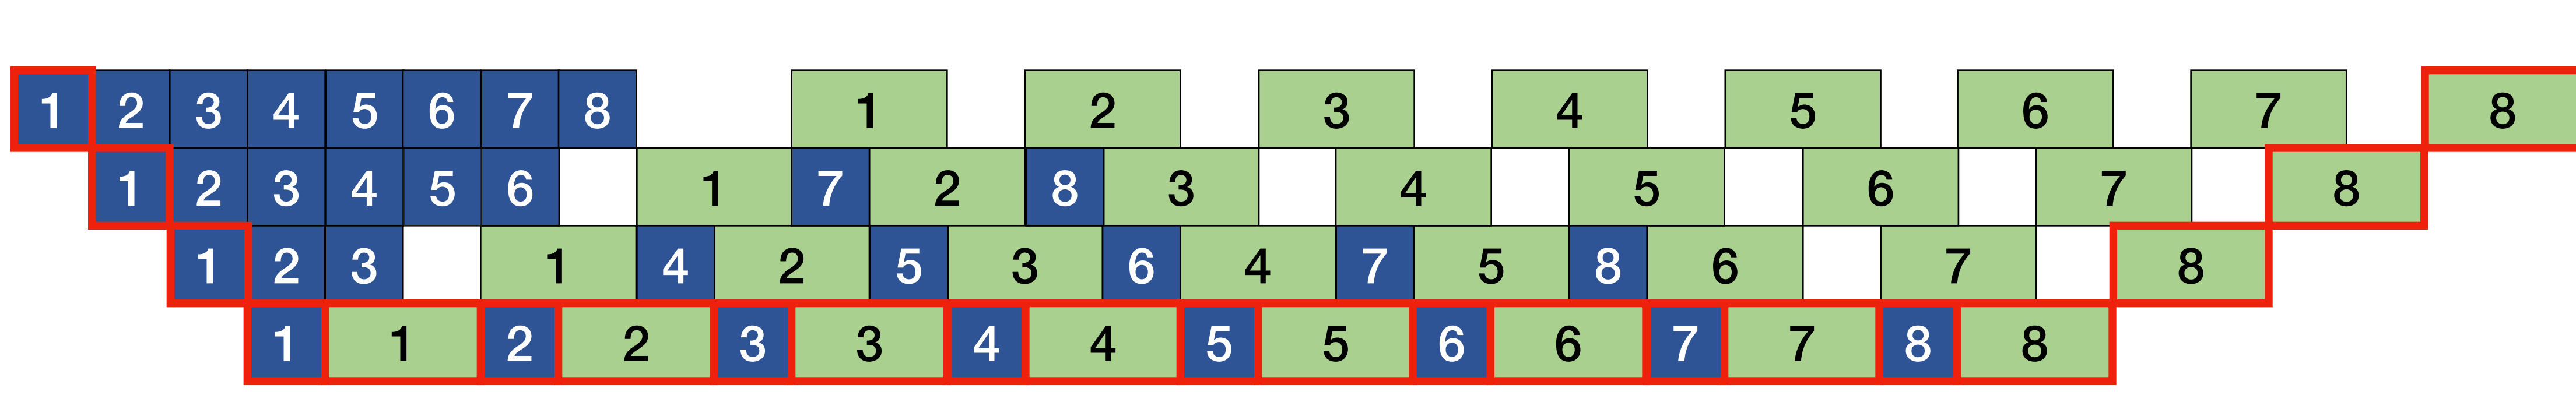

1. Profiler

# 2. Scheduler

3. Freq. Planner

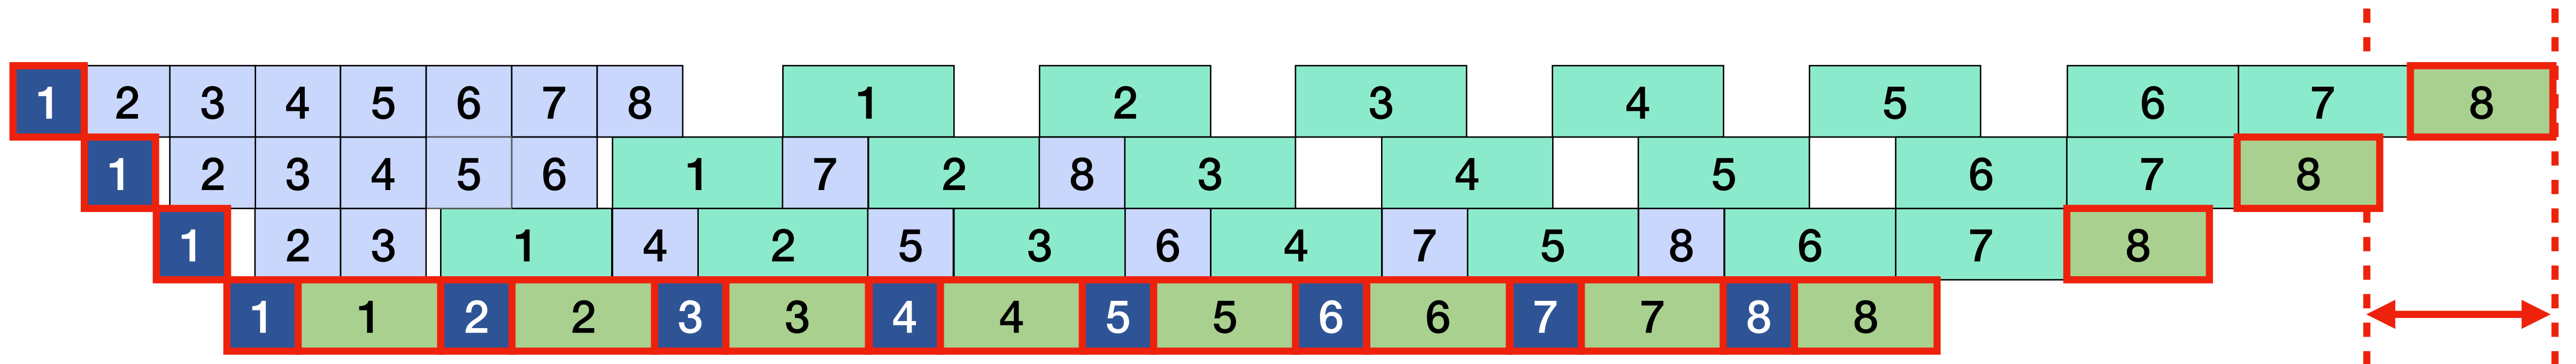

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

## Stretch all pipeline units inside the envelope



Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

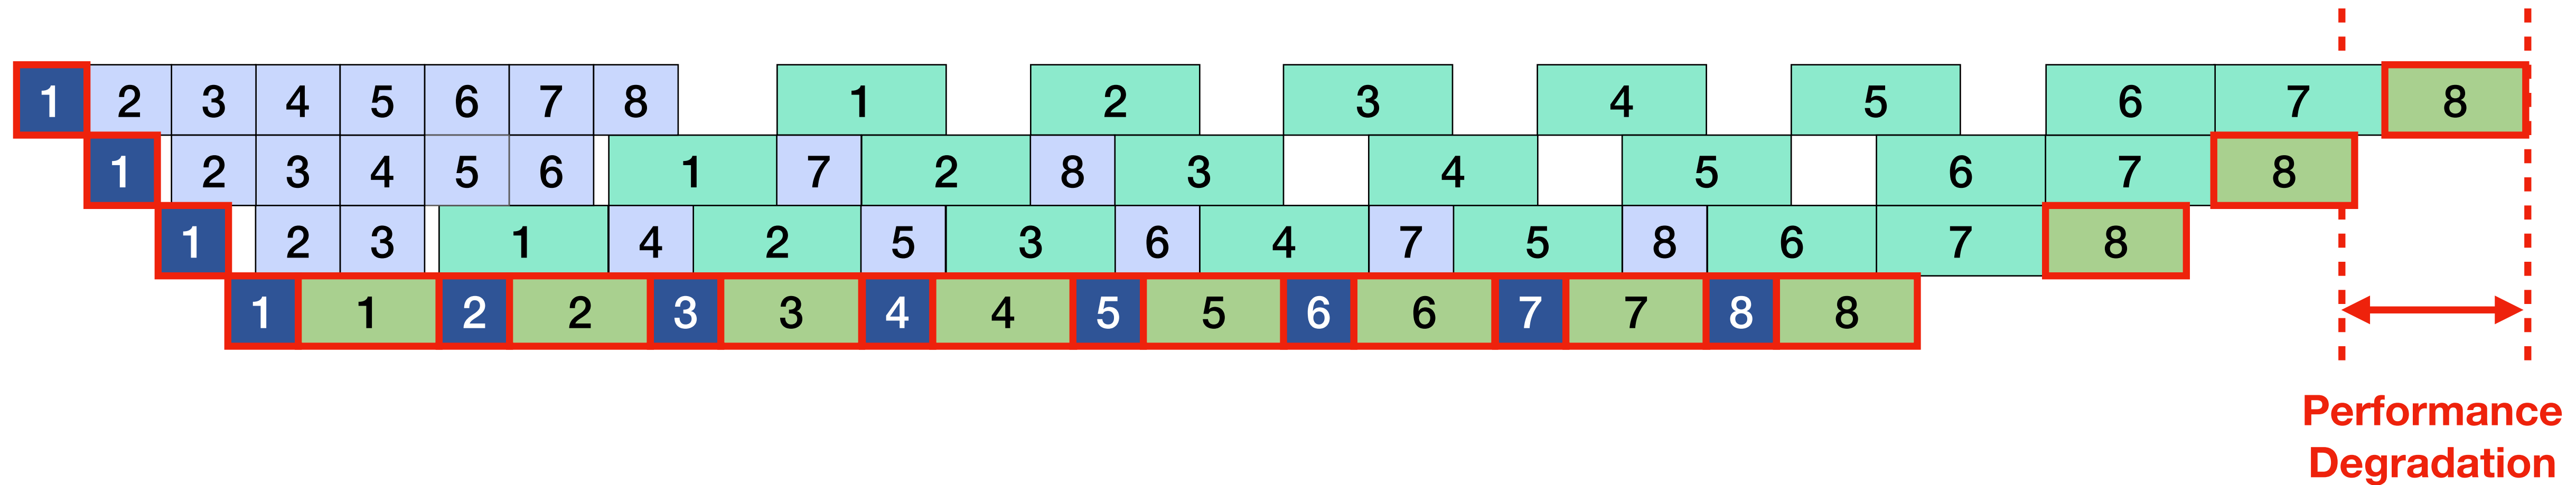

## Stretch all pipeline units inside the envelope





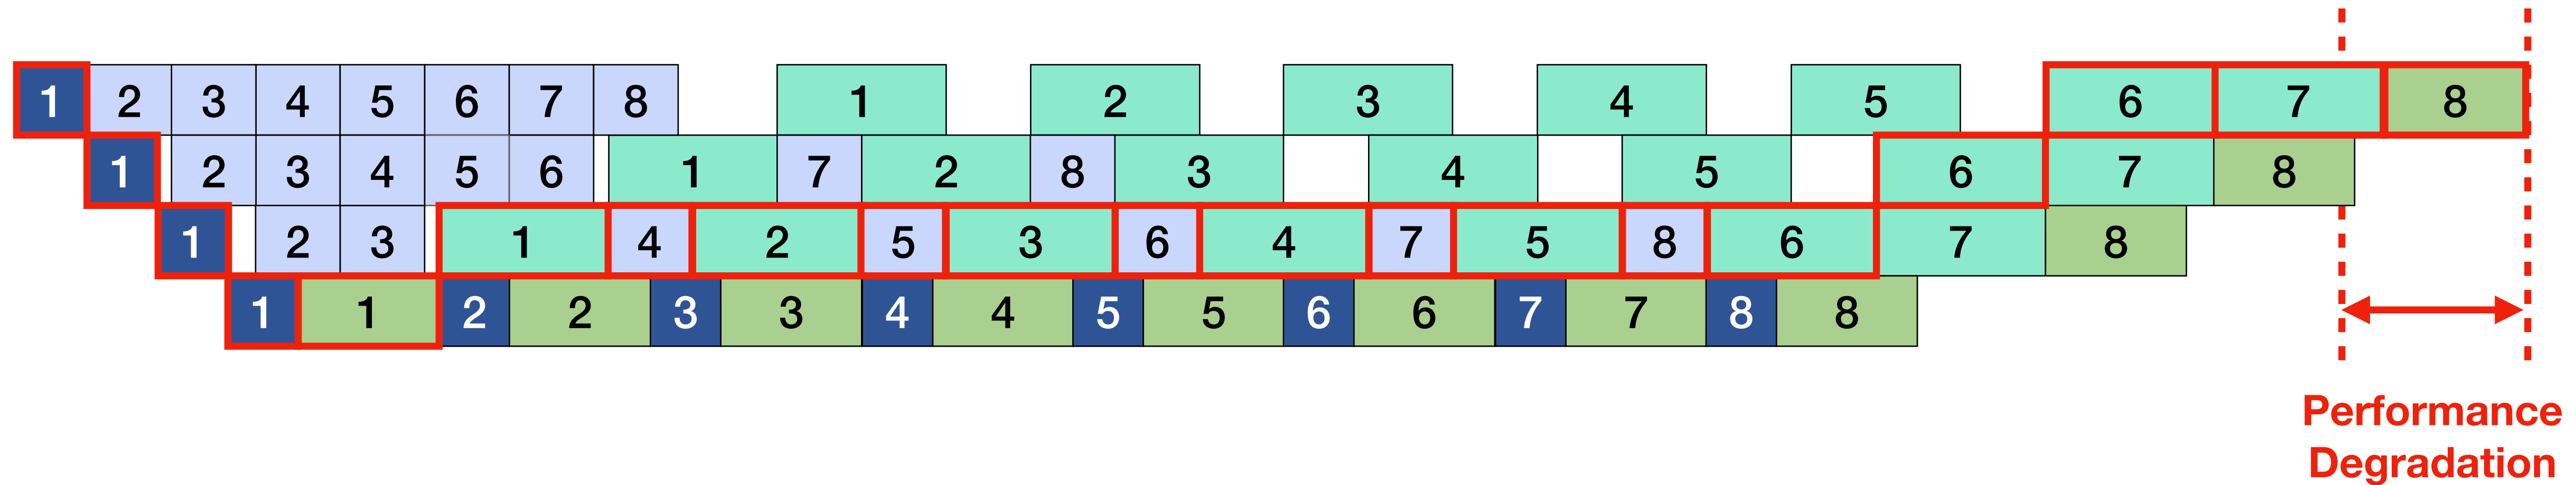

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

## Stretch all pipeline units inside the envelope

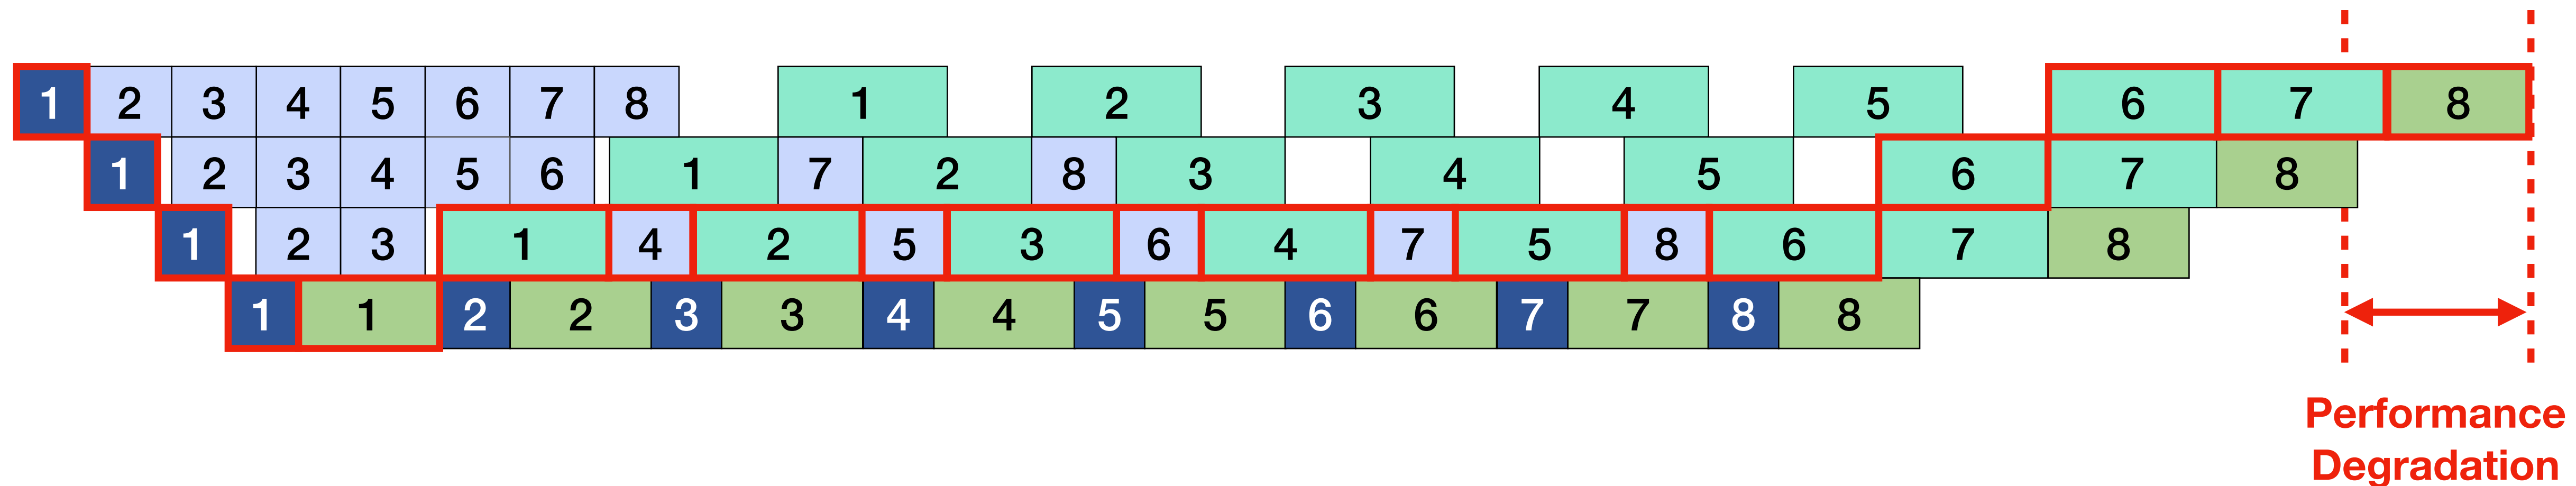



Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

## Stretch all pipeline units inside the envelope



## Stretch all pipeline units inside the envelope



**Naively stretching all pipeline units inside the envelope degrades performance**

1. Profiler

2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

- Stretching all pipeline units inside the envelope leads to performance degradation
- To avoid performance degradation, SM frequency should be **reconfigured**


1. Profiler

2. Scheduler

# 3. Freq. Planner


Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

- Stretching all pipeline units inside the envelope leads to performance degradation
- To avoid performance degradation, SM frequency should be **reconfigured**



Max Freq. Forward

Max Freq. Backward



1. Profiler

2. Scheduler

# 3. Freq. Planner


Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

- Stretching all pipeline units inside the envelope leads to performance degradation
- To avoid performance degradation, SM frequency should be **reconfigured**



Optimal Freq. Forward

Optimal Freq. Backward




1. Profiler

2. Scheduler

# 3. Freq. Planner


Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

- Stretching all pipeline units inside the envelope leads to performance degradation
- To avoid performance degradation, SM frequency should be **reconfigured**



## Reconfigured Forward

## Reconfigured Backward



1. Profiler

2. Scheduler

## 3. Freq. Planner

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

- Stretching all pipeline units inside the envelope leads to performance degradation
- To avoid performance degradation, SM frequency should be **reconfigured**

### How much reconfiguration?

- Critical path should be outer envelope to minimize performance degradation
- Increasing SM frequency too much will reduce energy saving



1. Profiler

2. Scheduler

**3. Freq. Planner**

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

- Stretching all pipeline units inside the envelope leads to performance degradation
- To avoid performance degradation, SM frequency should be **reconfigured**

### How much reconfiguration?

- Critical path should be outer envelope to minimize performance degradation
- Increasing SM frequency too much will reduce energy saving

**Reconfigure just the right amount to maximize energy saving**

1. Profiler

2. Scheduler

**3. Freq. Planner**

**Maximize Energy Saving**  
**No Accuracy Degradation**  
**Minimize Perf. Degradation**

**Initial**



**Iteration**



**Termination**

1. Profiler

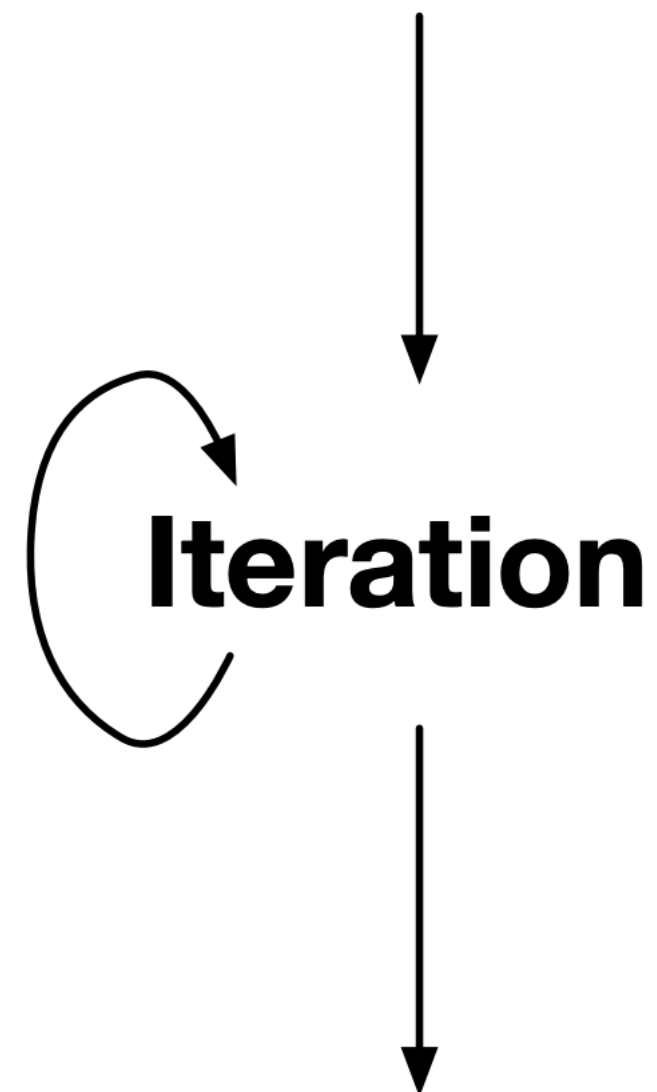
2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

**Initial**

Based on the scheduling decision, stretches pipeline units inside the envelope



**Termination**

1. Profiler

2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

**Initial**

Based on the scheduling decision, stretches pipeline units inside the envelope



**Iteration**

1) **Executes** single pipeline step

2) **Finds** changed critical path

3) **Reconfigures** SM frequency on the critical path in small steps



**Termination**

1. Profiler

2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

**Initial**

Based on the scheduling decision, stretches pipeline units inside the envelope



**Iteration**

1) **Executes** single pipeline step

2) **Finds** changed critical path

3) **Reconfigures** SM frequency on the critical path in small steps



**Termination**

Terminates when critical path is outer envelope

1. Profiler


2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

**Initial**

Based on the scheduling decision, stretches pipeline units inside the envelope



**Iteration**

1) **Executes** single pipeline step

2) **Finds** changed critical path

3) **Reconfigures** SM frequency on the critical path in small steps



**Termination**

Terminates when critical path is outer envelope

**How to find critical path?**


1. Profiler

2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

How to *find* critical path?



Which pipeline unit is delaying **8** ?

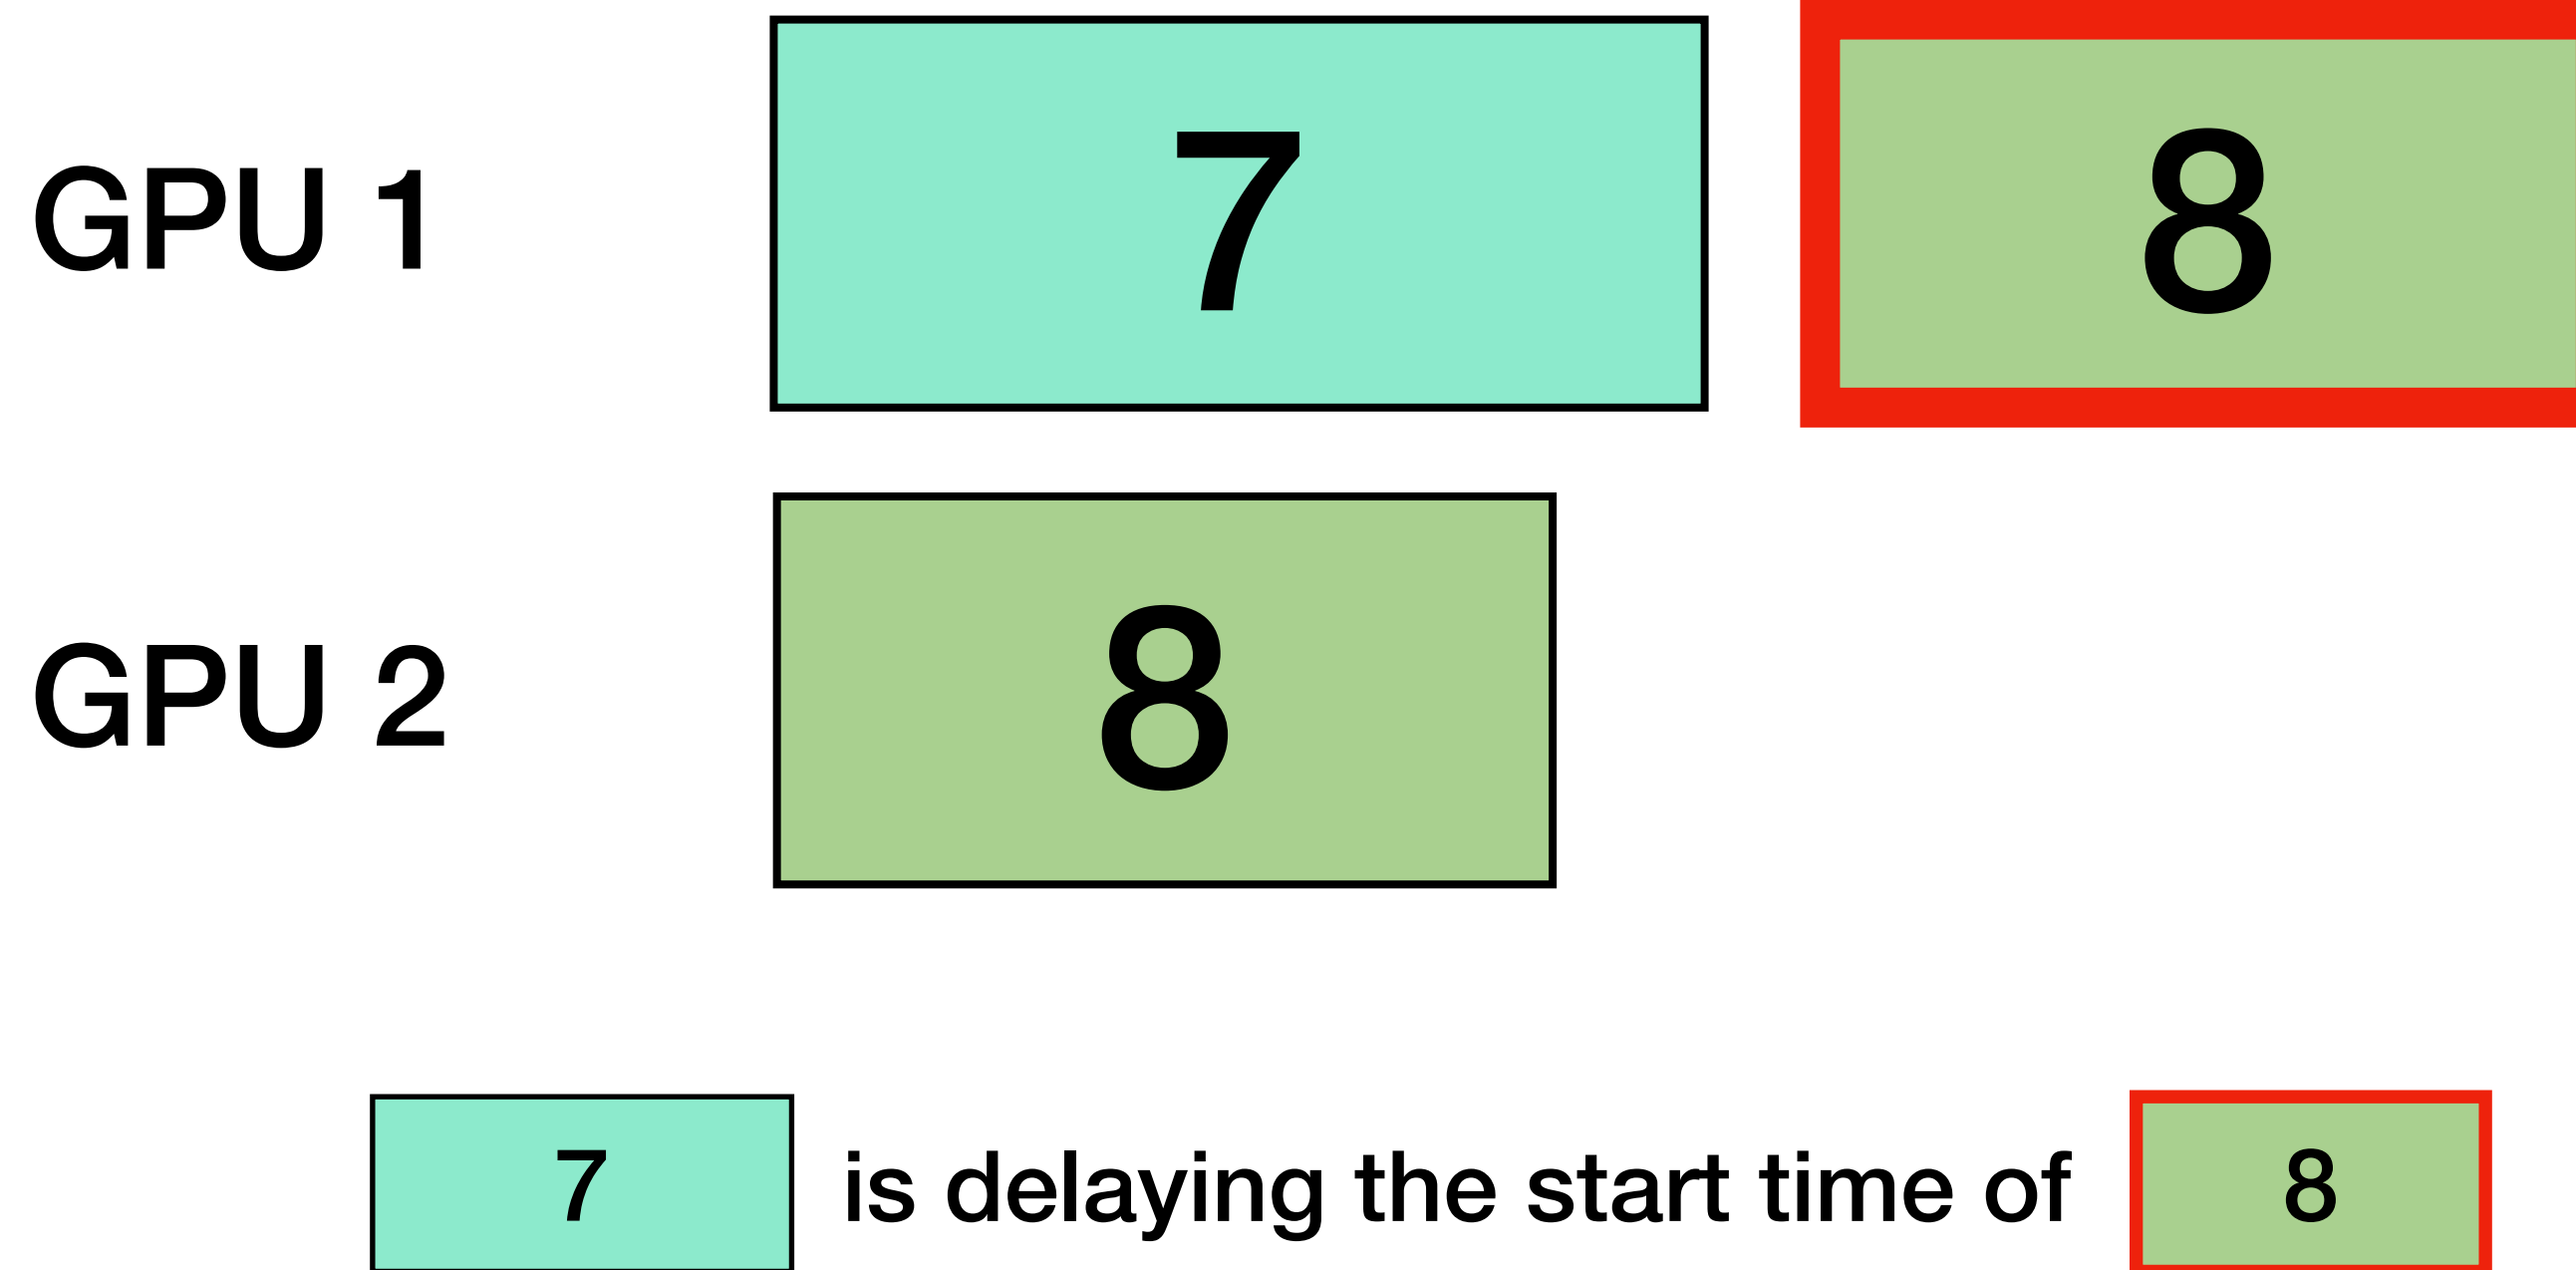
1. Profiler

2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

How to *find* critical path?





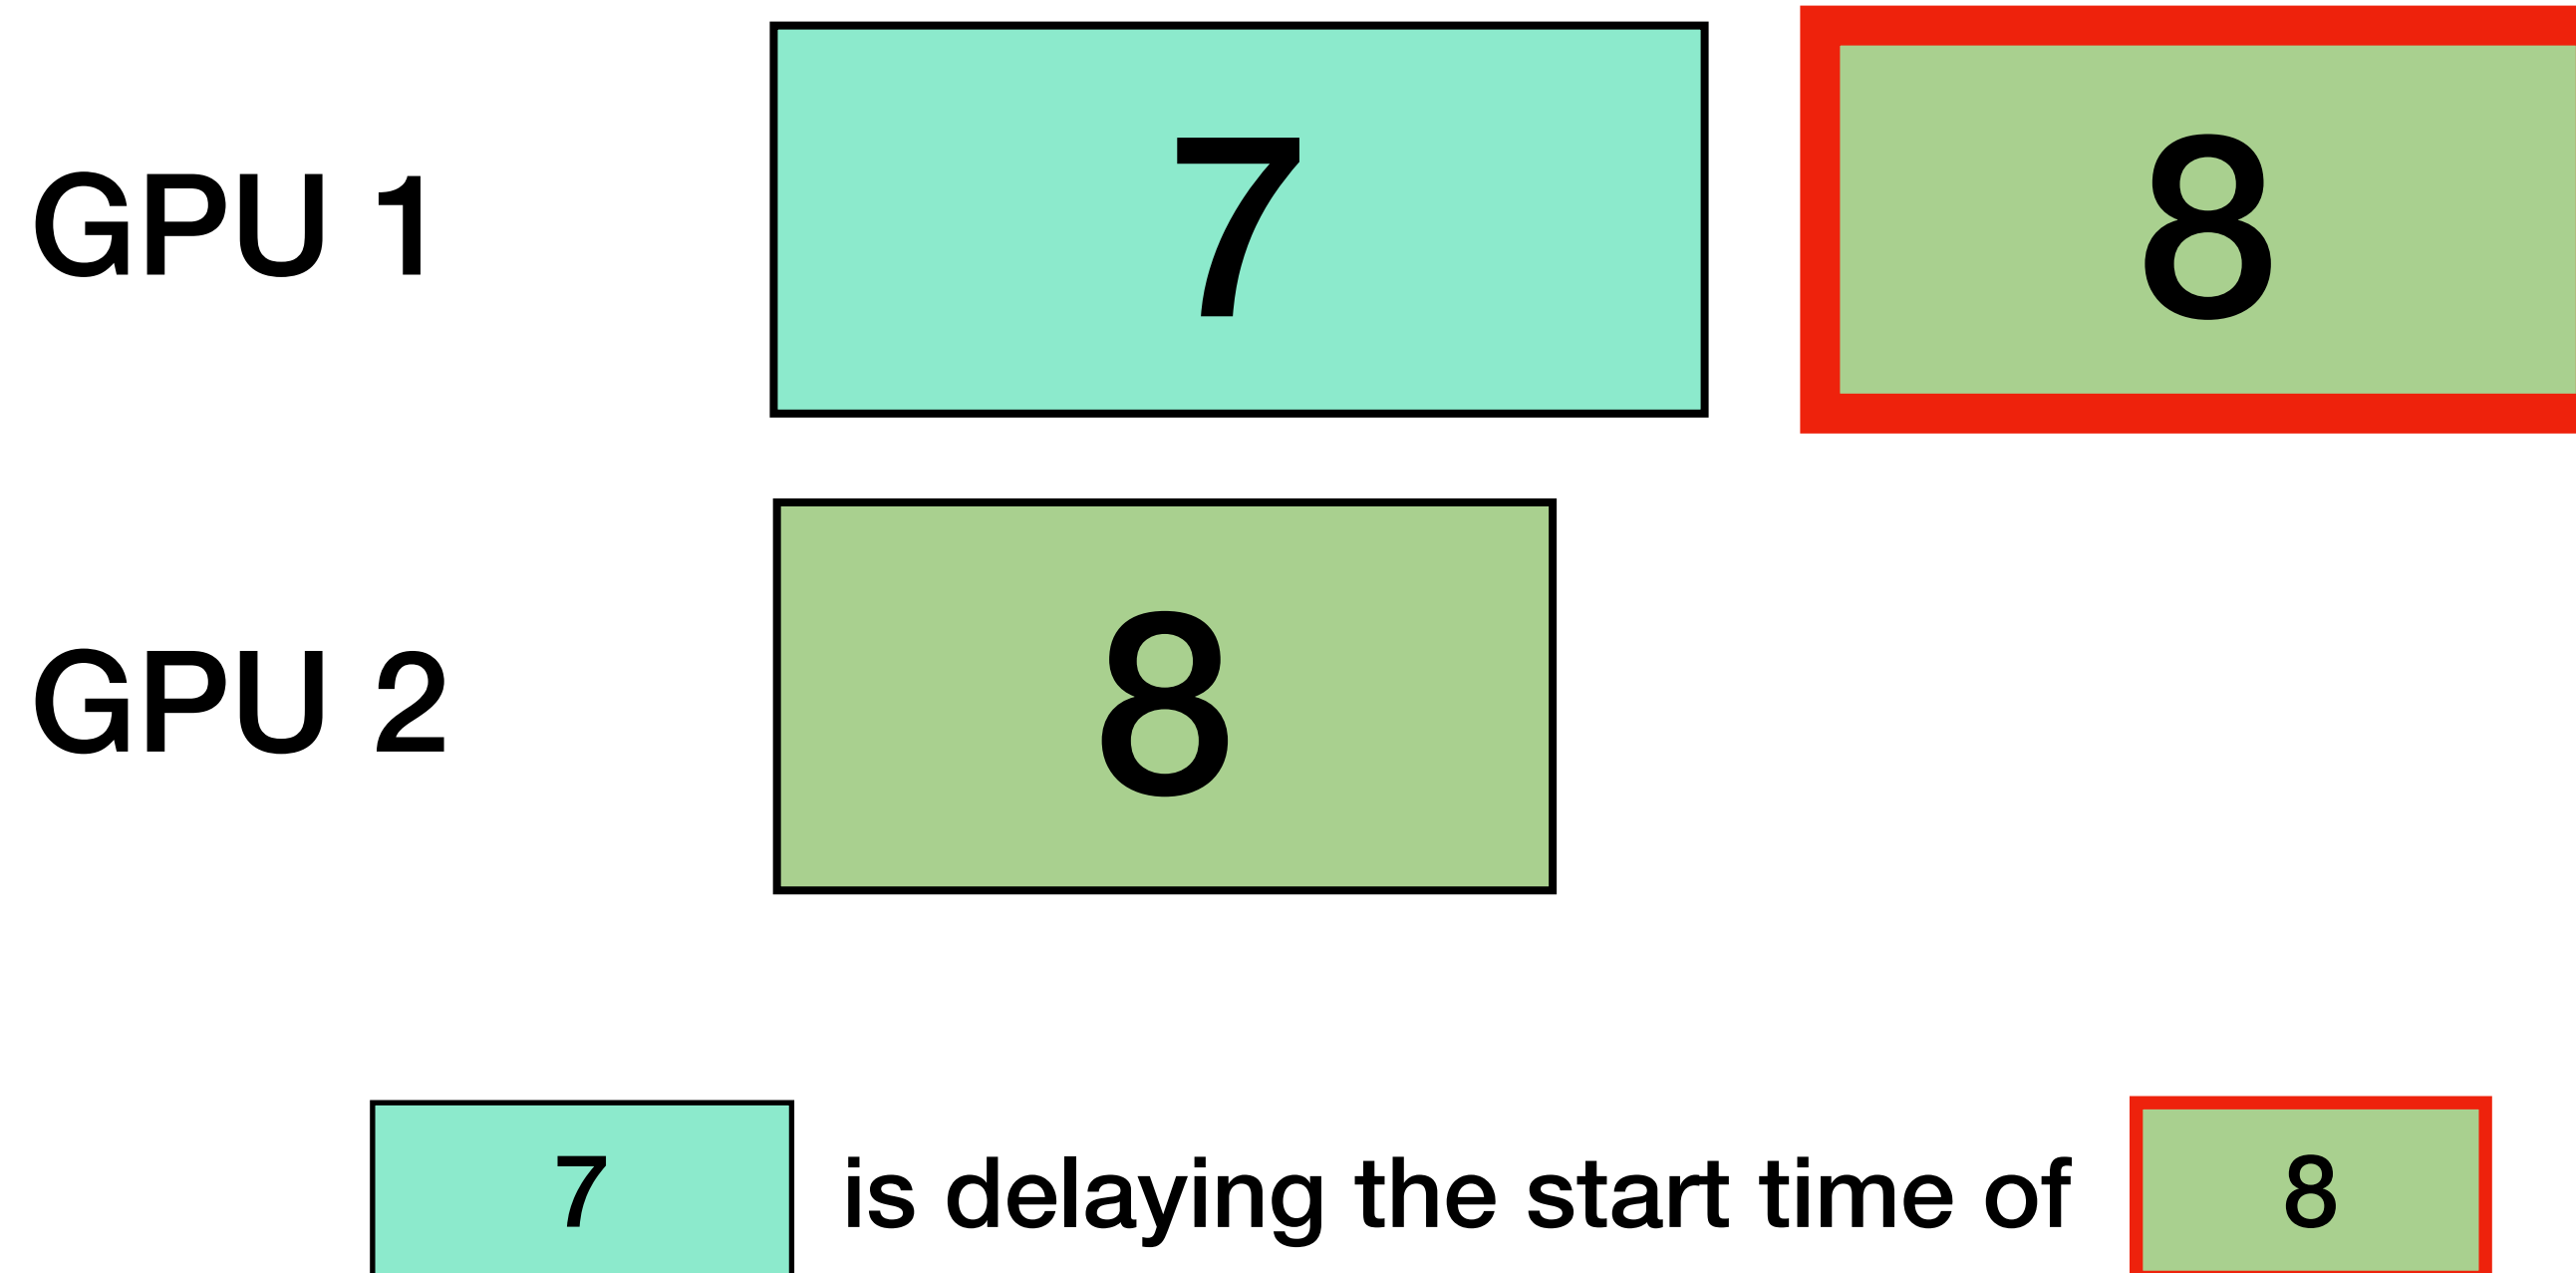
1. Profiler

2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

How to *find* critical path?



1. Profiler

2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

How to *find* critical path?

GPU 1

7

8

GPU 2

8

7

is delaying the start time of

8

, thus critical path extends to

7

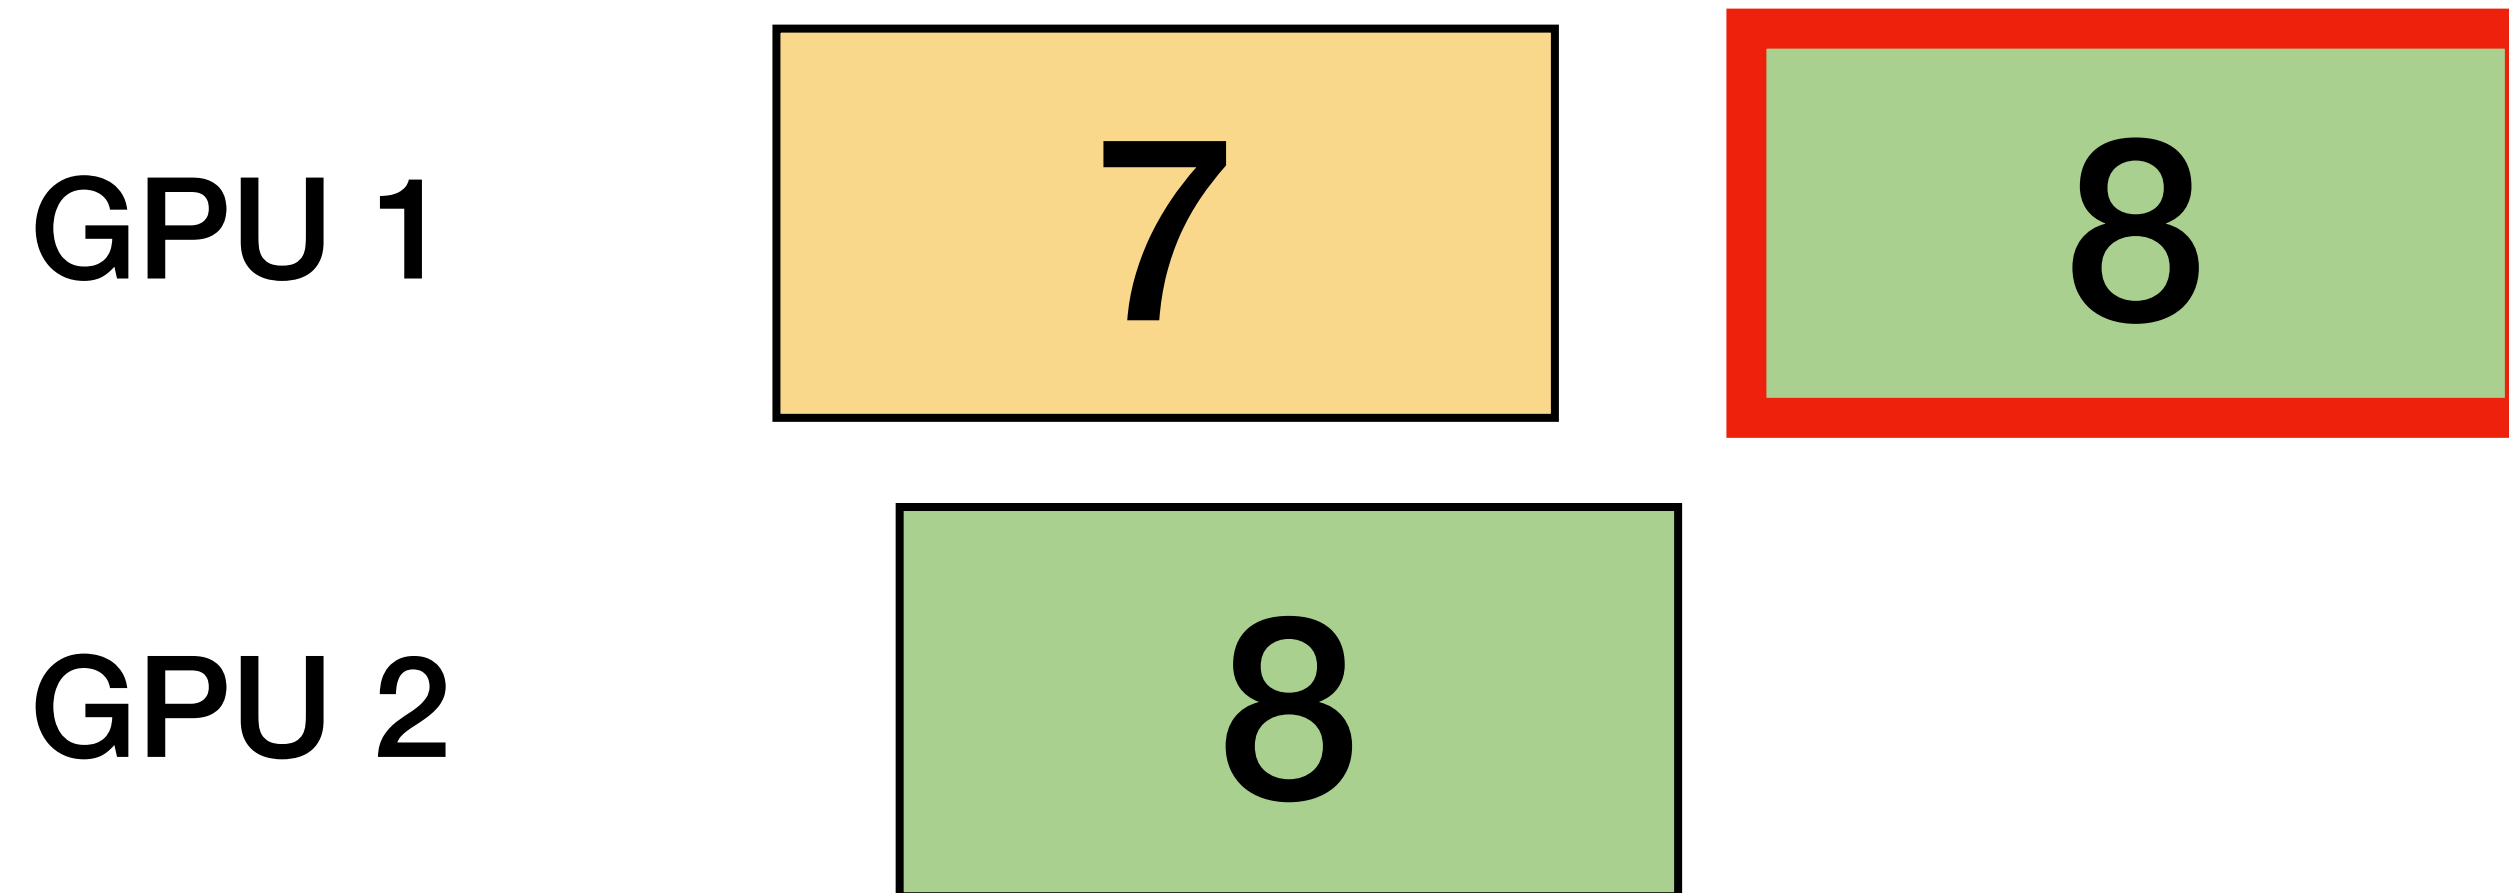
1. Profiler

2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

How to *find* critical path?



Which pipeline unit is delaying **8** ?

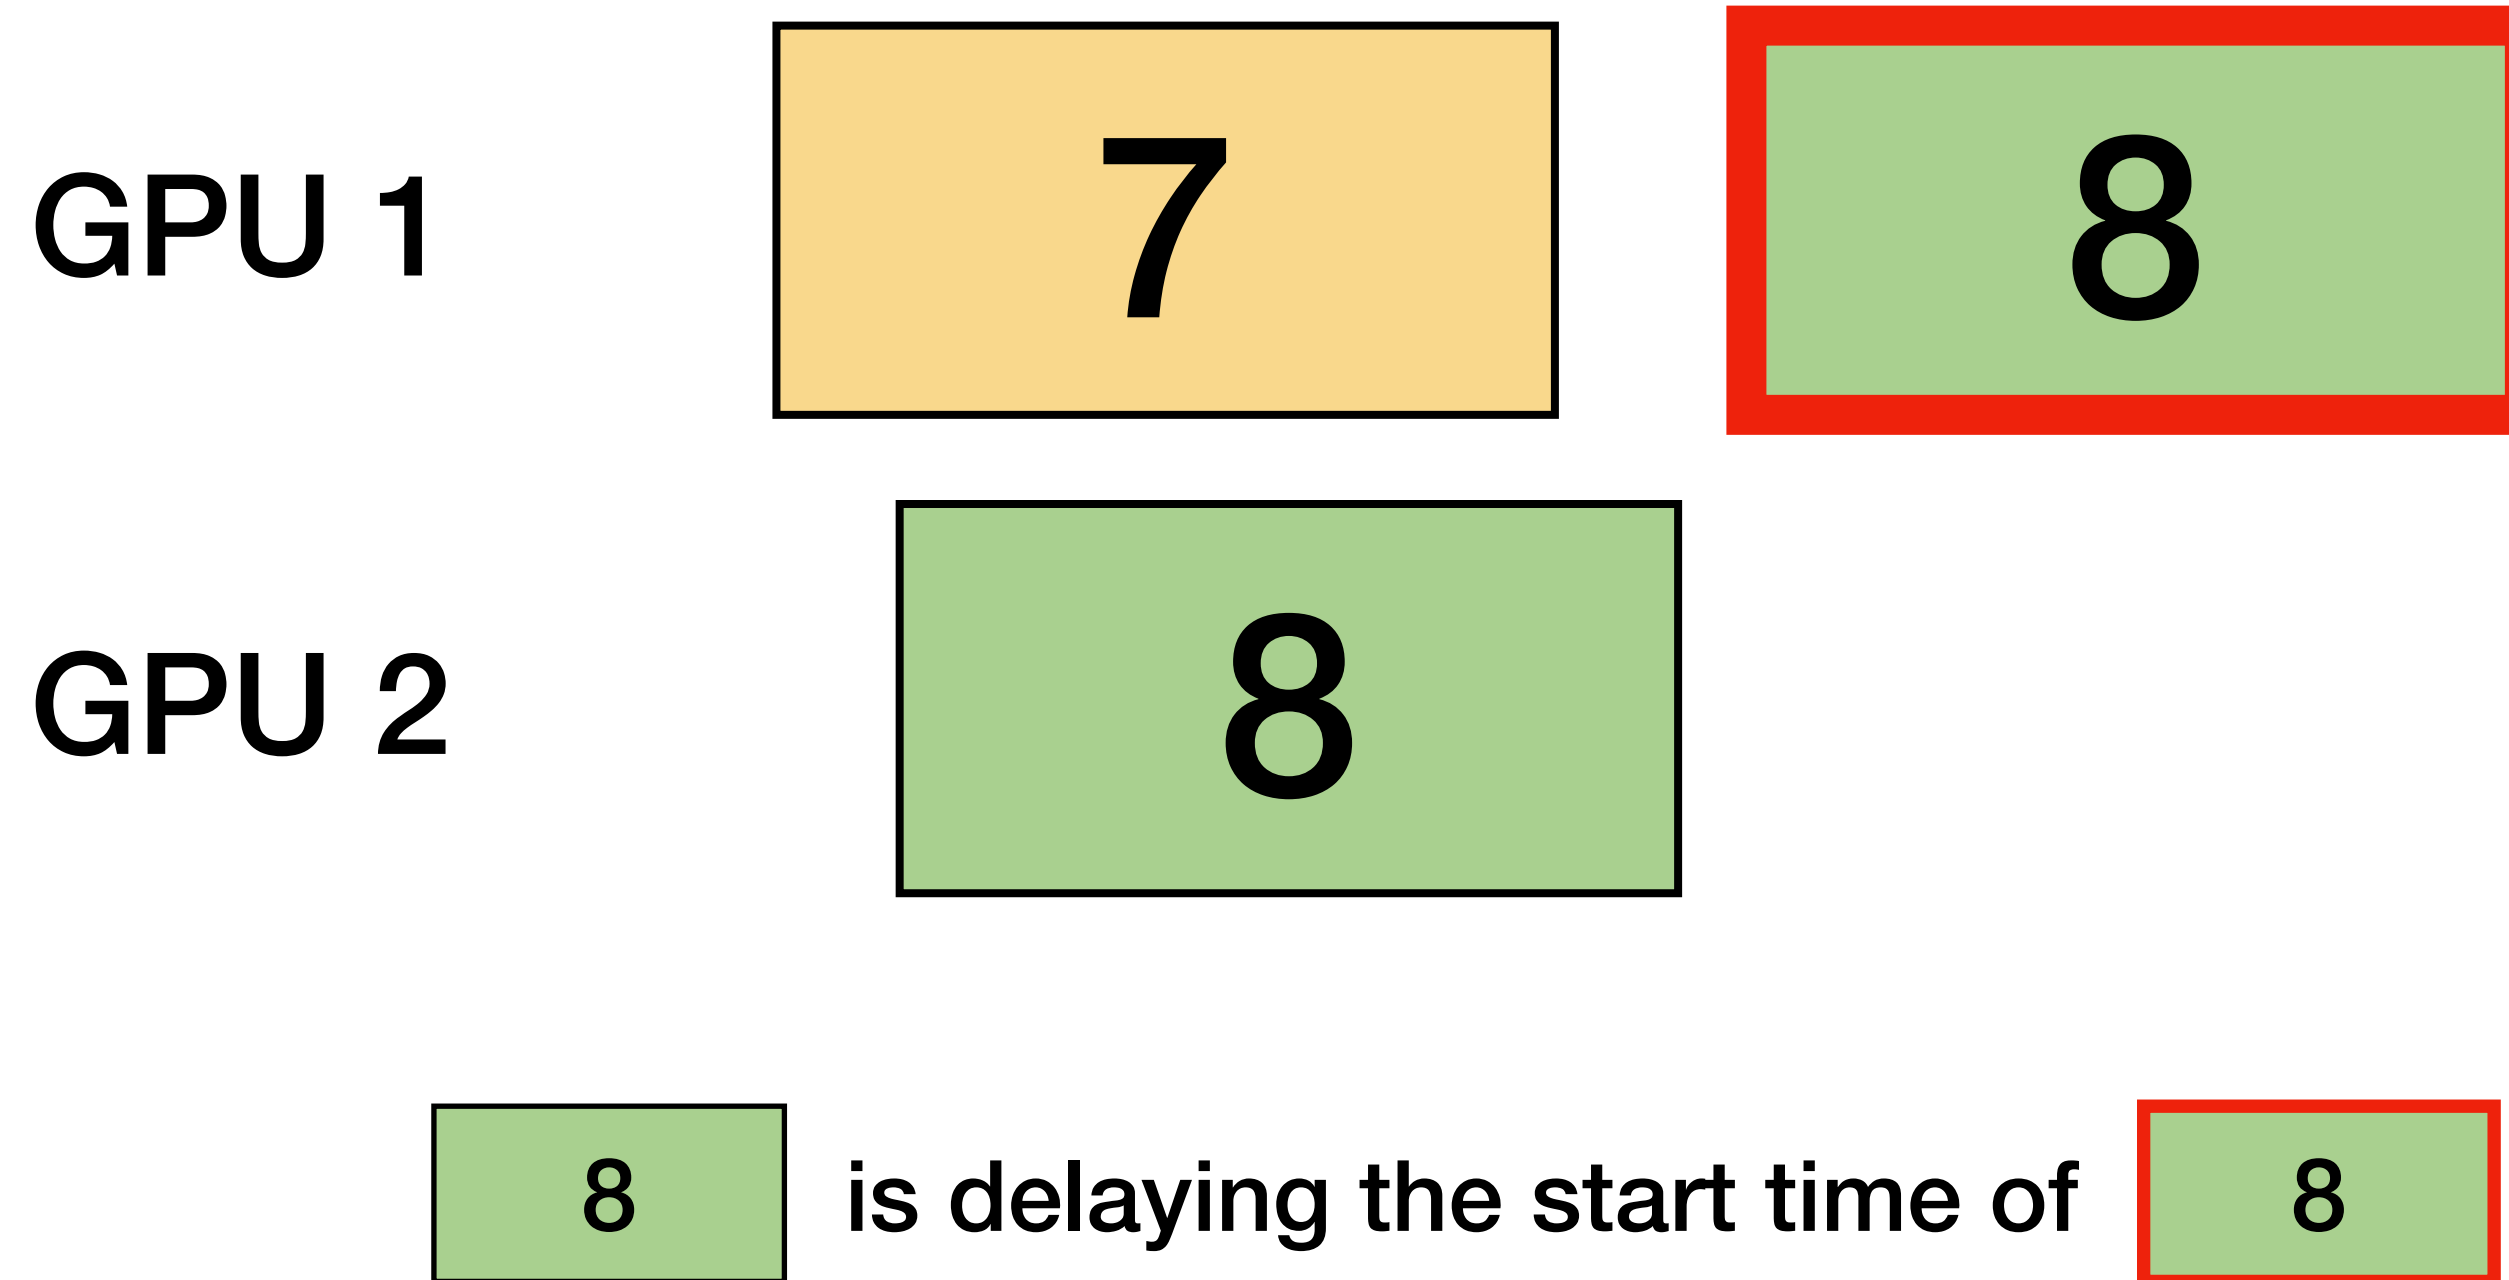
1. Profiler

2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

How to *find* critical path?




1. Profiler

2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

How to *find* critical path?



1. Profiler

2. Scheduler

# 3. Freq. Planner



Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

How to *find* critical path?

GPU 1



GPU 2



is delaying the start time of



, thus critical path extends to



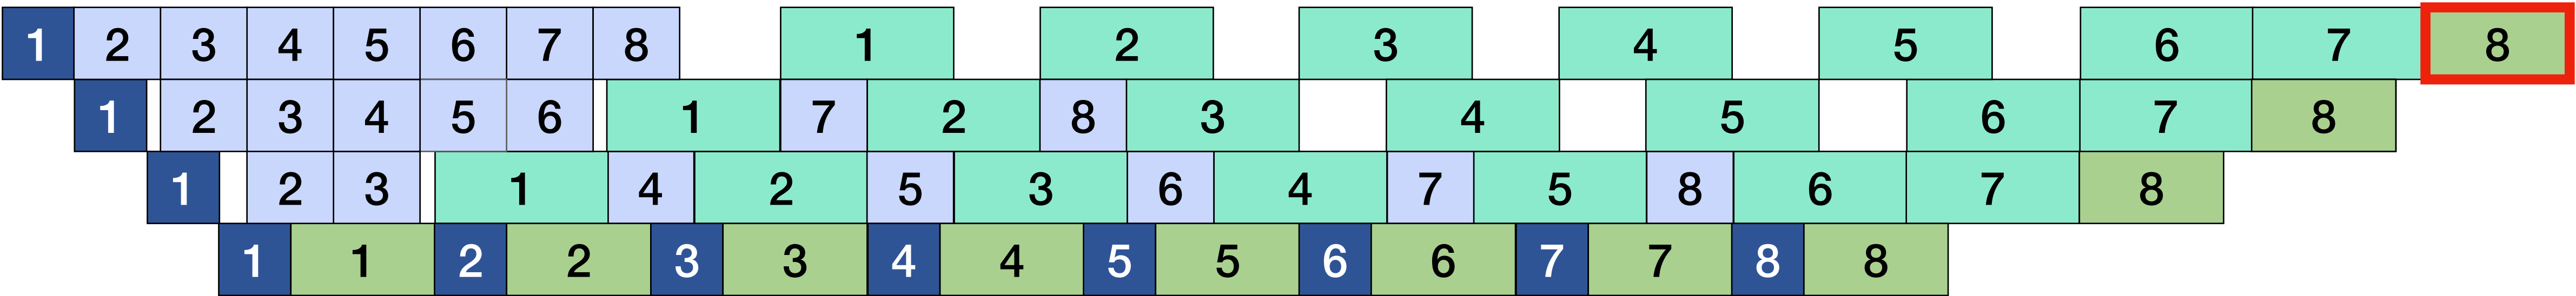
- 1. Profiler
- 2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
 No Accuracy Degradation  
 Minimize Perf. Degradation



- Max Freq Forward
- Reconfigured Forward
- Optimal Freq Forward
- Max Freq Backward
- Reconfigured Backward
- Optimal Freq Backward



**Incrementally build critical path backwards**

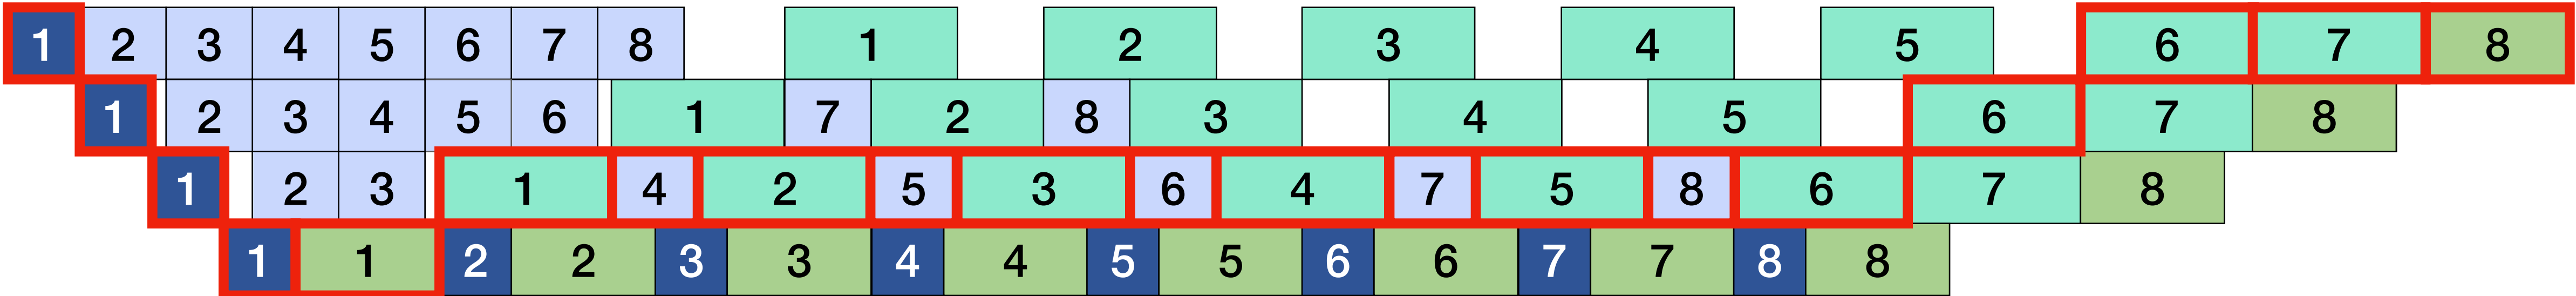
- 1. Profiler
- 2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
 No Accuracy Degradation  
 Minimize Perf. Degradation



<span style="display: inline-block; width: 15px; height: 15px; background-color: #1a3d54; border: 1px solid black;"></span> Max Freq Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #add8e6; border: 1px solid black;"></span> Reconfigured Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #b0c4de; border: 1px solid black;"></span> Optimal Freq Forward
<span style="display: inline-block; width: 15px; height: 15px; background-color: #90ee90; border: 1px solid black;"></span> Max Freq Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #ffcc99; border: 1px solid black;"></span> Reconfigured Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #70e0e0; border: 1px solid black;"></span> Optimal Freq Backward



**Incrementally build critical path backwards**




- 1. Profiler
- 2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
 No Accuracy Degradation  
 Minimize Perf. Degradation



<span style="display: inline-block; width: 15px; height: 15px; background-color: #1a3d54; border: 1px solid black;"></span> Max Freq Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #00aaff; border: 1px solid black;"></span> Reconfigured Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #a0c0ff; border: 1px solid black;"></span> Optimal Freq Forward
<span style="display: inline-block; width: 15px; height: 15px; background-color: #90c060; border: 1px solid black;"></span> Max Freq Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #ffcc66; border: 1px solid black;"></span> Reconfigured Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #66ffcc; border: 1px solid black;"></span> Optimal Freq Backward




**Iteratively finds critical path and reconfigures SM frequency on the critical path**


- 1. Profiler
- 2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
 No Accuracy Degradation  
 Minimize Perf. Degradation



<span style="display: inline-block; width: 15px; height: 15px; background-color: #1a3d54; border: 1px solid black;"></span> Max Freq Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #00aaff; border: 1px solid black;"></span> Reconfigured Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #a6c9ff; border: 1px solid black;"></span> Optimal Freq Forward
<span style="display: inline-block; width: 15px; height: 15px; background-color: #90c070; border: 1px solid black;"></span> Max Freq Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #ffcc66; border: 1px solid black;"></span> Reconfigured Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #66ffcc; border: 1px solid black;"></span> Optimal Freq Backward



**Iteratively finds critical path and reconfigures SM frequency on the critical path**

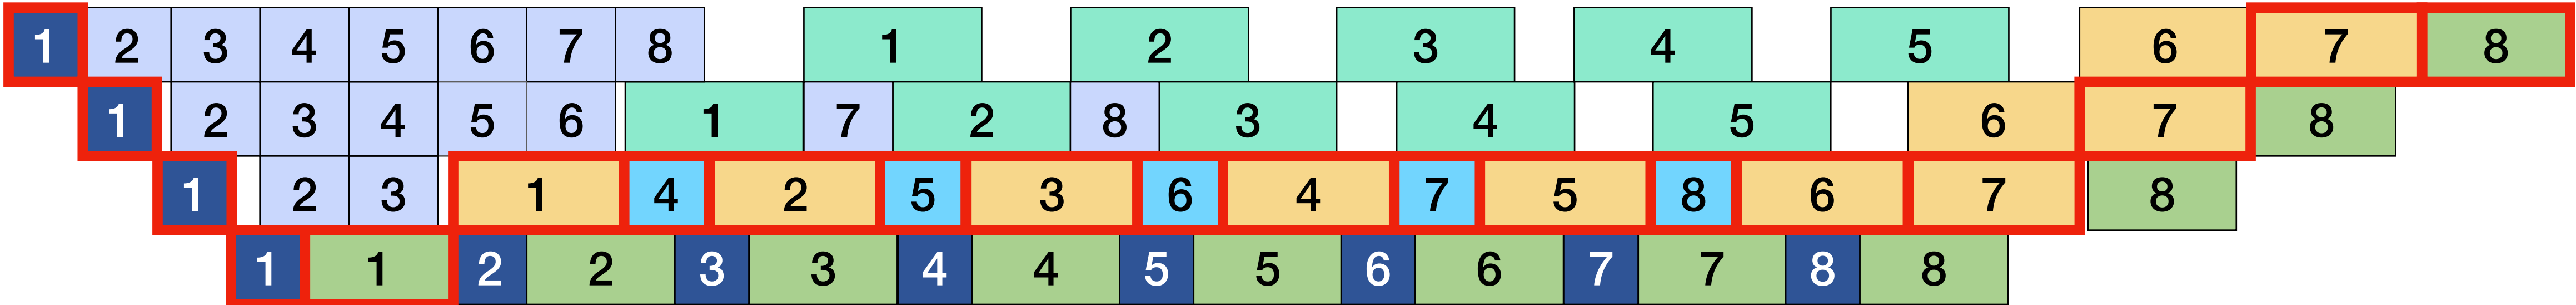
- 1. Profiler
- 2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
 No Accuracy Degradation  
 Minimize Perf. Degradation



<span style="display: inline-block; width: 15px; height: 15px; background-color: #1a3d54; border: 1px solid black;"></span> Max Freq Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #00aaff; border: 1px solid black;"></span> Reconfigured Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #a6c9ff; border: 1px solid black;"></span> Optimal Freq Forward
<span style="display: inline-block; width: 15px; height: 15px; background-color: #90c070; border: 1px solid black;"></span> Max Freq Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #ffcc66; border: 1px solid black;"></span> Reconfigured Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #66ffcc; border: 1px solid black;"></span> Optimal Freq Backward




**Iteratively finds critical path and reconfigures SM frequency on the critical path**


- 1. Profiler
- 2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
 No Accuracy Degradation  
 Minimize Perf. Degradation



<span style="display: inline-block; width: 15px; height: 15px; background-color: #1a3d54; border: 1px solid black;"></span> Max Freq Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #00aaff; border: 1px solid black;"></span> Reconfigured Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #a6c1ff; border: 1px solid black;"></span> Optimal Freq Forward
<span style="display: inline-block; width: 15px; height: 15px; background-color: #90c060; border: 1px solid black;"></span> Max Freq Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #ffcc66; border: 1px solid black;"></span> Reconfigured Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #66ffcc; border: 1px solid black;"></span> Optimal Freq Backward



**Iteratively finds critical path and reconfigures SM frequency on the critical path**

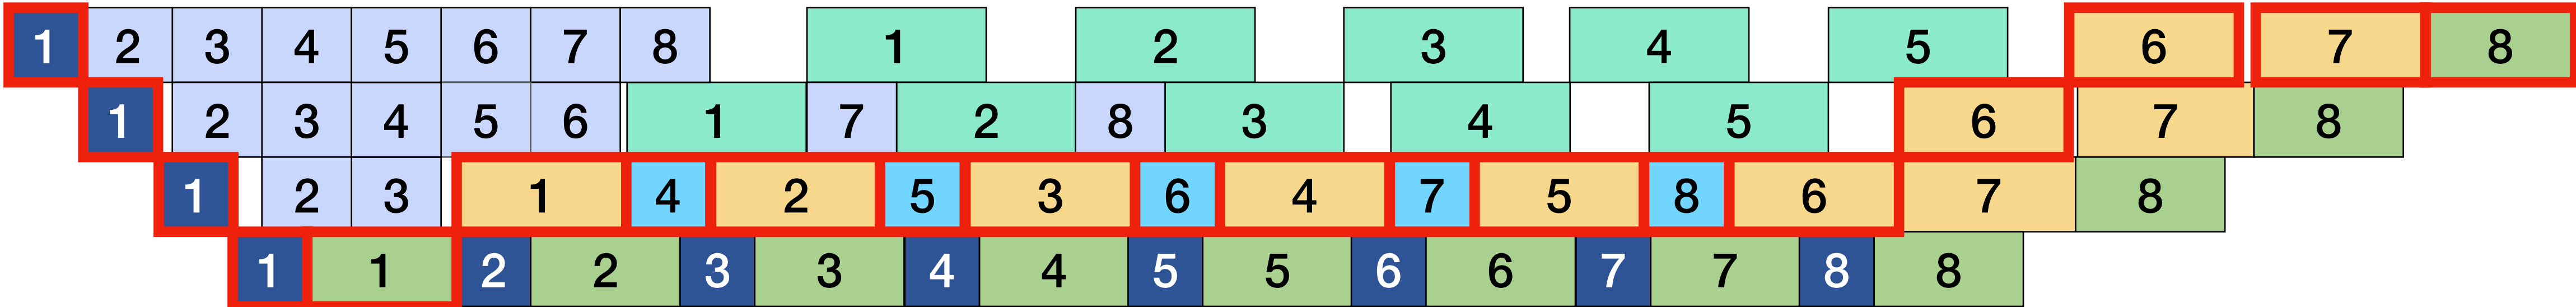
- 1. Profiler
- 2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
 No Accuracy Degradation  
 Minimize Perf. Degradation



<span style="display: inline-block; width: 15px; height: 15px; background-color: #1a3d54; border: 1px solid black;"></span> Max Freq Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #00aaff; border: 1px solid black;"></span> Reconfigured Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #a6c9ff; border: 1px solid black;"></span> Optimal Freq Forward
<span style="display: inline-block; width: 15px; height: 15px; background-color: #90c070; border: 1px solid black;"></span> Max Freq Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #ffcc66; border: 1px solid black;"></span> Reconfigured Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #66ffcc; border: 1px solid black;"></span> Optimal Freq Backward




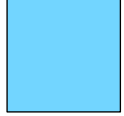
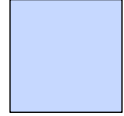

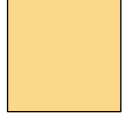
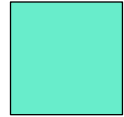
**Iteratively finds critical path and reconfigures SM frequency on the critical path**

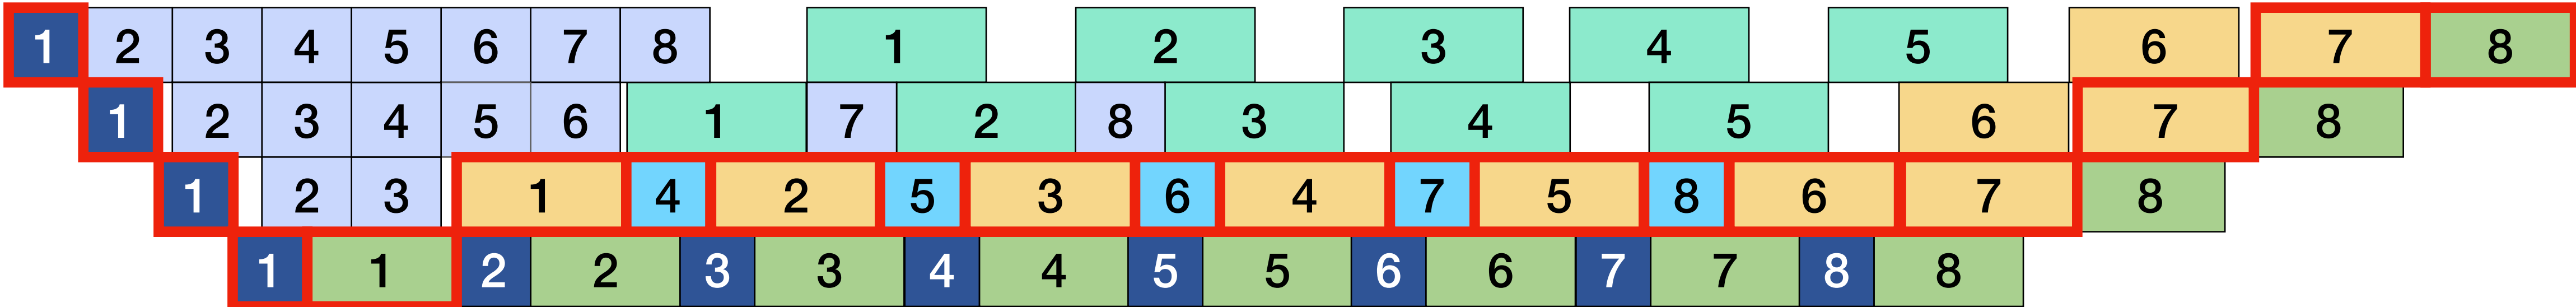
- 1. Profiler
- 2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
 No Accuracy Degradation  
 Minimize Perf. Degradation



	Max Freq Forward		Reconfigured Forward		Optimal Freq Forward
	Max Freq Backward		Reconfigured Backward		Optimal Freq Backward




**Iteratively finds critical path and reconfigures SM frequency on the critical path**

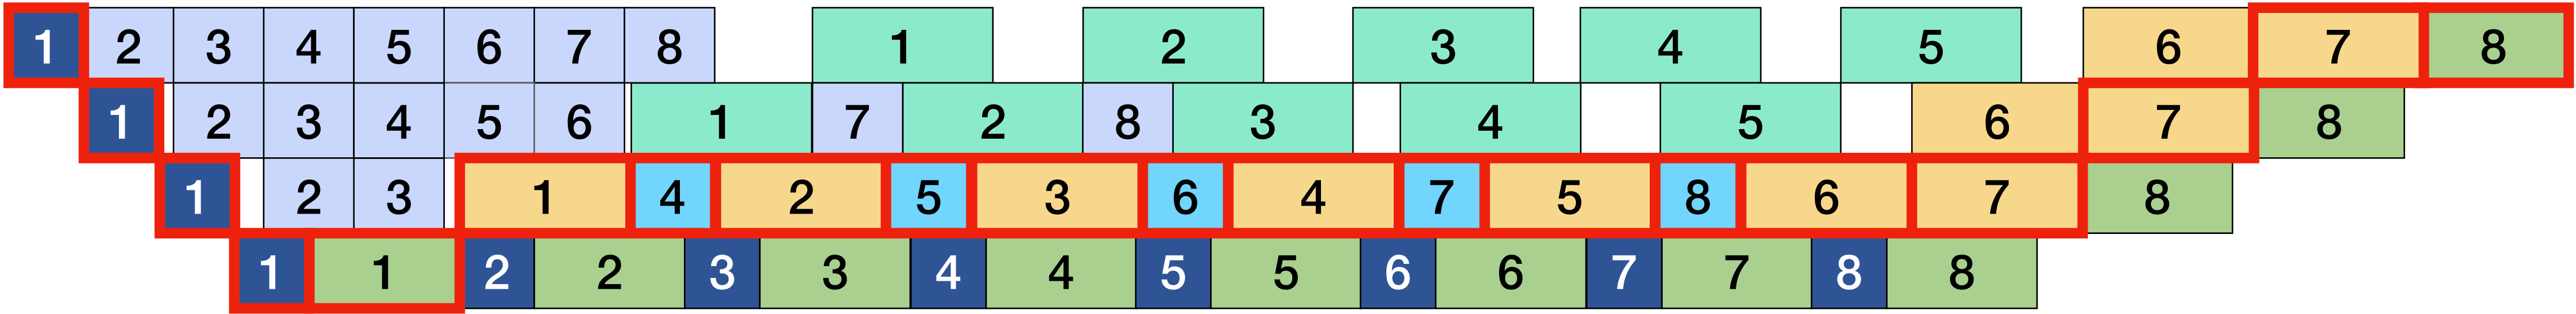
- 1. Profiler
- 2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
 No Accuracy Degradation  
 Minimize Perf. Degradation



<span style="display: inline-block; width: 15px; height: 15px; background-color: #1a3d54; border: 1px solid black;"></span> Max Freq Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #00aaff; border: 1px solid black;"></span> Reconfigured Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #a6c9ff; border: 1px solid black;"></span> Optimal Freq Forward
<span style="display: inline-block; width: 15px; height: 15px; background-color: #90c060; border: 1px solid black;"></span> Max Freq Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #ffcc66; border: 1px solid black;"></span> Reconfigured Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #66ffcc; border: 1px solid black;"></span> Optimal Freq Backward




**Iteratively finds critical path and reconfigures SM frequency on the critical path**

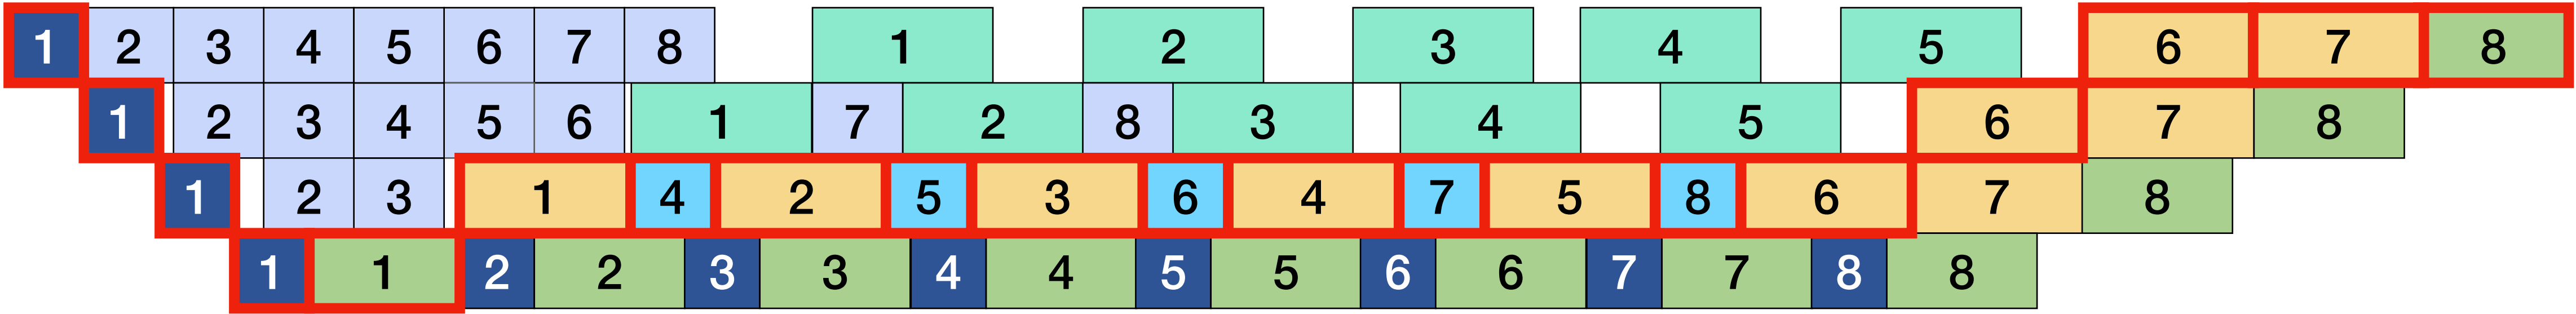
- 1. Profiler
- 2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
 No Accuracy Degradation  
 Minimize Perf. Degradation



<span style="display: inline-block; width: 15px; height: 15px; background-color: #1a3d54; border: 1px solid black;"></span> Max Freq Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #00aaff; border: 1px solid black;"></span> Reconfigured Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #a6c1ff; border: 1px solid black;"></span> Optimal Freq Forward
<span style="display: inline-block; width: 15px; height: 15px; background-color: #90c060; border: 1px solid black;"></span> Max Freq Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #ffcc66; border: 1px solid black;"></span> Reconfigured Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #66ffcc; border: 1px solid black;"></span> Optimal Freq Backward



**Iteratively finds critical path and reconfigures SM frequency on the critical path**



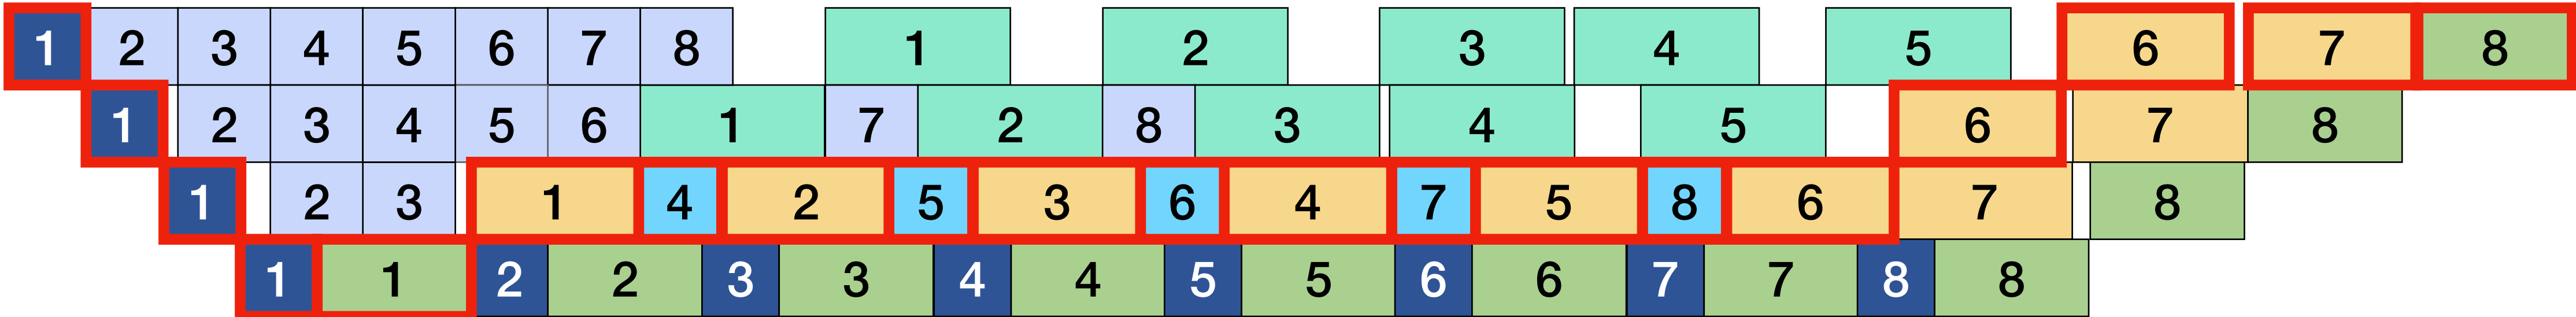
- 1. Profiler
- 2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
 No Accuracy Degradation  
 Minimize Perf. Degradation



<span style="display: inline-block; width: 15px; height: 15px; background-color: #1a3d54; border: 1px solid black;"></span> Max Freq Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #00aaff; border: 1px solid black;"></span> Reconfigured Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #cfe2f3; border: 1px solid black;"></span> Optimal Freq Forward
<span style="display: inline-block; width: 15px; height: 15px; background-color: #90c070; border: 1px solid black;"></span> Max Freq Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #f4b084; border: 1px solid black;"></span> Reconfigured Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #76d7c4; border: 1px solid black;"></span> Optimal Freq Backward



**Iteratively finds critical path and reconfigures SM frequency on the critical path**

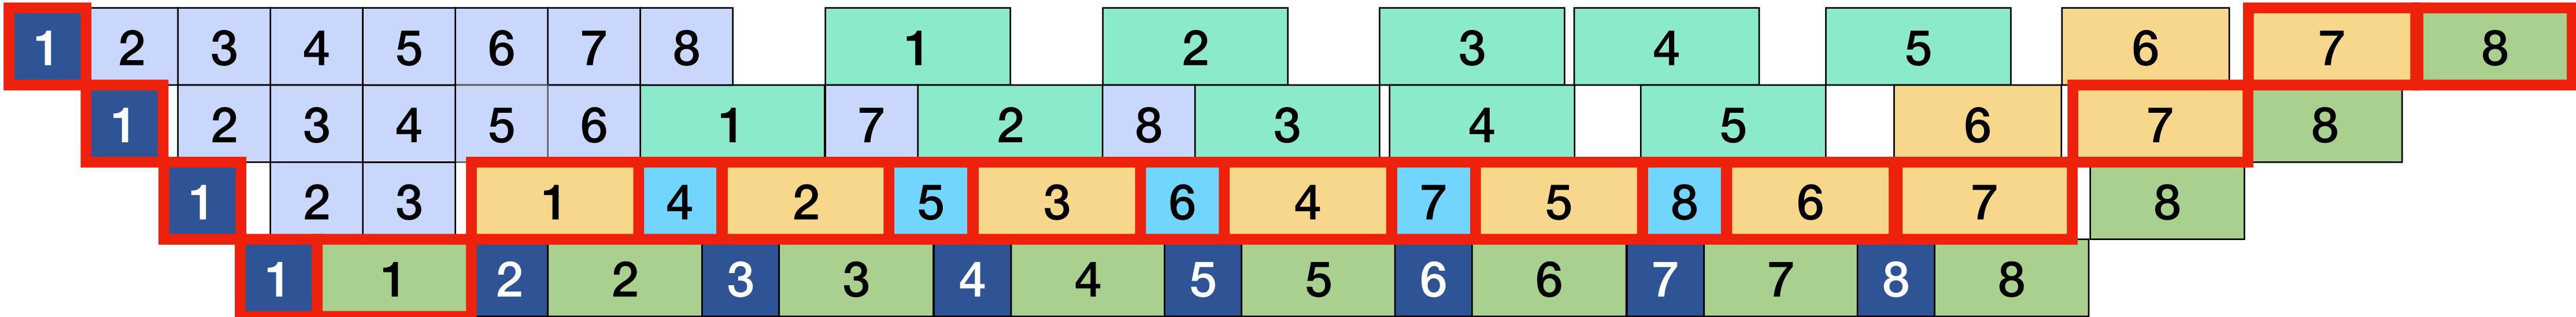
- 1. Profiler
- 2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
 No Accuracy Degradation  
 Minimize Perf. Degradation



<span style="display: inline-block; width: 15px; height: 15px; background-color: #1a3d54; border: 1px solid black;"></span> Max Freq Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #00aaff; border: 1px solid black;"></span> Reconfigured Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #a6c9ff; border: 1px solid black;"></span> Optimal Freq Forward
<span style="display: inline-block; width: 15px; height: 15px; background-color: #90c060; border: 1px solid black;"></span> Max Freq Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #ffcc66; border: 1px solid black;"></span> Reconfigured Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #66ffcc; border: 1px solid black;"></span> Optimal Freq Backward



**Iteratively finds critical path and reconfigures SM frequency on the critical path**

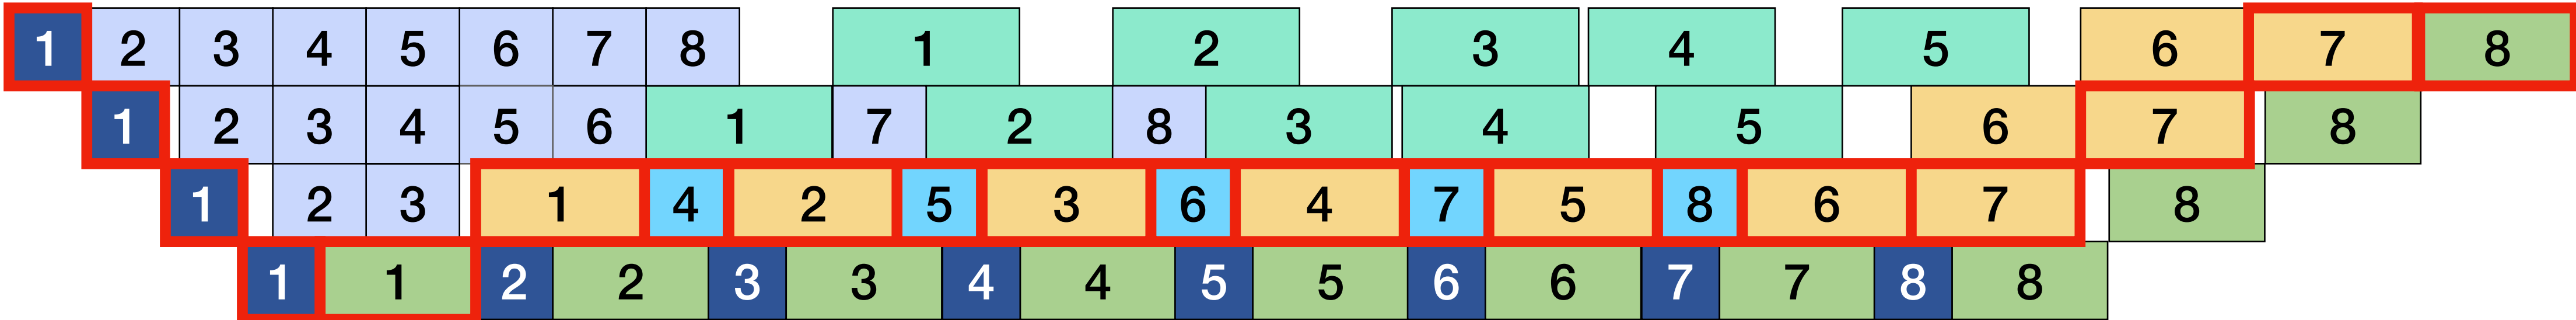
- 1. Profiler
- 2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
 No Accuracy Degradation  
 Minimize Perf. Degradation



<span style="display: inline-block; width: 15px; height: 15px; background-color: #1a3d54; border: 1px solid black;"></span> Max Freq Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #00aaff; border: 1px solid black;"></span> Reconfigured Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #a0c4ff; border: 1px solid black;"></span> Optimal Freq Forward
<span style="display: inline-block; width: 15px; height: 15px; background-color: #90c060; border: 1px solid black;"></span> Max Freq Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #ffcc66; border: 1px solid black;"></span> Reconfigured Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #66ffcc; border: 1px solid black;"></span> Optimal Freq Backward



**Iteratively finds critical path and reconfigures SM frequency on the critical path**

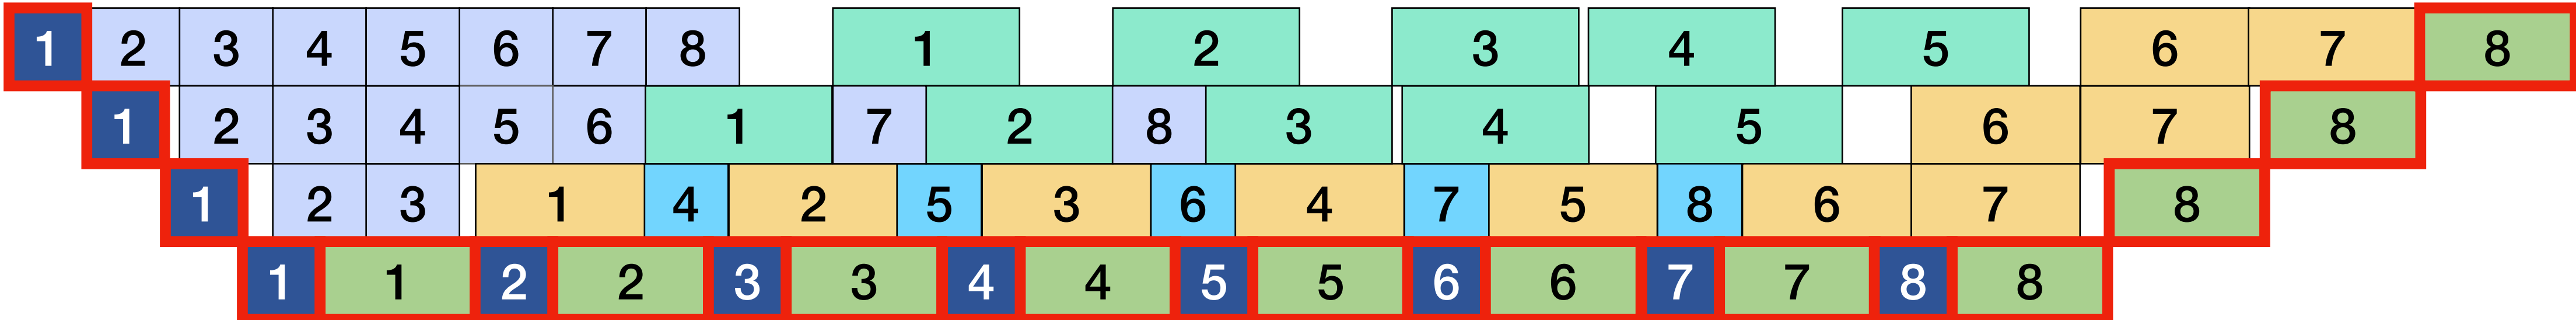
- 1. Profiler
- 2. Scheduler

# 3. Freq. Planner

Maximize Energy Saving  
 No Accuracy Degradation  
 Minimize Perf. Degradation



<span style="display: inline-block; width: 15px; height: 15px; background-color: #1a3d54; border: 1px solid black;"></span> Max Freq Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #00aaff; border: 1px solid black;"></span> Reconfigured Forward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #a6c9ff; border: 1px solid black;"></span> Optimal Freq Forward
<span style="display: inline-block; width: 15px; height: 15px; background-color: #90c070; border: 1px solid black;"></span> Max Freq Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #ffcc66; border: 1px solid black;"></span> Reconfigured Backward	<span style="display: inline-block; width: 15px; height: 15px; background-color: #00ffcc; border: 1px solid black;"></span> Optimal Freq Backward



**Terminates when critical path is outer envelope**

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

# EnvPipe

Performance-preserving DNN Training Framework for Saving Energy

**1. Profiler**

**2. Scheduler**

**3. Freq. Planner**

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

# EnvPipe

Performance-preserving DNN Training Framework for Saving Energy

## 1. Profiler

Profiles performance-energy trend for each stage

## 2. Scheduler

## 3. Freq. Planner

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

# EnvPipe

Performance-preserving DNN Training Framework for Saving Energy

## 1. Profiler

Profiles performance-energy trend for each stage

## 2. Scheduler

Schedules pipeline units to maximize energy saving

## 3. Freq. Planner

Maximize Energy Saving  
No Accuracy Degradation  
Minimize Perf. Degradation

# EnvPipe

Performance-preserving DNN Training Framework for Saving Energy

## 1. Profiler

Profiles performance-energy trend for each stage

## 2. Scheduler

Schedules pipeline units to maximize energy saving

## 3. Freq. Planner

Reconfigures to minimize performance degradation



# Evaluation: Single-Node

- Single AWS P3.8xLarge with 4 V100 GPUs
- Baseline: run all GPUs with maximum frequency

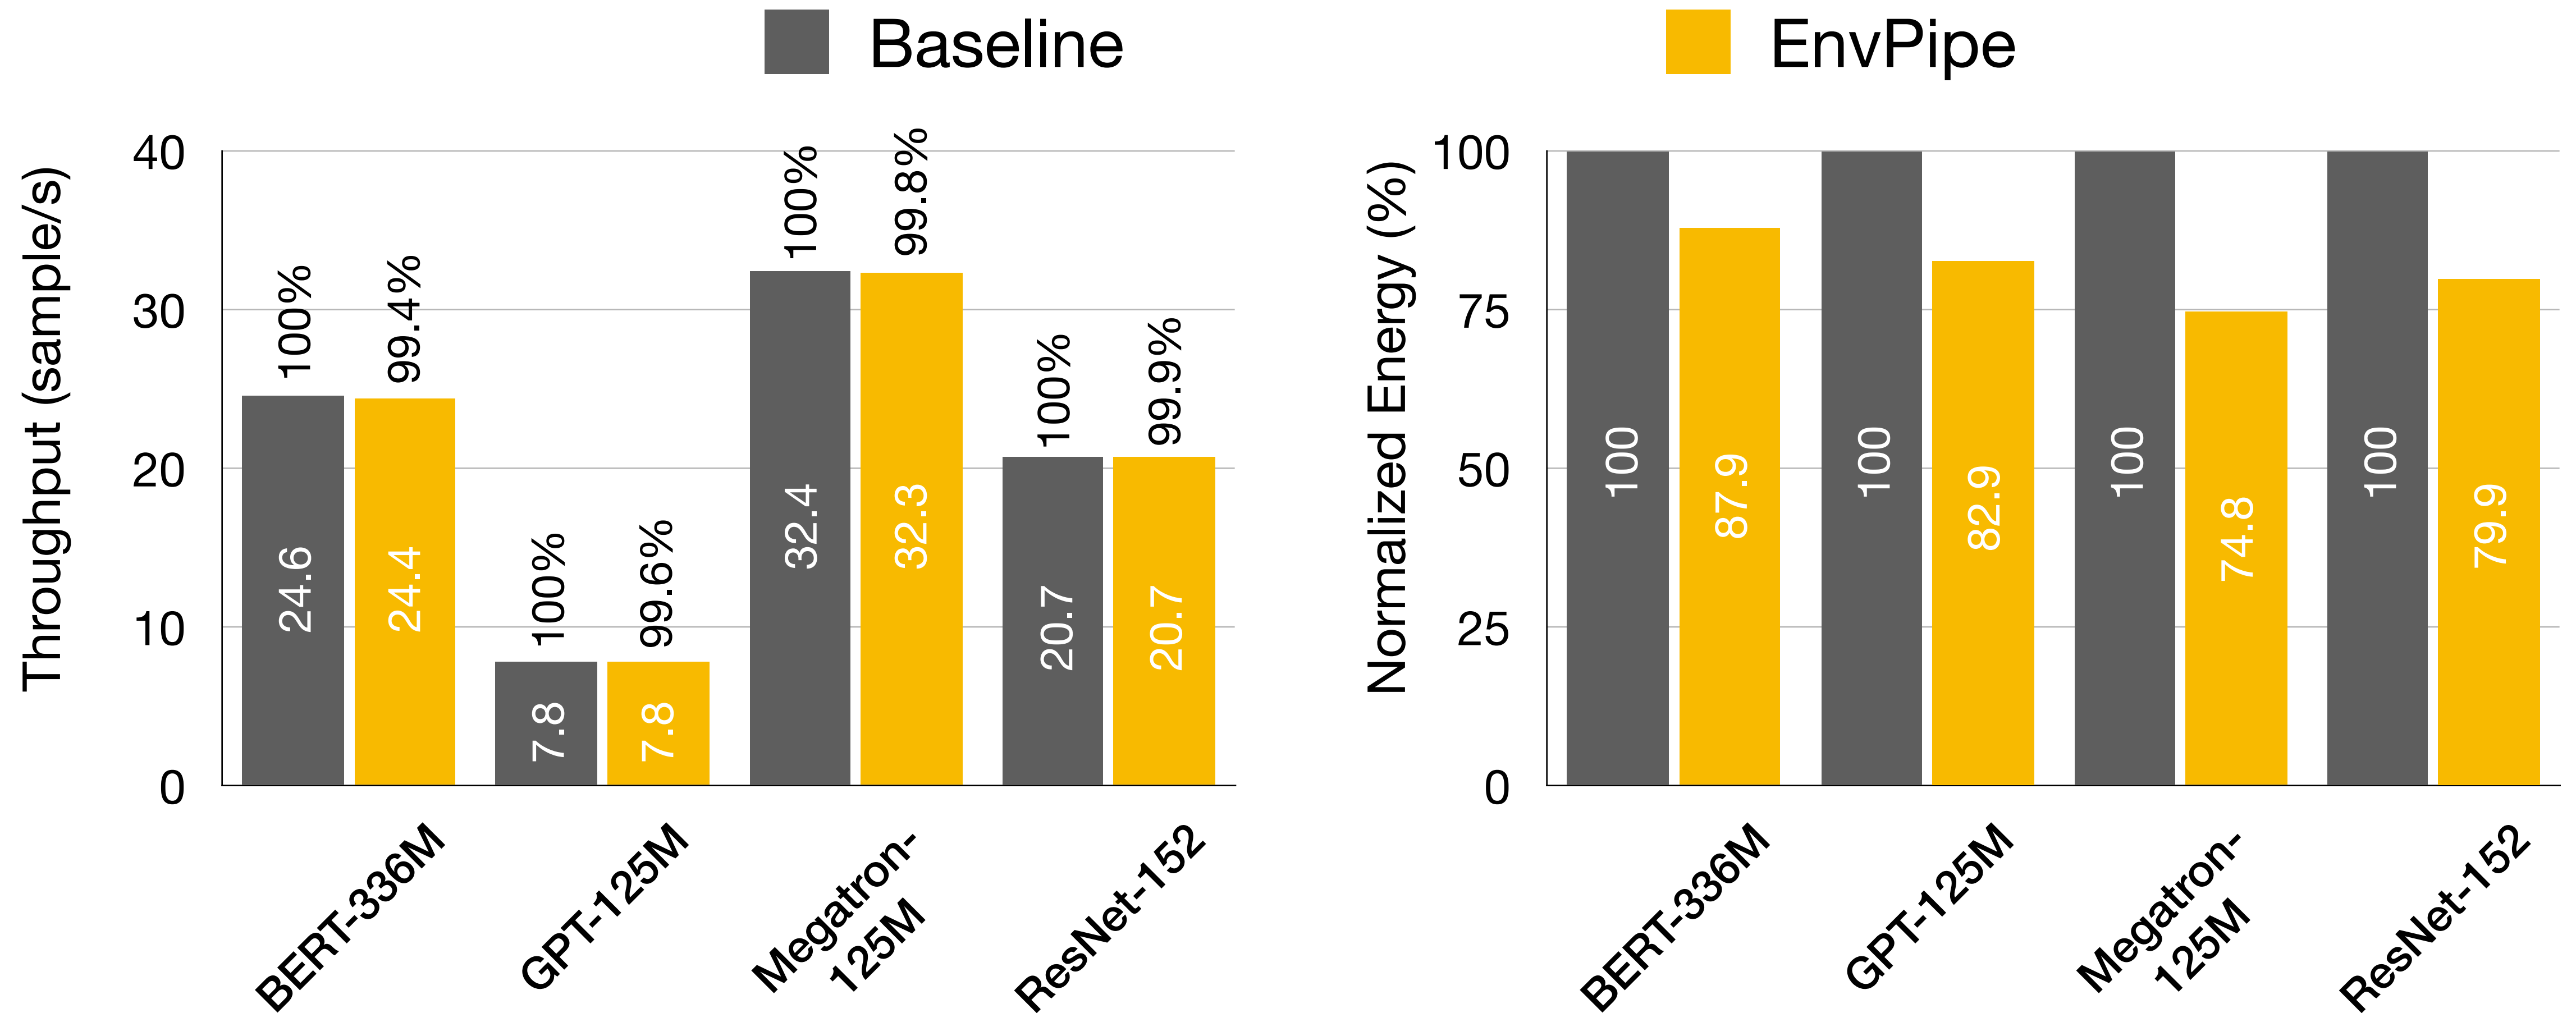
**Maximize Energy Saving**

**No Accuracy Degradation**

**Minimize Perf. Degradation**


Model	Microbatch	Minibatch
BERT-336M	4	64
GPT-125M	2	32
Megatron-125M	4	64
ResNet-152	2	32

# Evaluation: Single-Node




Throughput and energy consumption of single-node training with 4 V100

# Evaluation: Single-Node



Throughput and energy consumption of single-node training with 4 V100

# Evaluation: Single-Node



Throughput and energy consumption of single-node training with 4 V100

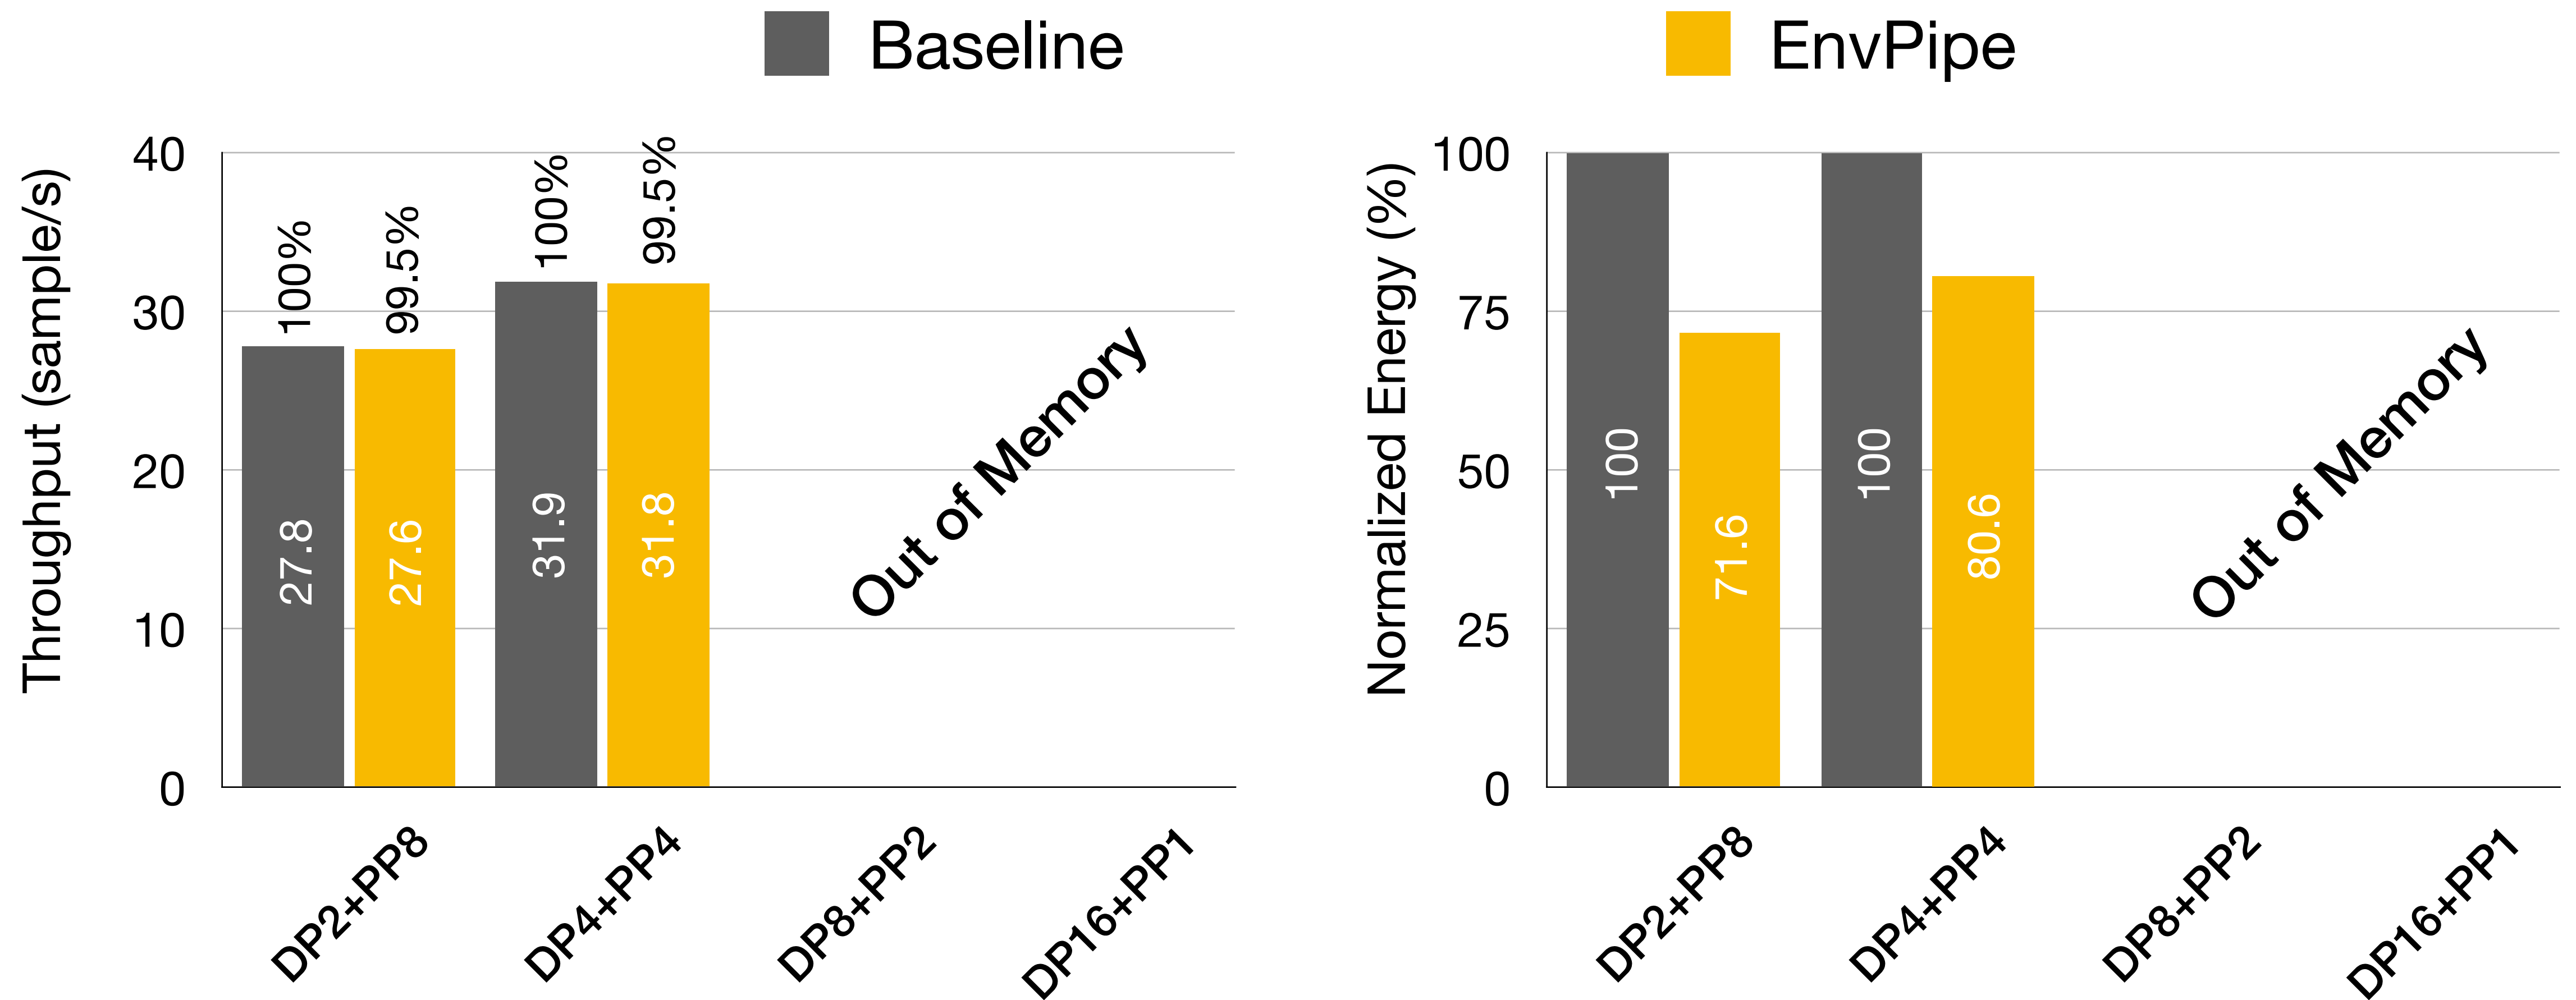
# Evaluation: Multi-Node

- Two AWS P3.16xLarge instances each with 8 V100 GPUs
- Connected with 25Gbps network

DP	PP	Microbatch	Minibatch	Num of Microbatch
2	8	2	512	128
4	4	2	512	64
8	2	2	512	32
16	1	2	512	X

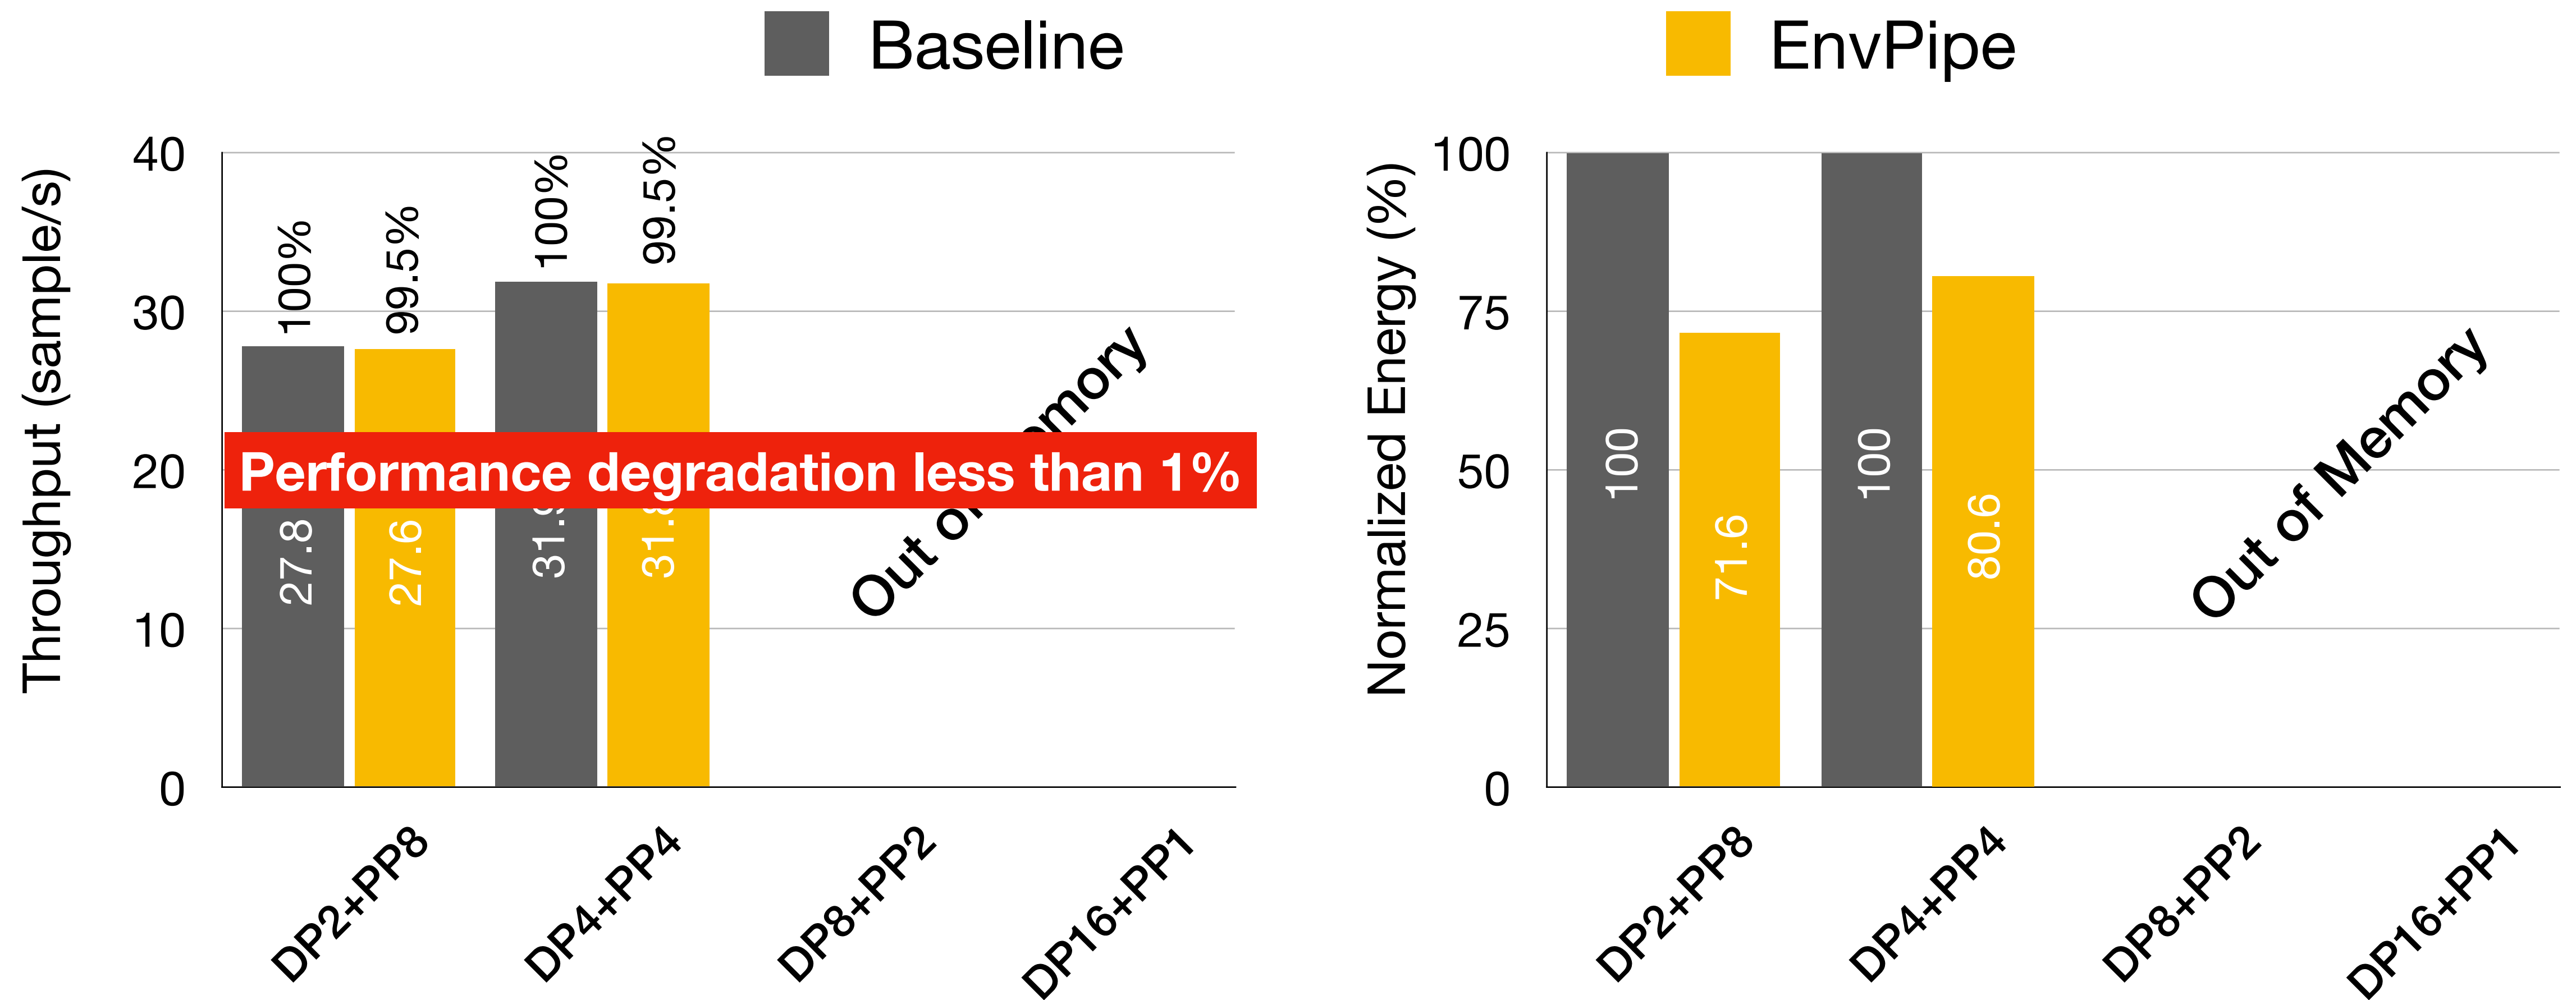
Train with various parallelism dimensions of DP+PP for Megatron-1.3B

# Evaluation: Multi-Node



Throughput and energy consumption of Megatron-1.3B in multi-node training

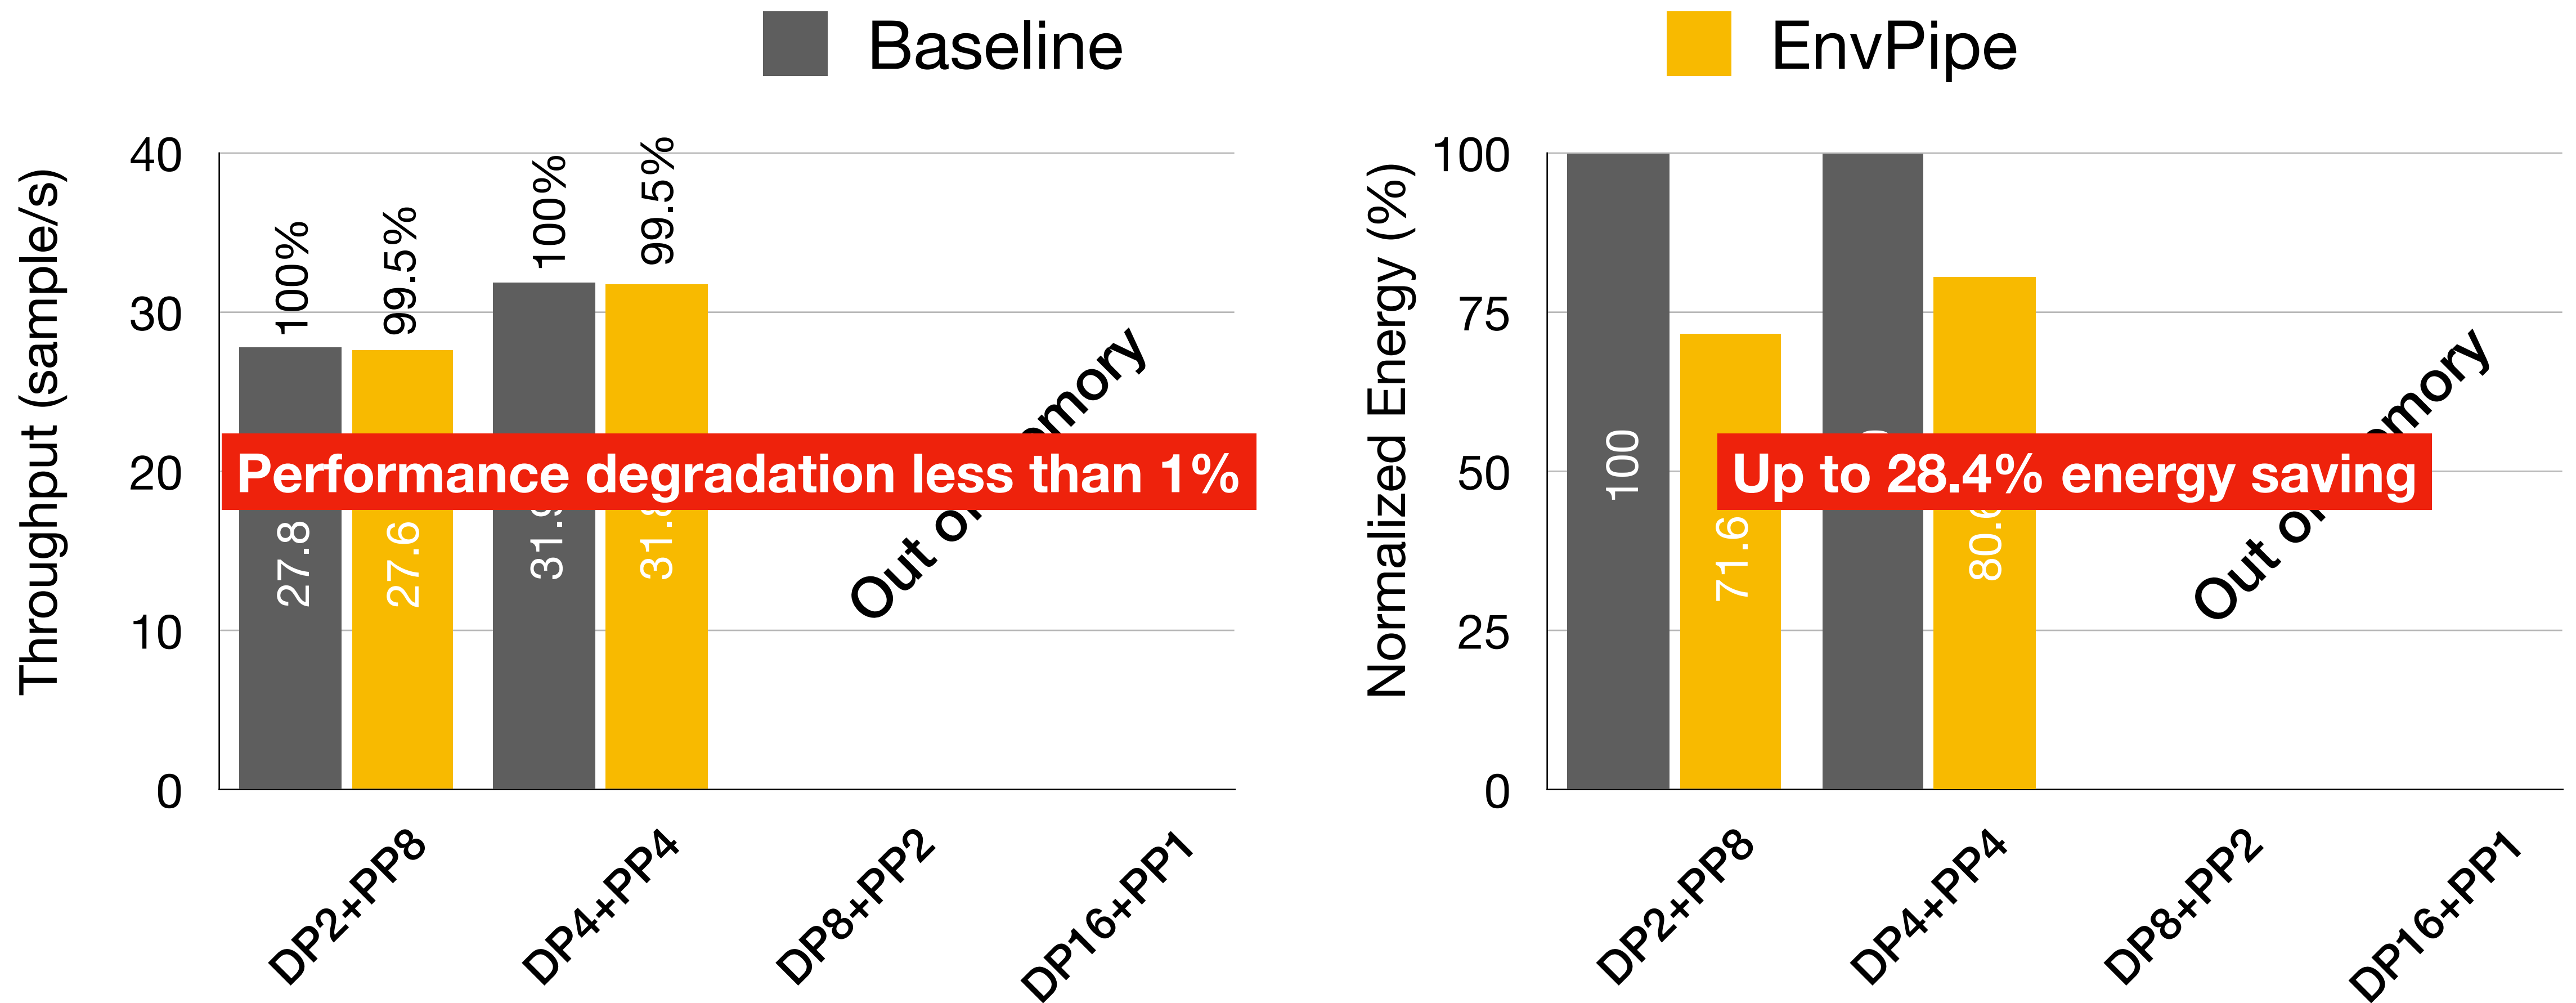
# Evaluation: Multi-Node



Throughput and energy consumption of Megatron-1.3B in multi-node training



# Evaluation: Multi-Node



Throughput and energy consumption of Megatron-1.3B in multi-node training



# Conclusion

# Conclusion

**Goal**

DNN training framework for energy saving without degrading performance or accuracy by leveraging bubbles in pipeline parallelism

# Conclusion

## Goal

DNN training framework for energy saving without degrading performance or accuracy by leveraging bubbles in pipeline parallelism

## Approach

1. Fine-grained online profiler
2. Scheduler to maximize usable bubbles
3. Frequency planner to minimize performance degradation

# Conclusion

## Goal

DNN training framework for energy saving without degrading performance or accuracy by leveraging bubbles in pipeline parallelism

## Approach

1. Fine-grained online profiler
2. Scheduler to maximize usable bubbles
3. Frequency planner to minimize performance degradation

## Result

EnvPipe: Performance-preserving energy-saving DNN training framework  
maximum energy saving up to **28.4%**  
with all cases **less than 1%** performance degradation