sigcomm 2025
COIMBRA

DistTrain: Addressing Model and Data Heterogeneity with Disaggregated Training for Multimodal Large Language Models

Zili Zhang, Yinmin Zhong, Yimin Jiang, Hanpeng Hu, Jianjian Sun,

Zheng Ge, Yibo Zhu, Xin Jin

PEKING UNIVERSITY 北京大学 PEKING UNIVERSITY

StepFun

Multimodal LLM

Unimodal

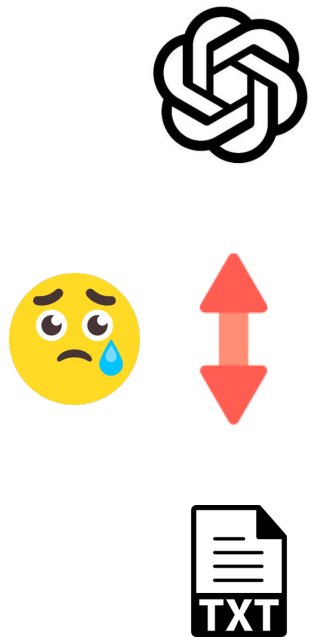

- *natural language*
- *code generation*

Multimodal

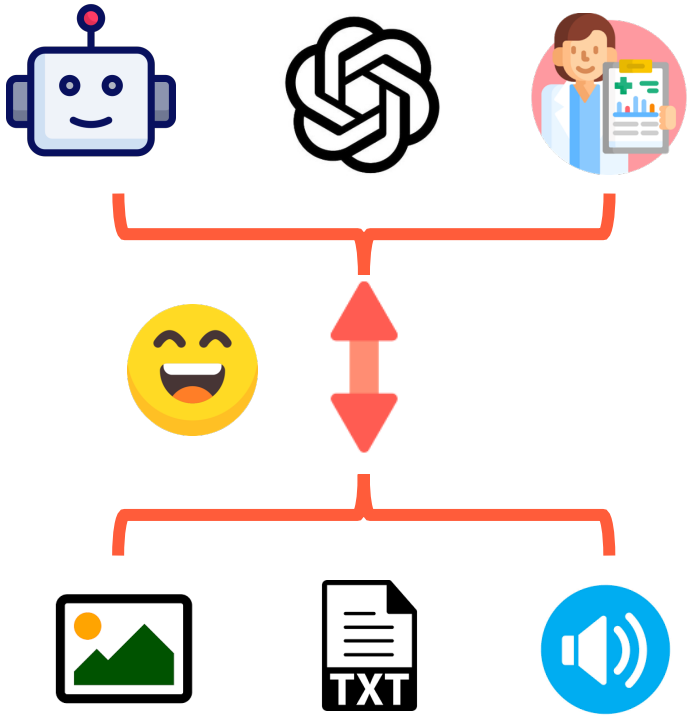

- *natural language, code*
- *photos, diagrams, charts*
- *speech, music, environmental sounds*

Multimodal LLM

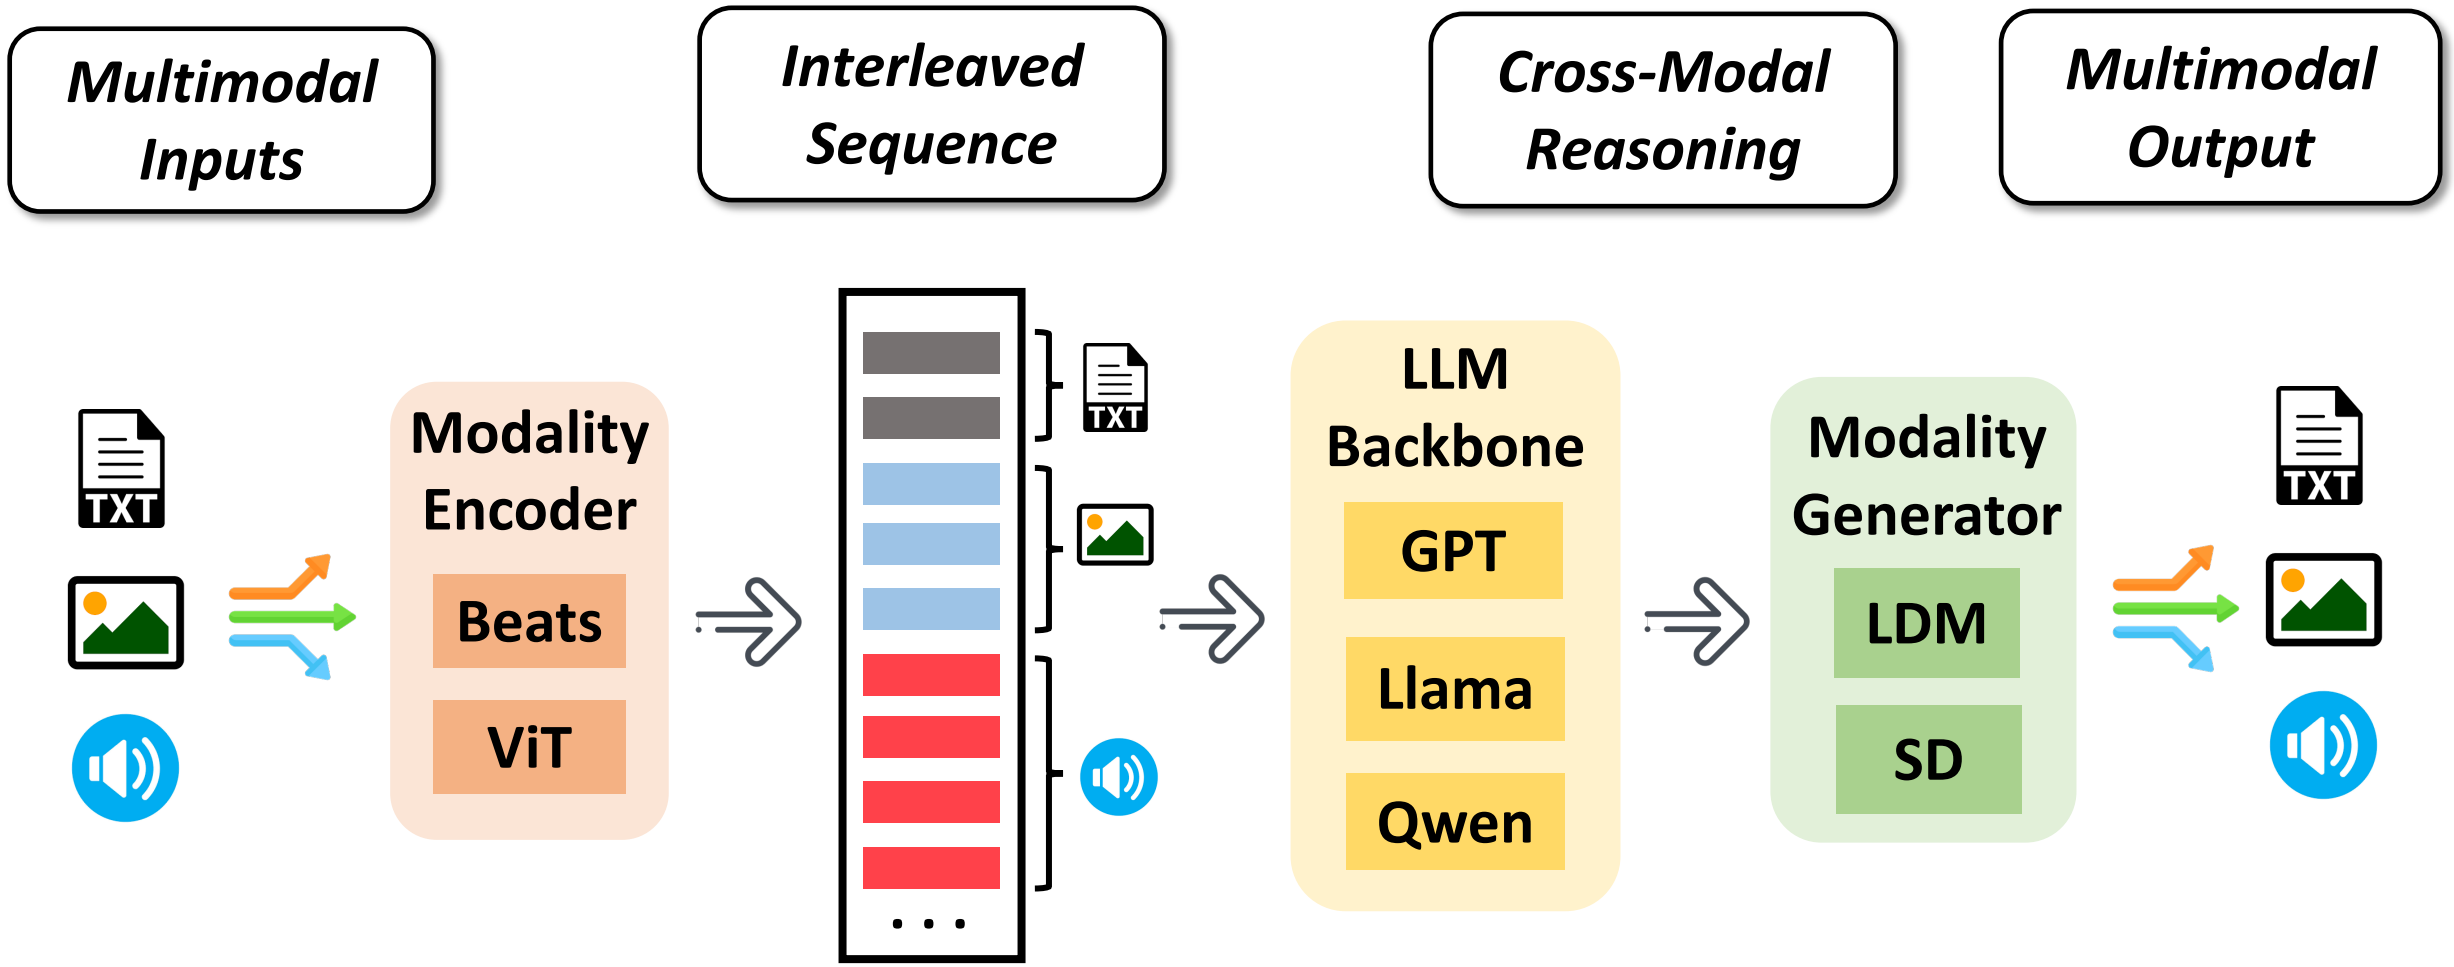

Model Heterogeneity

Modality Encoder

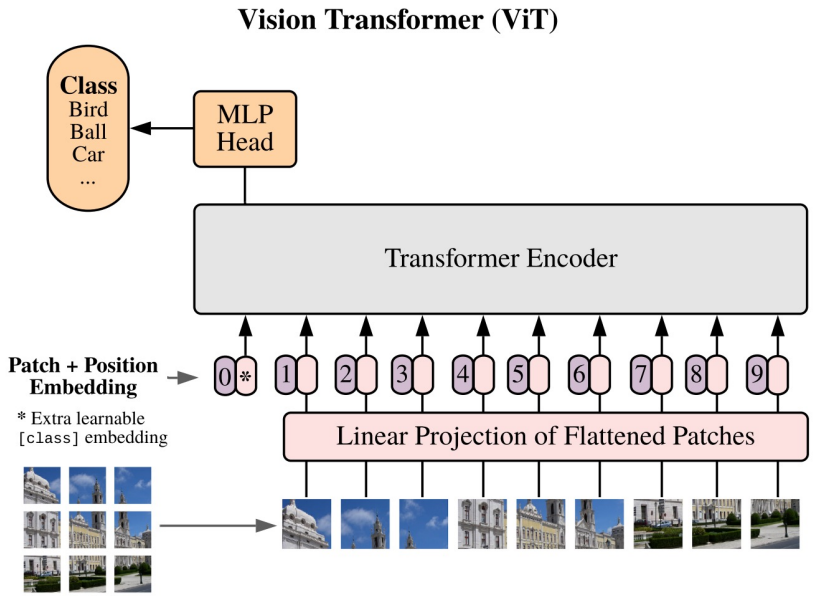

LLM Backbone

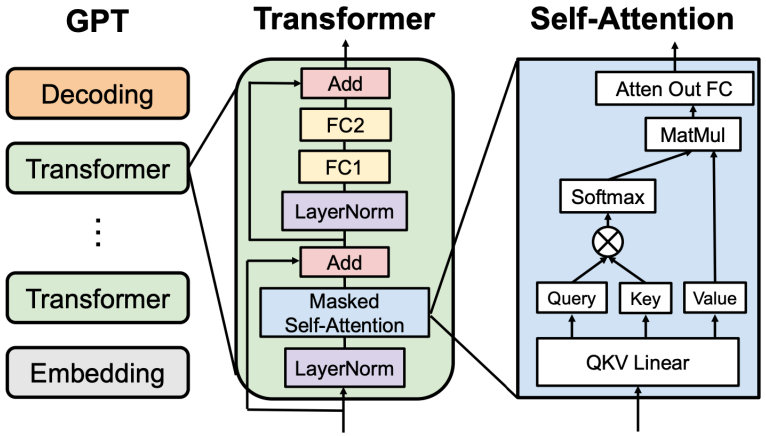

Modality Generator

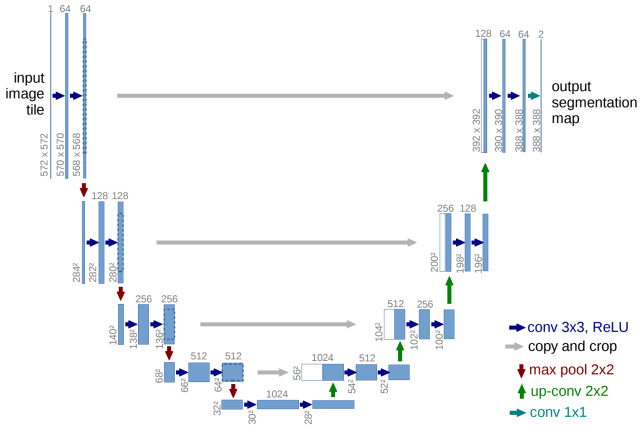

Lightweight Transformer

Heavyweight Transformer

Convolutional U-Net

Model Heterogeneity

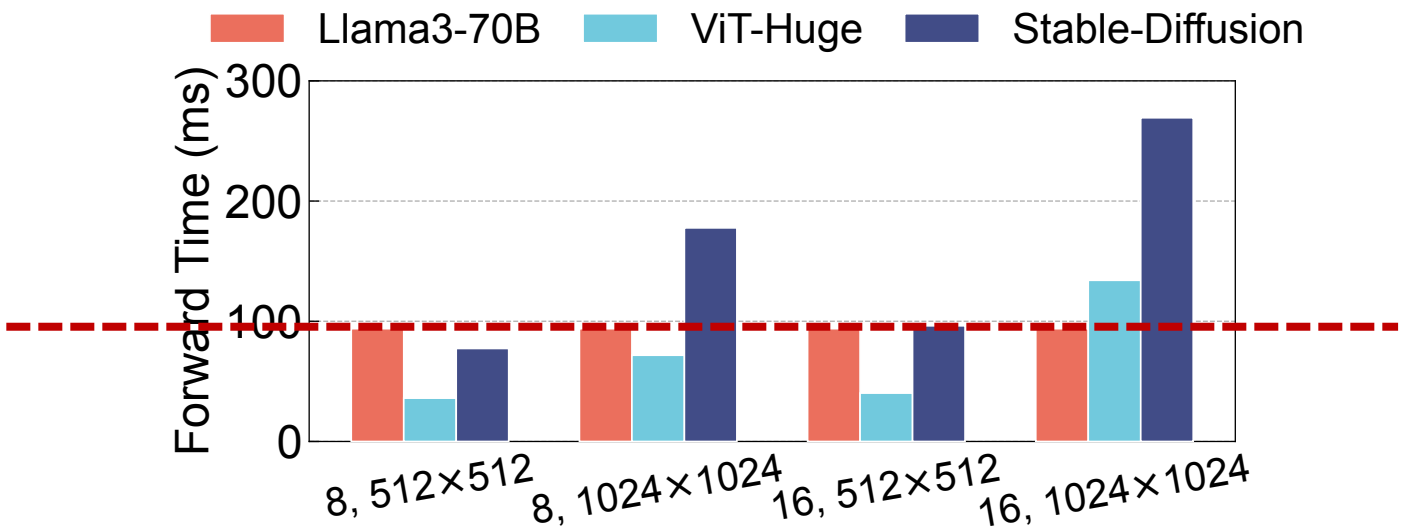

LLM Backbone is the PP straggler

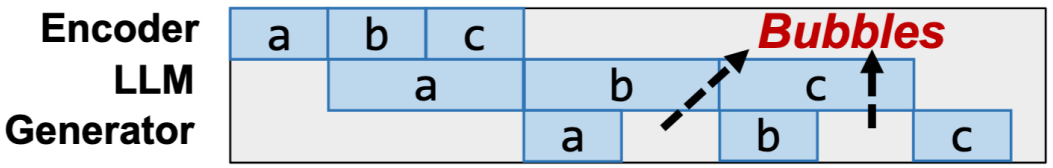

Modality encoder/generator is the PP straggler

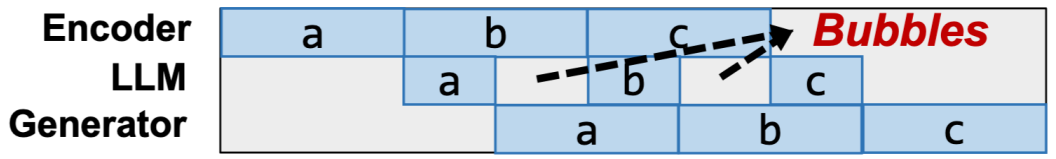

Data Heterogeneity

What is the landmark, and what plants surround it? Beside, remove the plants.

The landmark is the Great Wall, surrounded by flowers. Here is the edited image.

Training Sequence

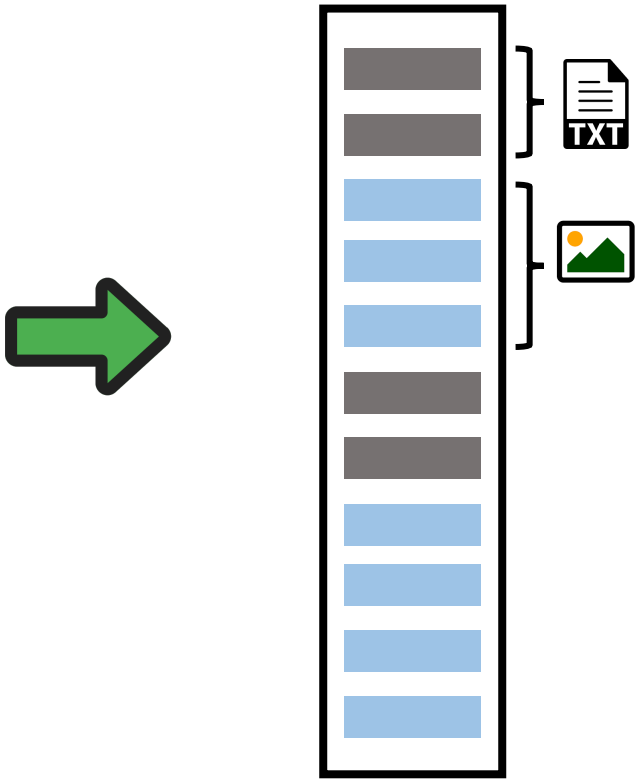

Data Heterogeneity

The sizes of text and image subsequences display highly skewed distributions

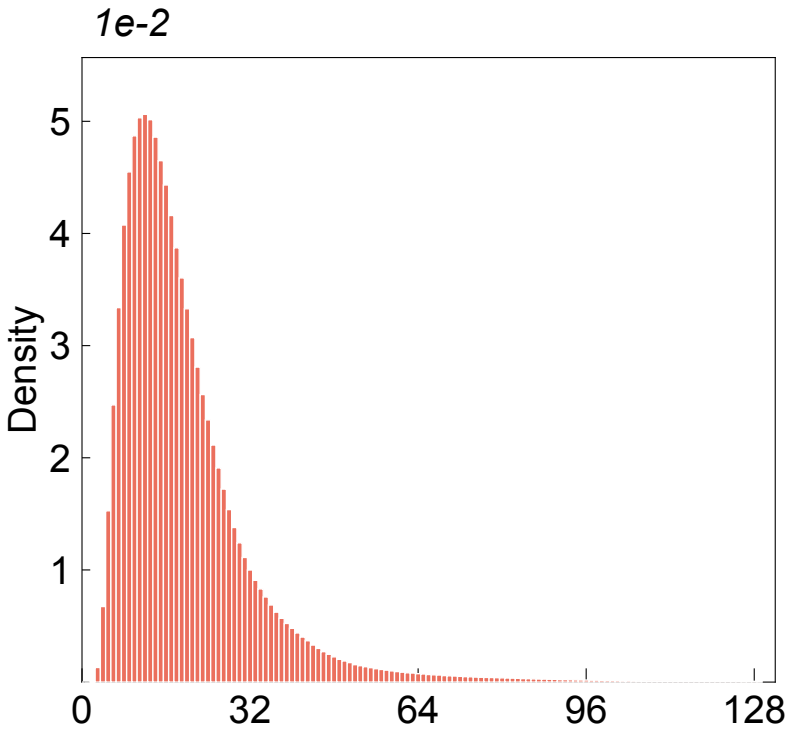

# of tokens for one text subsequence

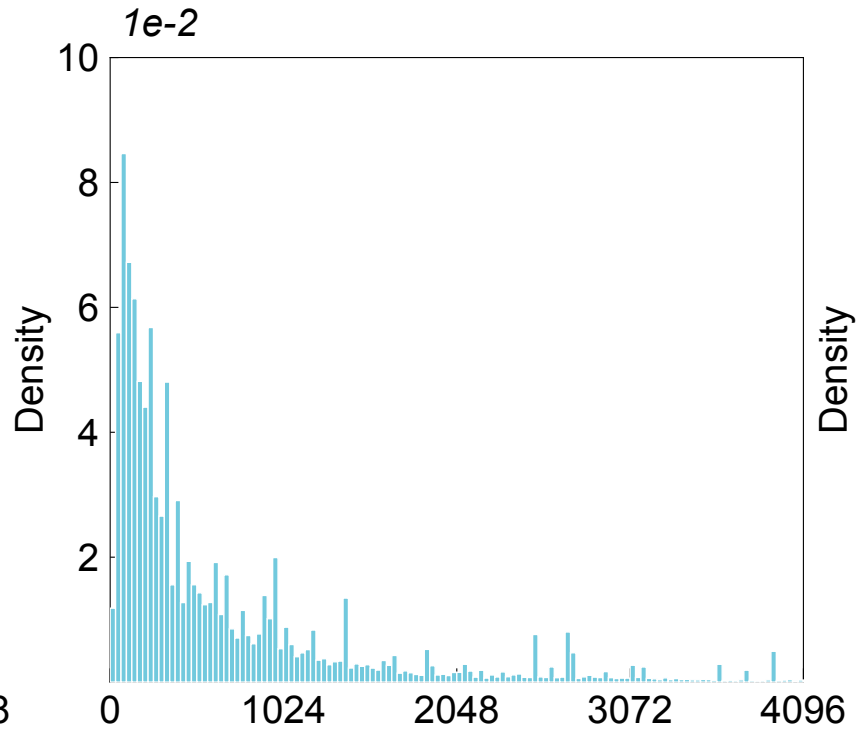

# of tokens for one image subsequence

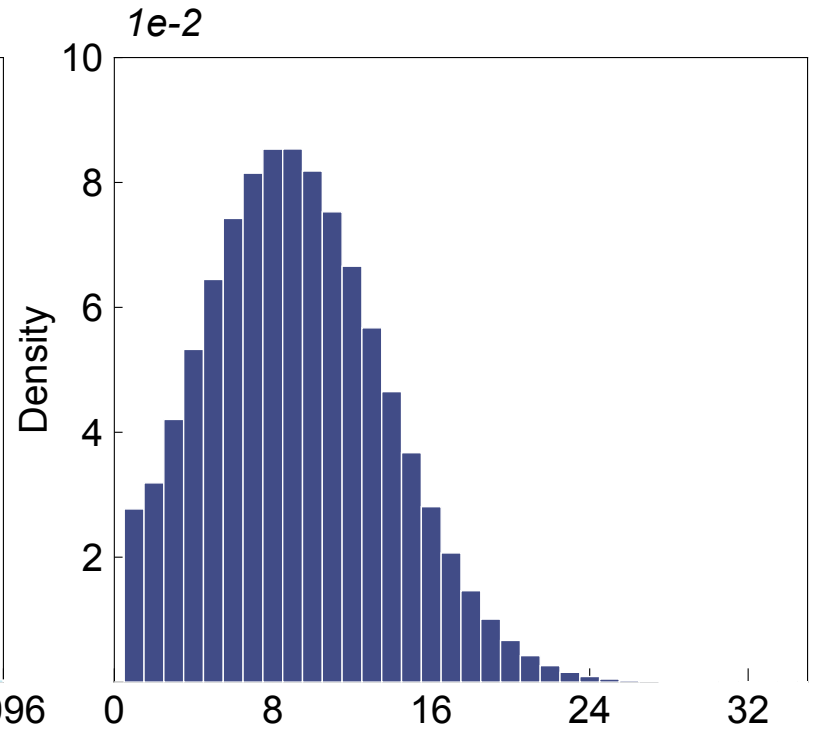

# of image subsequence in one 8K sequence

Data Heterogeneity

Intra-microbatch straggler

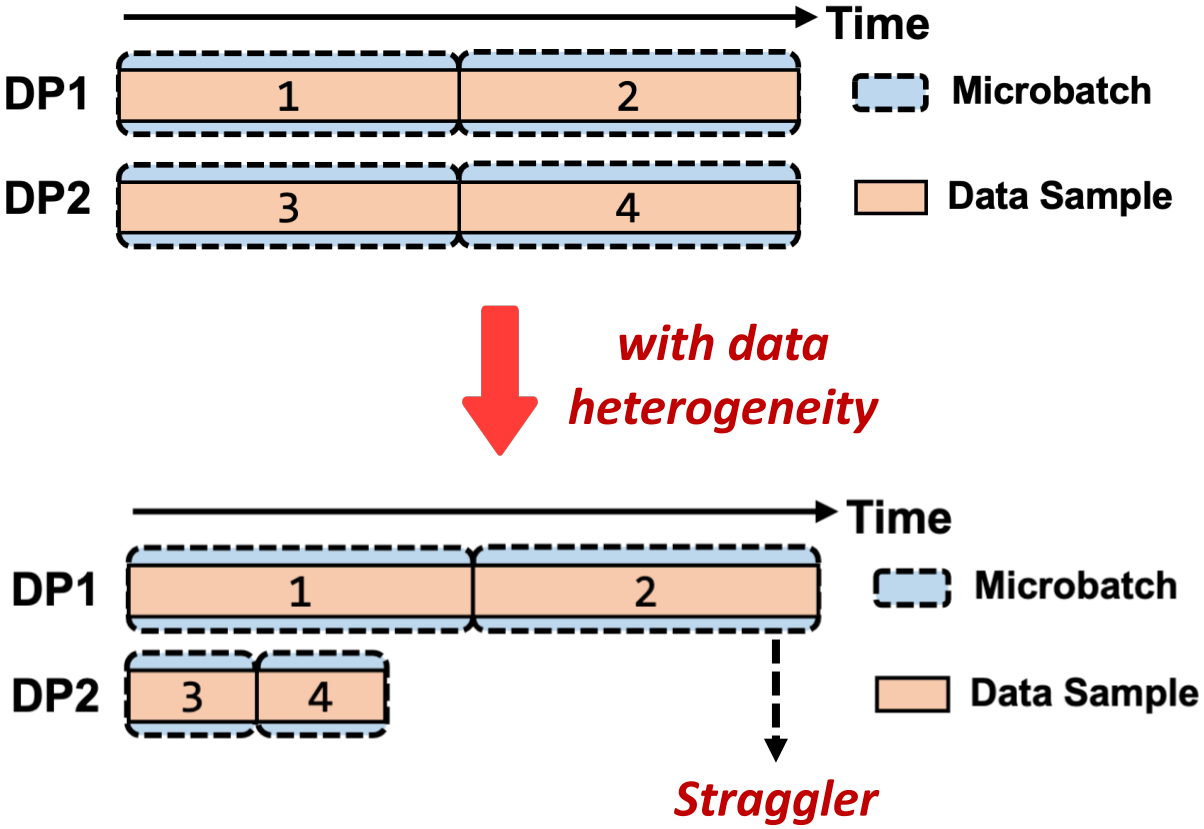

Inter-microbatch straggler

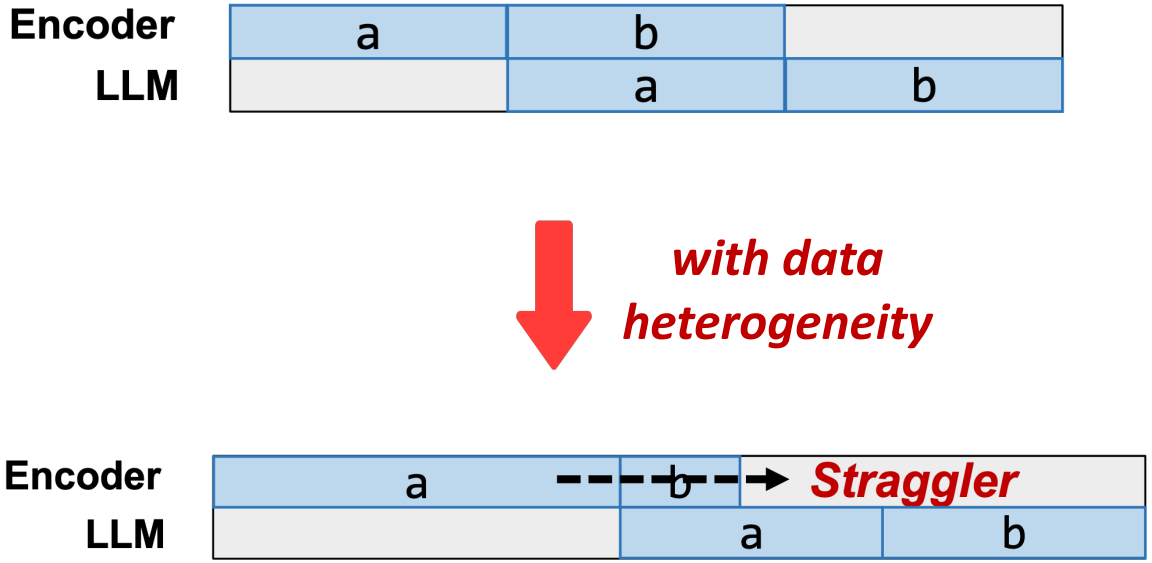

DistTrain design outline

Challenge 1: How to solve the model heterogeneity?

- Disaggregated model orchestration → Assign customized parallelism strategies to different modules
- Optimization formulation → Derive optimal parallelism configurations

Challenge 2: How to solve the data heterogeneity?

- Intra-microbatch reordering → Mitigate stragglers across data-parallel instances
- Inter-microbatch reordering → Mitigate stragglers within pipeline parallelism

Addressing Model Heterogeneity

The monolithic orchestration incurs severe imbalance across modules 😢

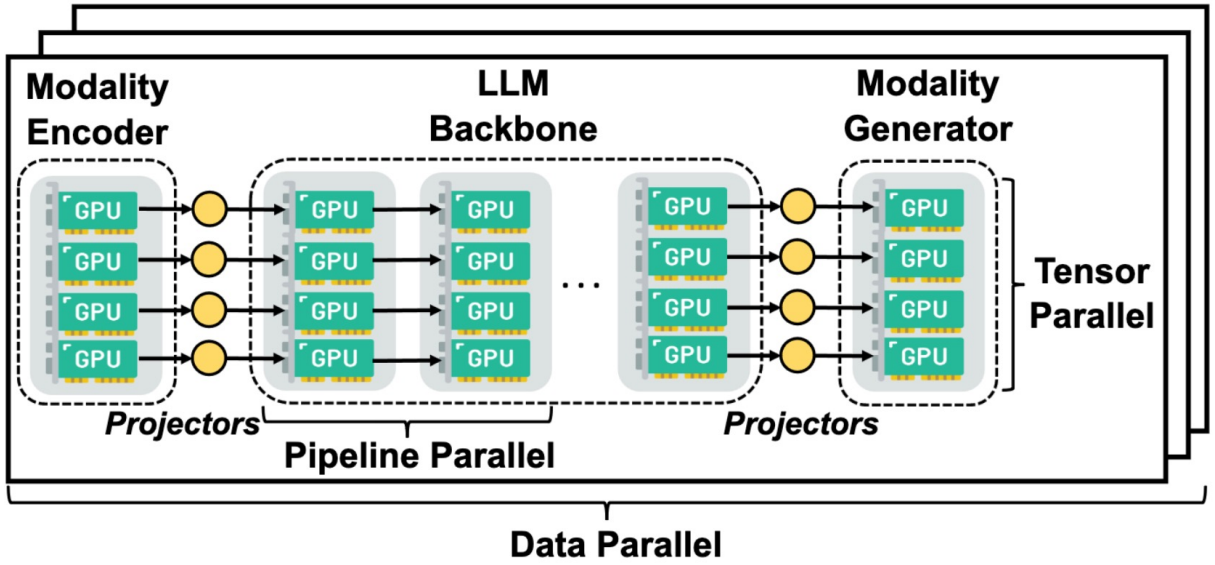

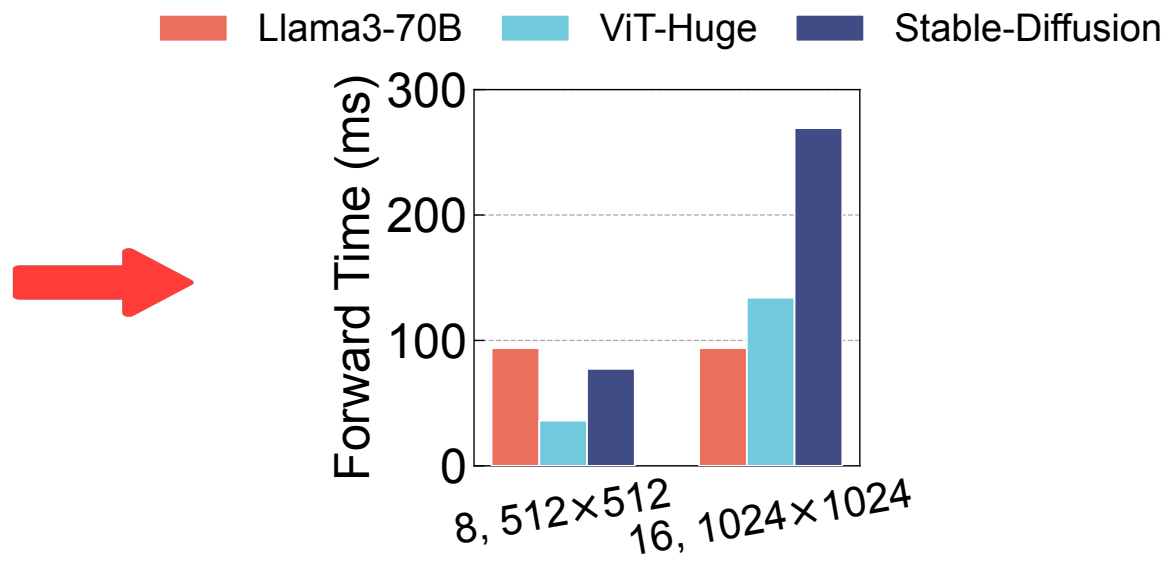

Addressing Model Heterogeneity

Our disaggregated orchestration balances computational loads 😄

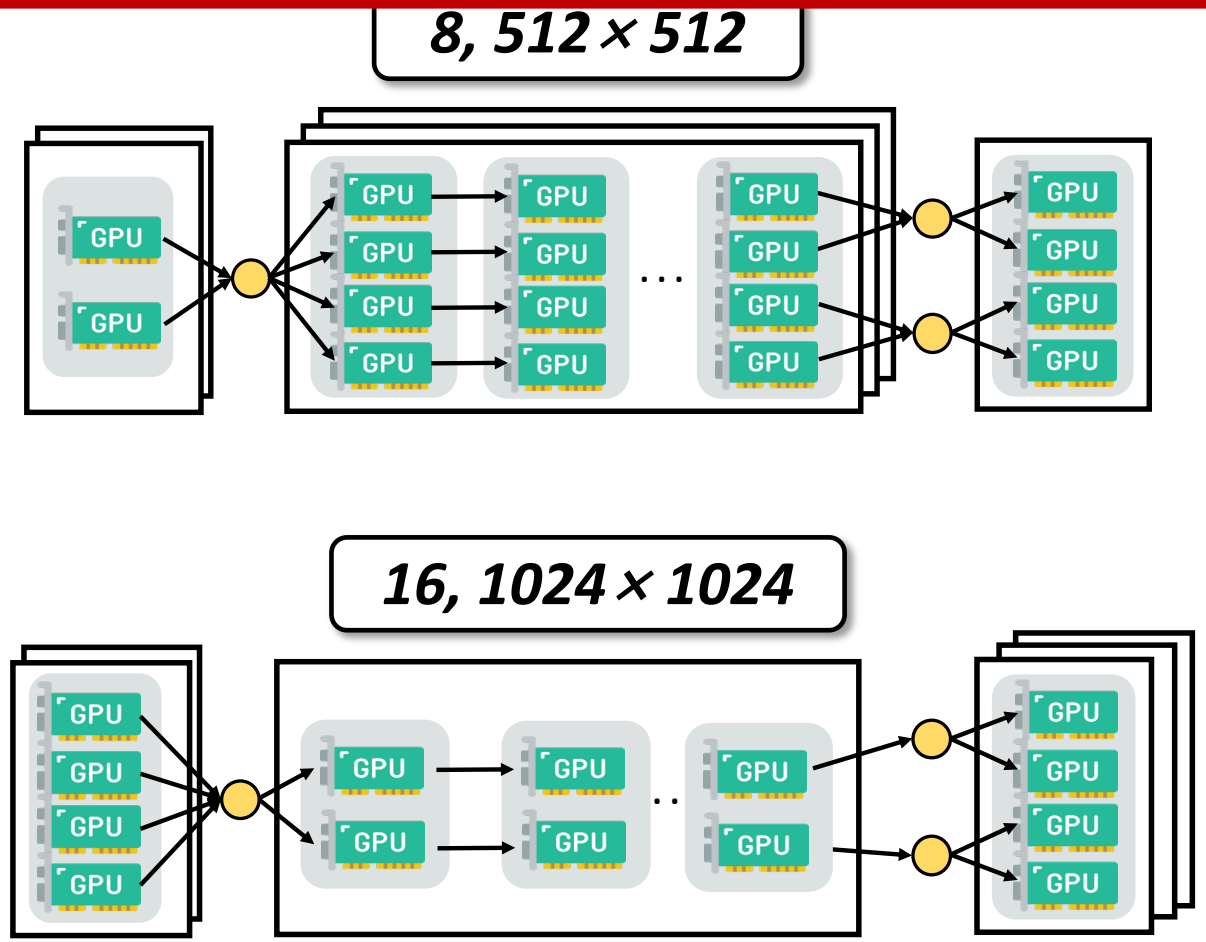

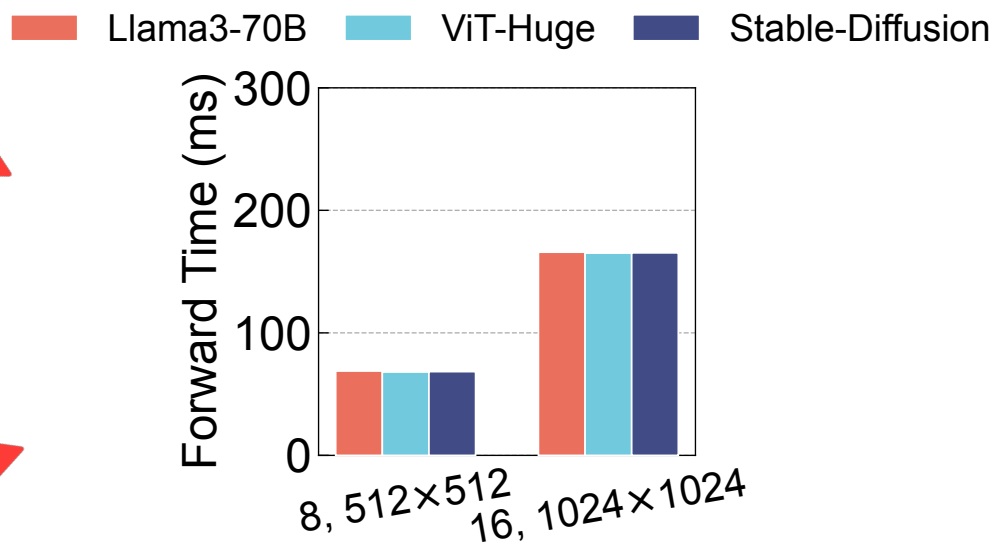

How to decide the detailed disaggregated parallelism strategy?

Addressing Model Heterogeneity

How to decide the detailed disaggregated parallelism strategy?

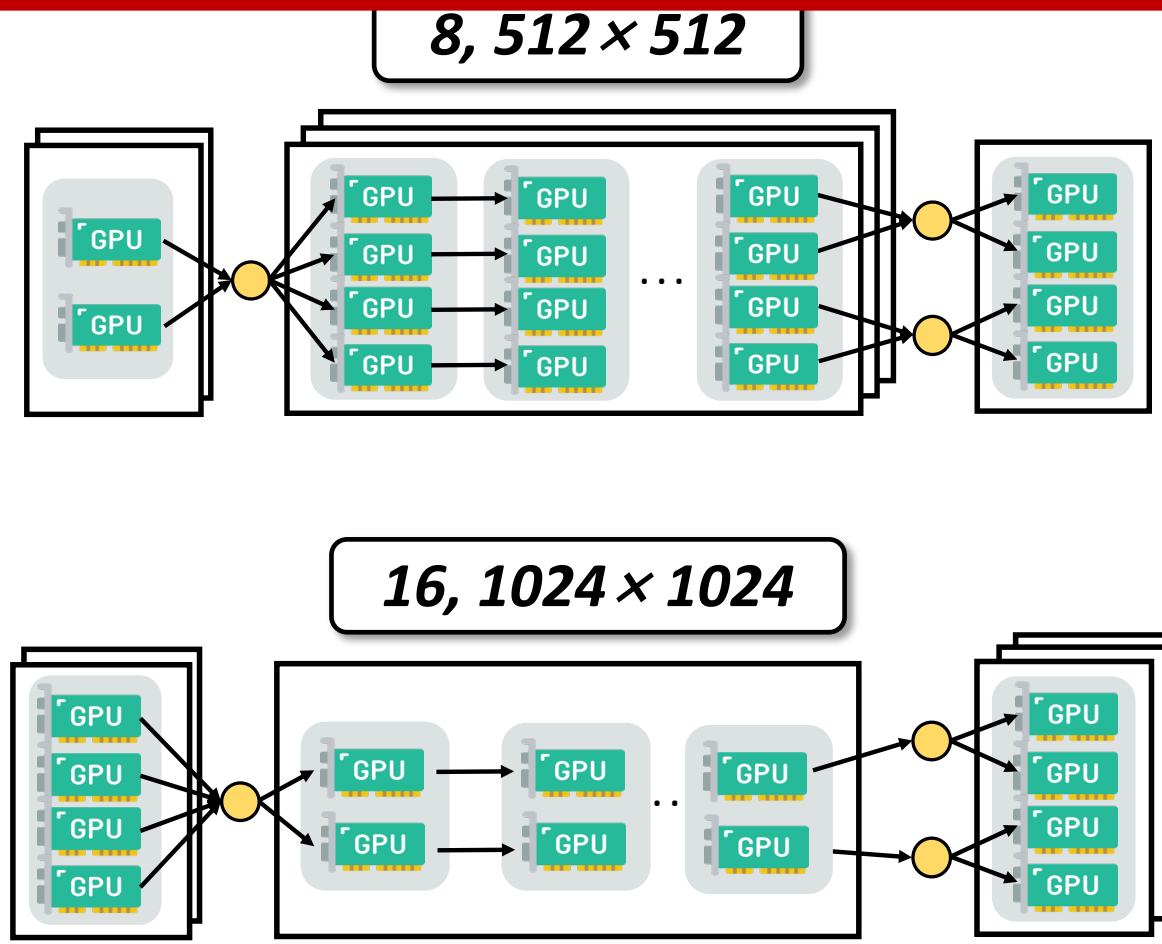

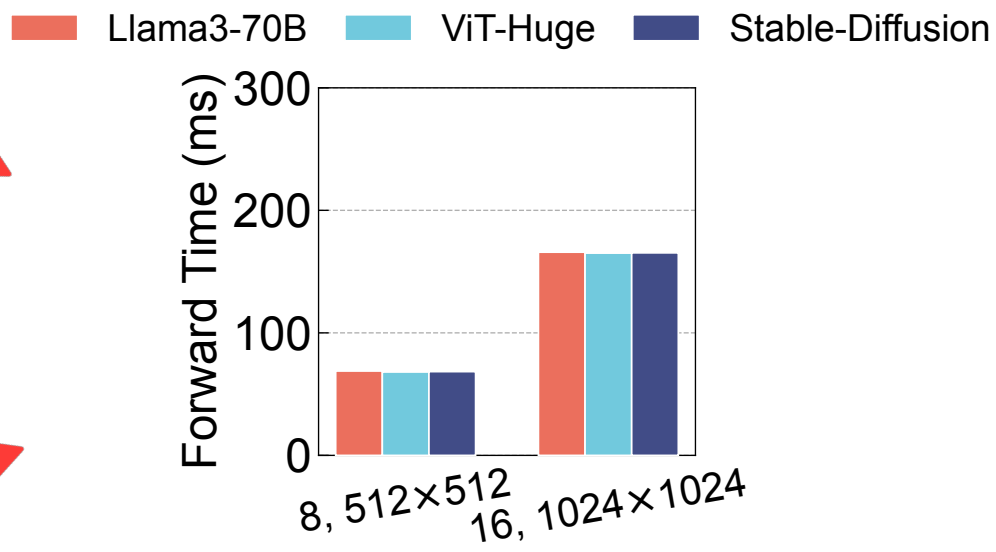

Addressing Model Heterogeneity

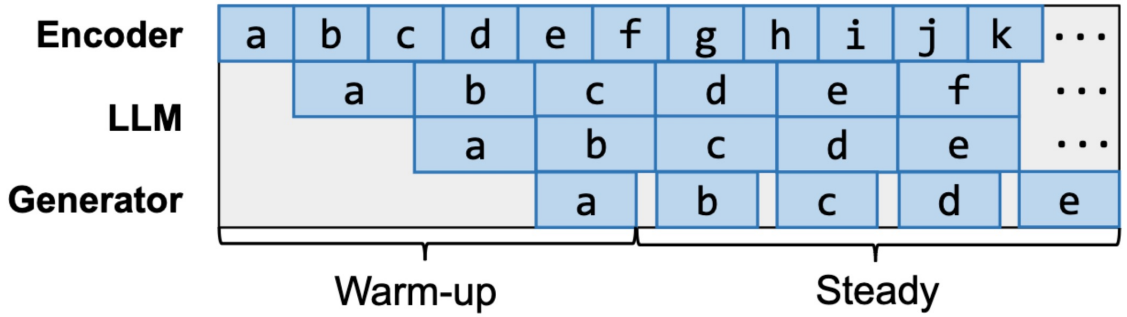

$$T_{warmup} = M \times C_{lm}(TP_{lm}) + \frac{DP_{lm} \times M}{DP_{me}} \times C_{me}(TP_{me}) + \frac{DP_{lm} \times M}{DP_{mg}} \times C_{mg}(TP_{mg})$$

Time for the first microbatch to complete

$$T_{steady} = max \left\{ \begin{array}{l} \frac{DP_{lm} \times TP_{lm} \times M}{y} \times C_{lm}(TP_{lm}), \\ \frac{DP_{lm} \times TP_{me} \times M}{x} \times C_{me}(TP_{me}), \\ \frac{DP_{lm} \times TP_{mg} \times M}{z} \times C_{mg}(TP_{mg}) \end{array} \right\} \times \left(\frac{BS}{DP_{lm} \times M} - 1 \right)$$

Time of the slowest pipeline stage

# of microbatches

Addressing Data Heterogeneity

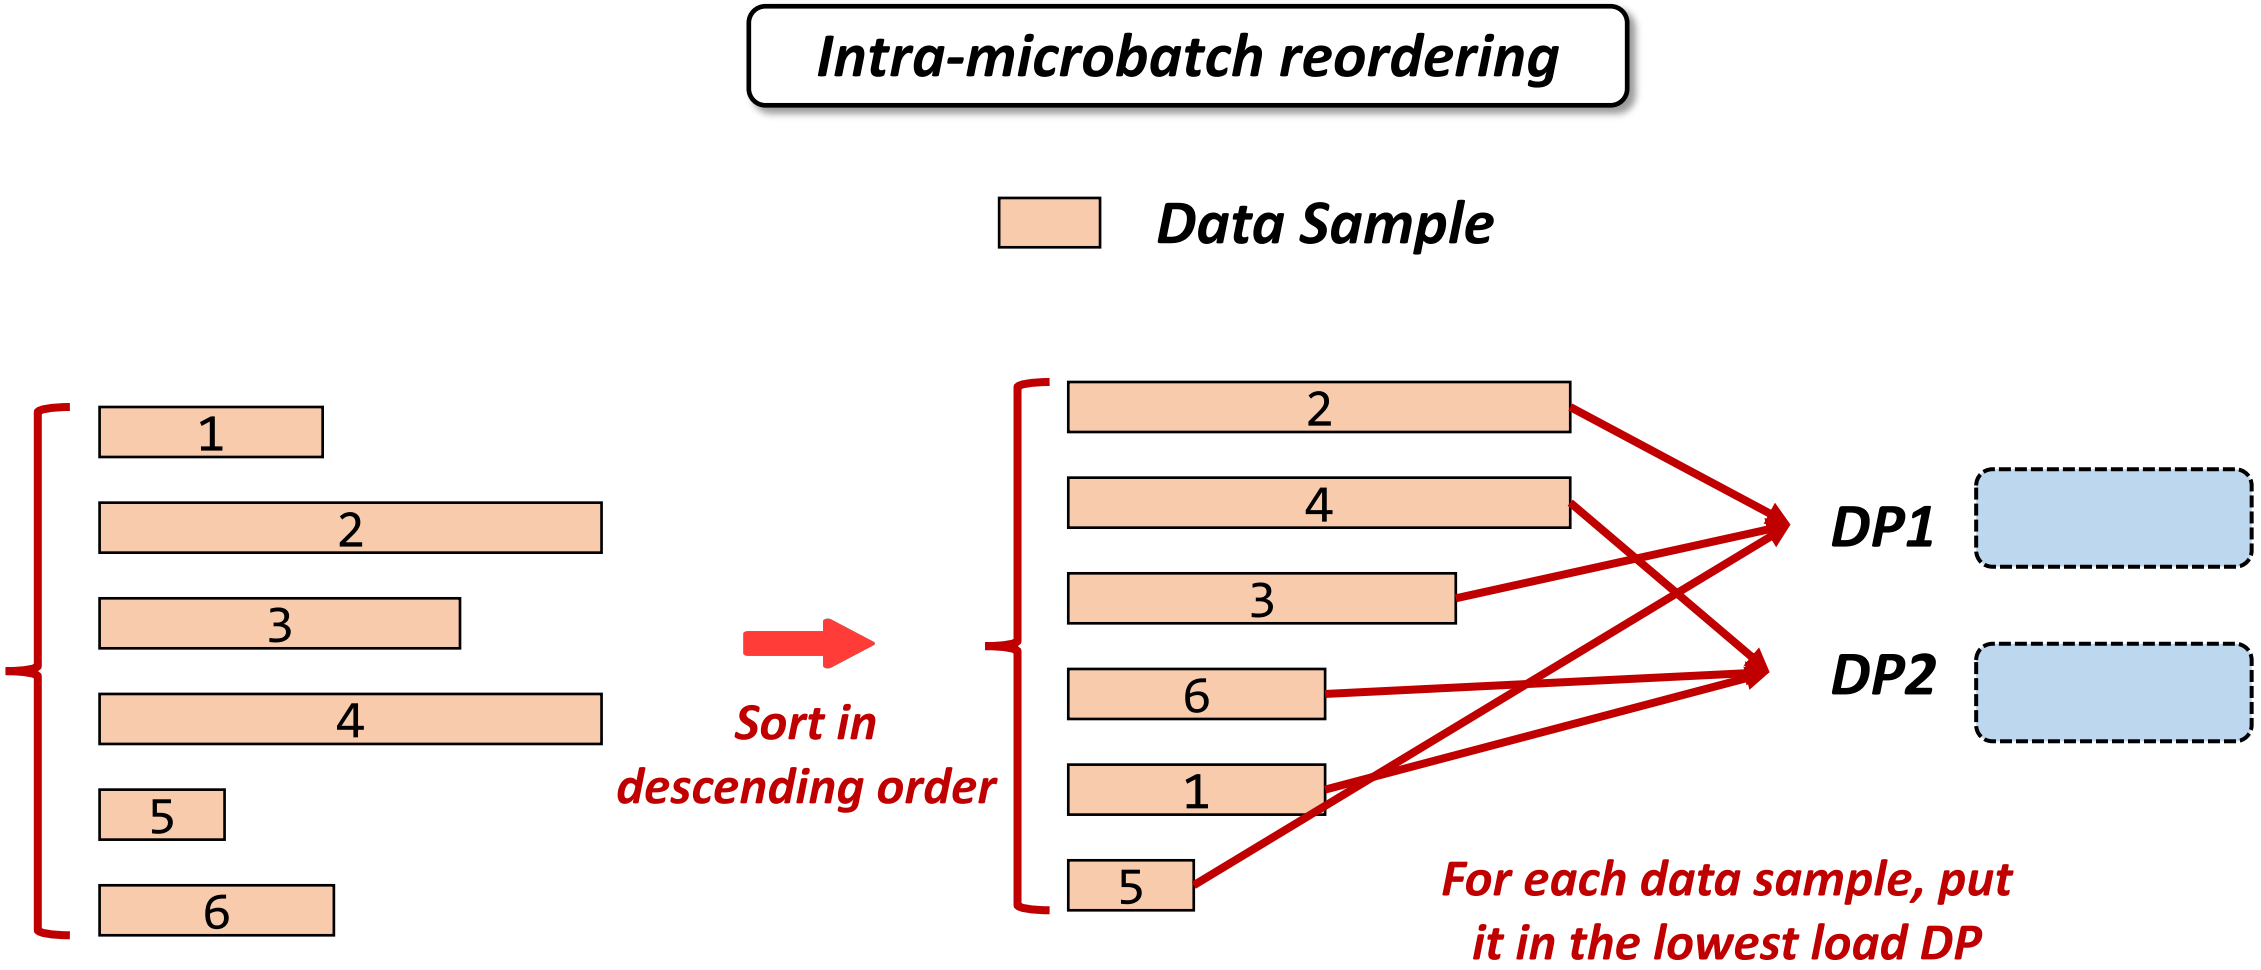

Addressing Data Heterogeneity

Intra-microbatch reordering

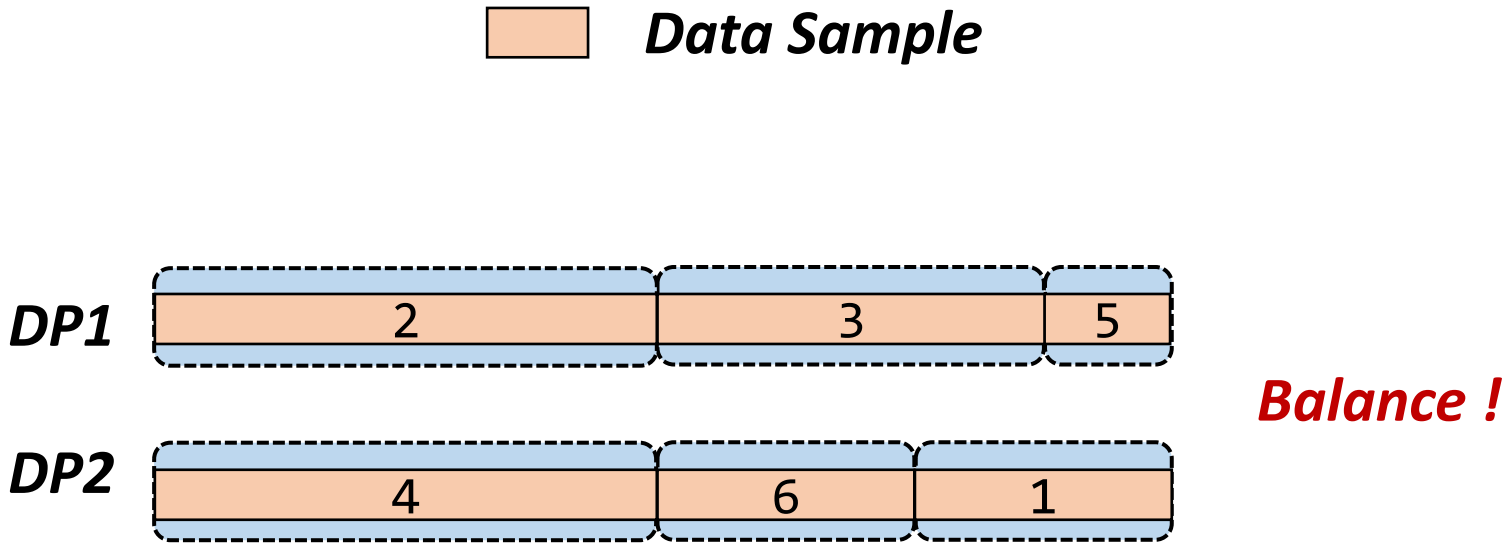

Addressing Data Heterogeneity

Inter-microbatch reordering

Hide the straggler within the large intervals in 1F1B pipeline

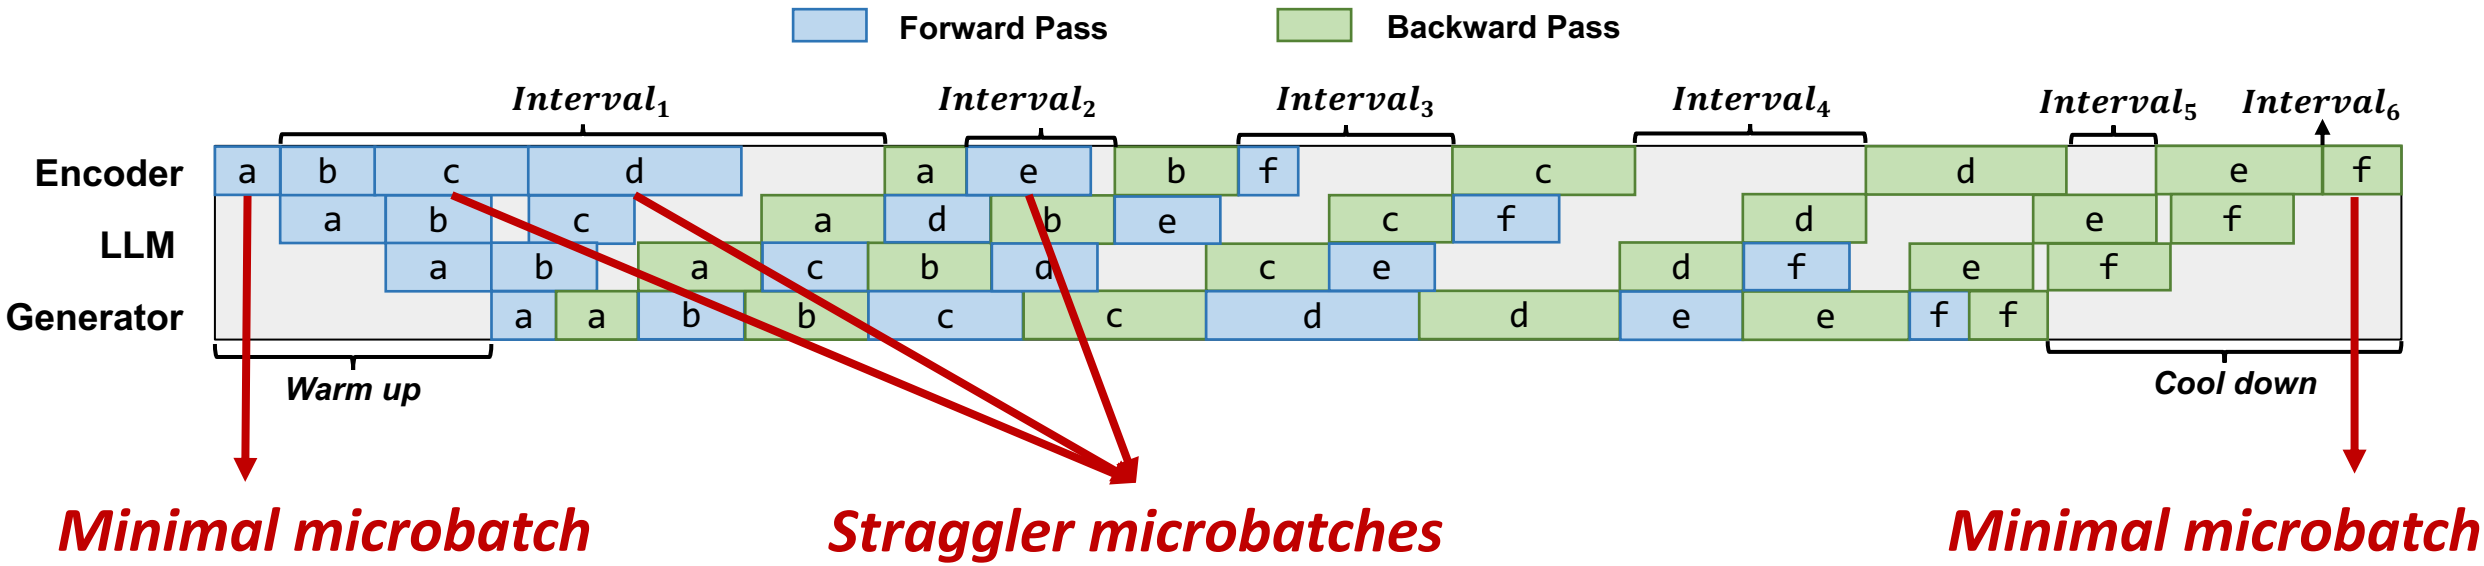

DistTrain System

StepCCL

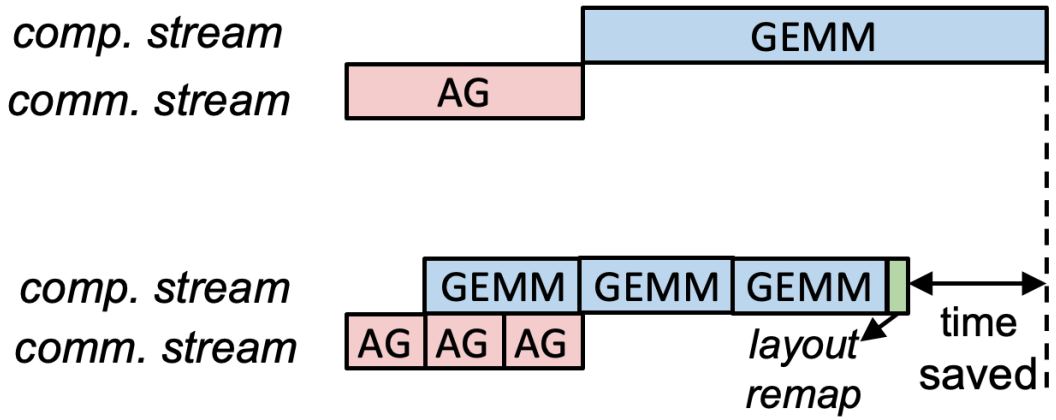

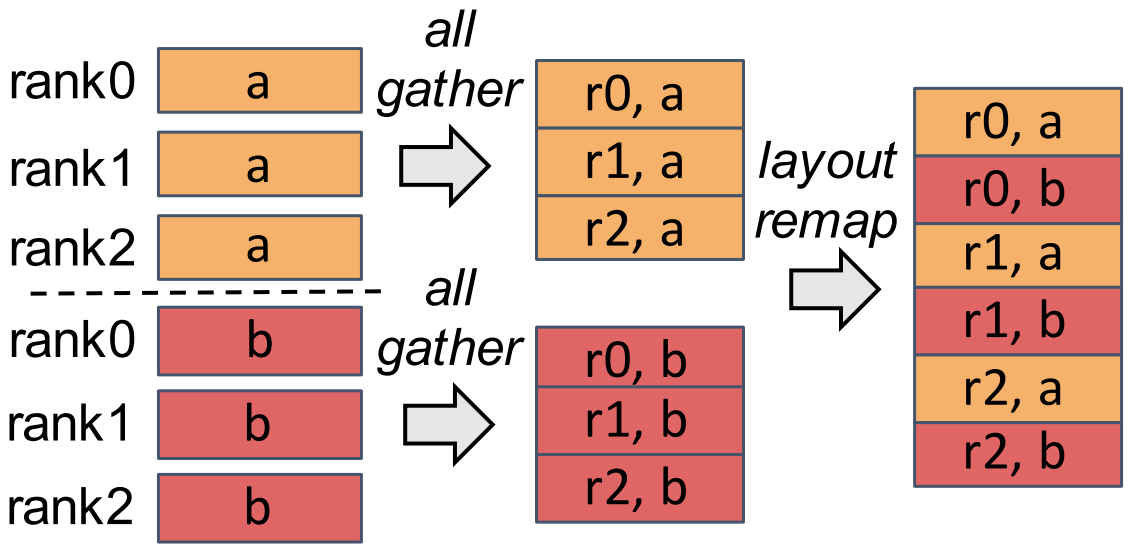

DistTrain System

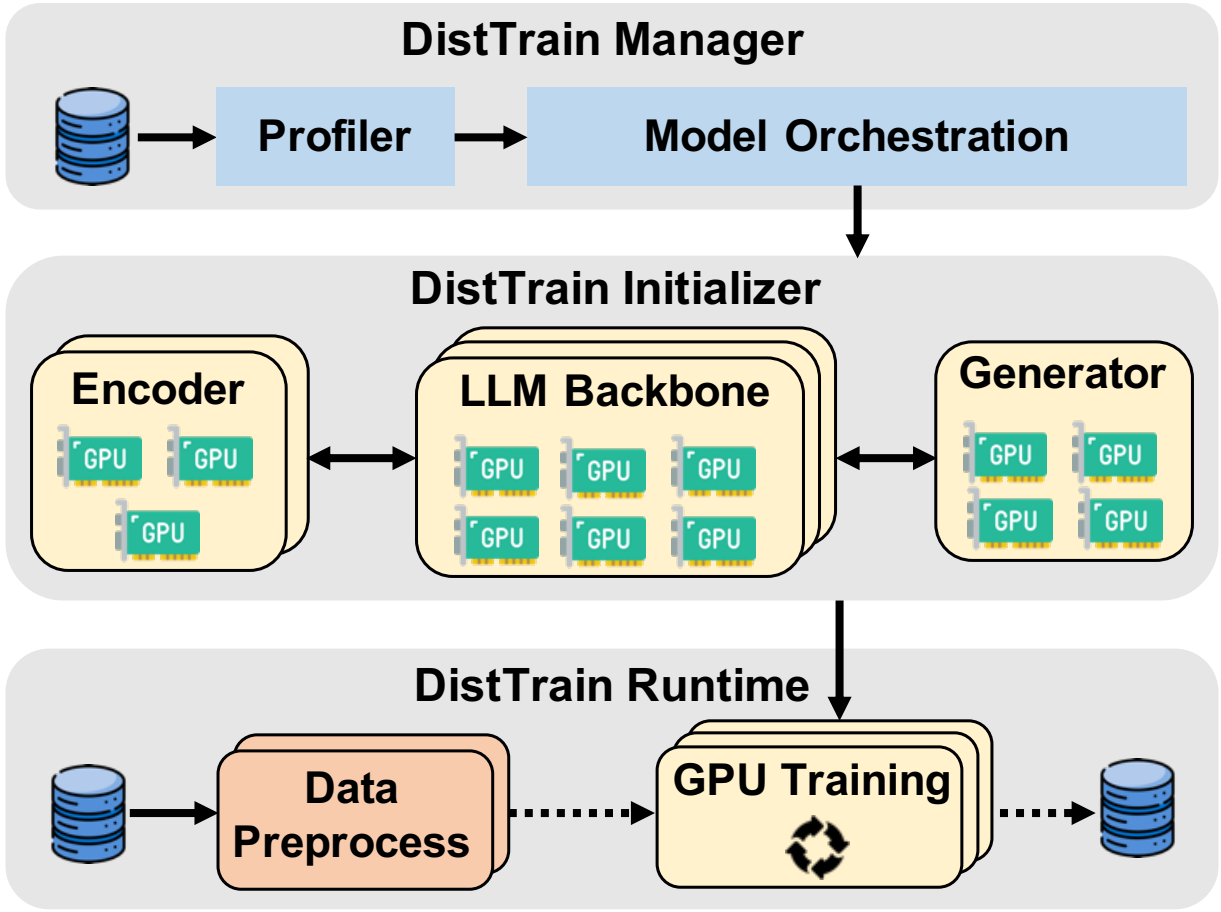

Evaluation

- Setup
  - Hardware: Production cluster with 1,296 A800 GPUs
  - Network: NVLink and RDMA based on RoCEv
  - Dataset: Production training data from LAION
  - Model: MLLM (Llama + ViT + Stable-Diffusion)

Models	# of Layers	Hidden Size	FFN Hidden Size	# of Heads	# of Groups
Llama3-7B	32	4096	11008	32	32
Llama3-13B	40	5120	13824	40	40
Llama3-70B	80	8192	28672	64	8

Evaluation

- DistTrain outperforms Megatron-LM on large-scale MLLM training

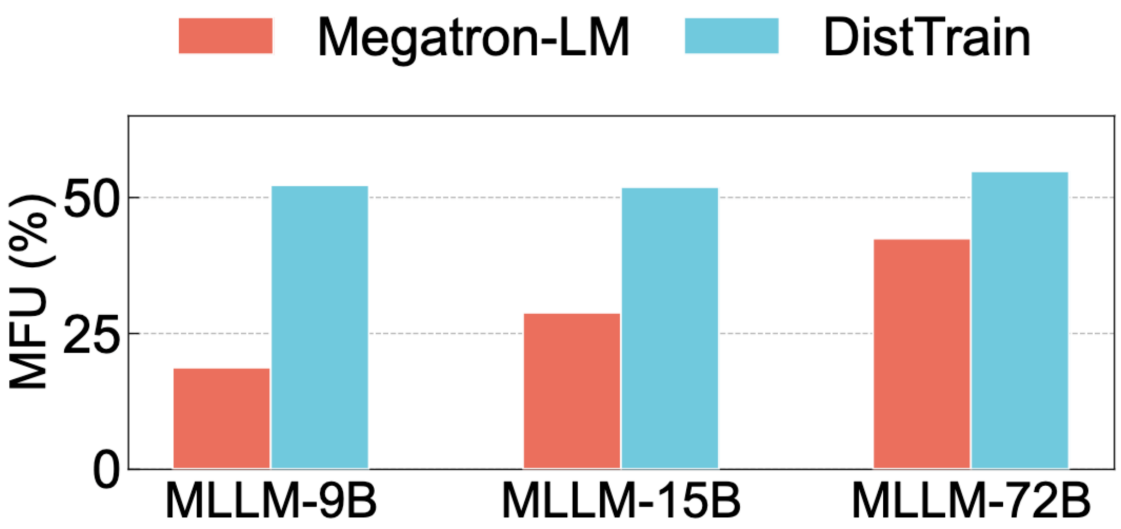

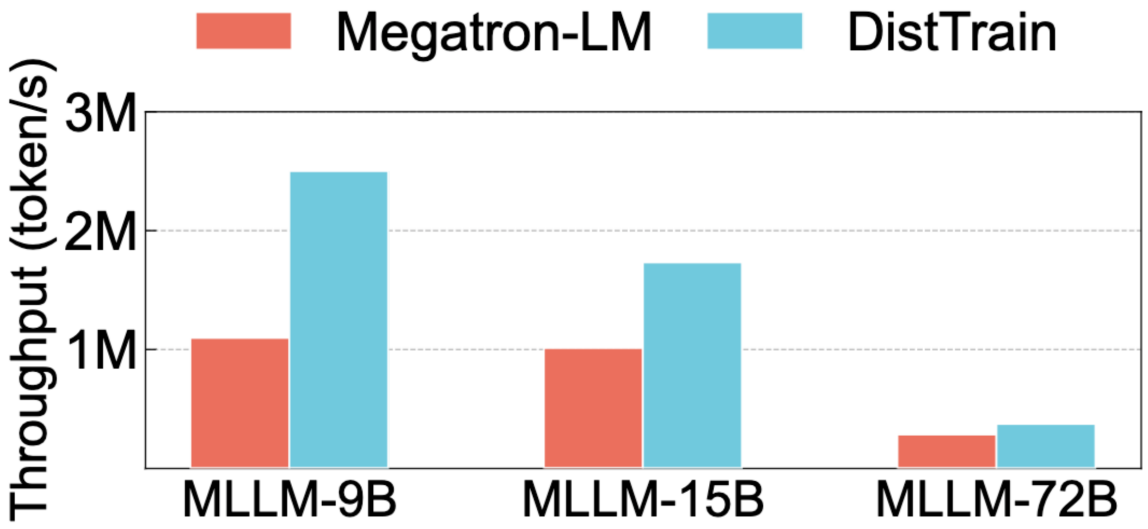

Evaluation

- DistTrain's disaggregated model orchestration is effective

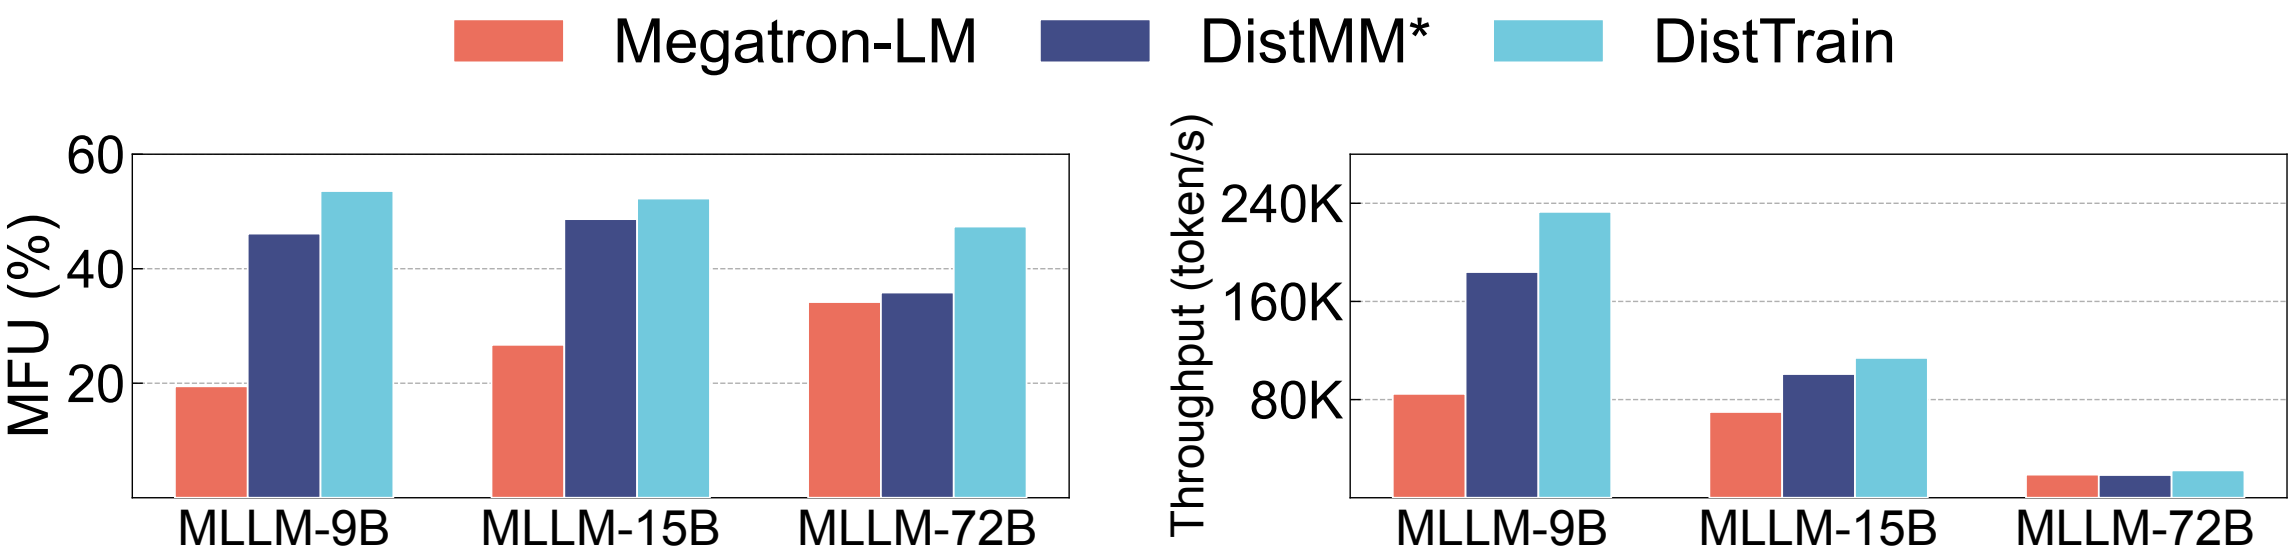

Evaluation

- DistTrain's disaggregated data processing (reordering) is effective

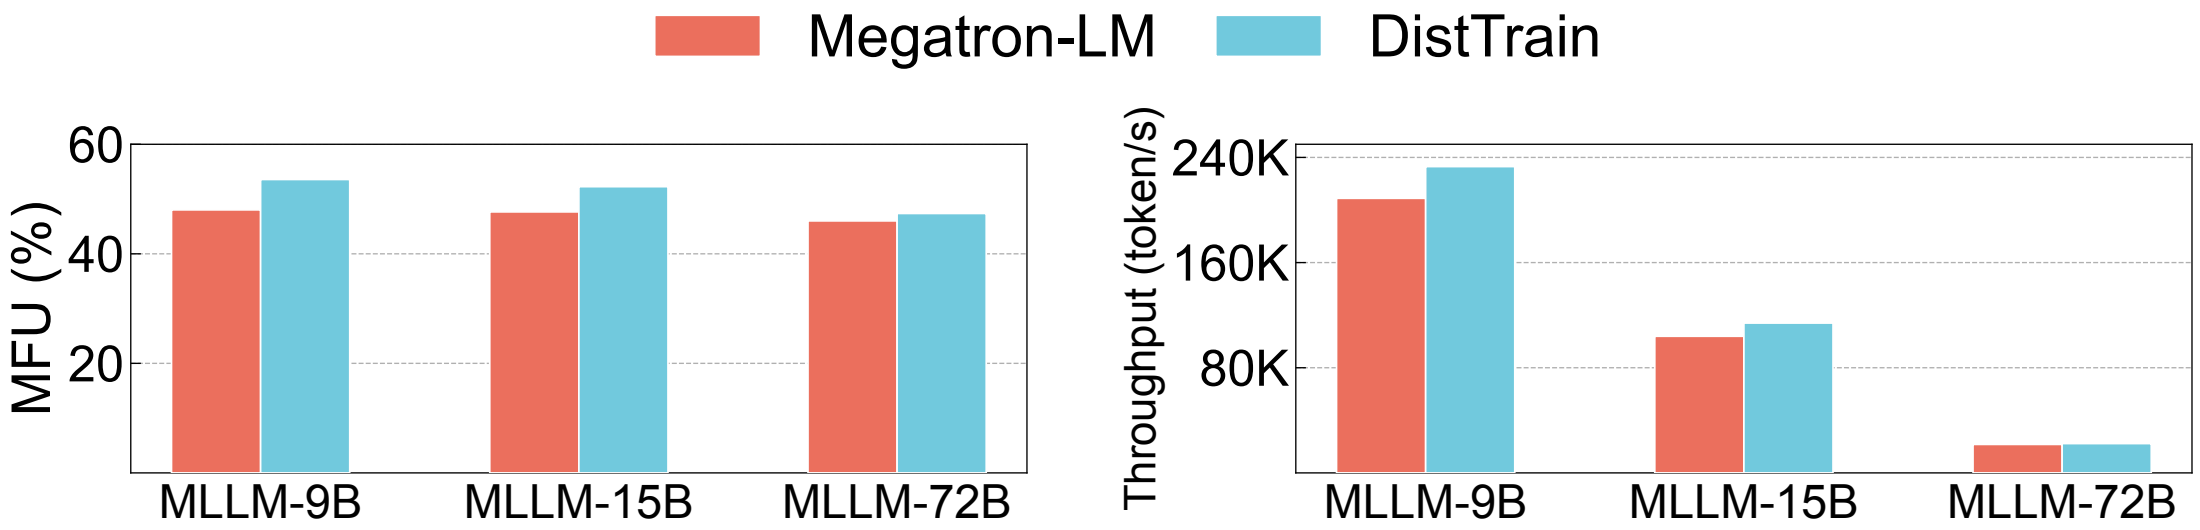

Evaluation

- Performance under module freezing
- Performance of StepCCL
- System overhead of DistTrain

Conclusion

- Multimodal LLM training suffers from **model** and **data heterogeneity**

- DistTrain introduces two techniques:
  - *Disaggregated model orchestration* → tackles model heterogeneity
  - *Disaggregated data preprocessing* → tackles data heterogeneity

- DistTrain outperforms Megatron-LM by up to 2.2× on training throughput and achieves 54.7% MFU

- DistTrain supports large-scale open-sourced **MLLM** (130B) training

Thank you!

zzlcs@pku.edu.cn