Parcae: Proactive, Liveput-Optimized DNN Training on Preemptible Instances

Jiangfei Duan*, **Ziang Song***, Xupeng Miao*, Xiaoli Xi,
Dahua Lin, Harry Xu, Minjia Zhang, Zhihao Jia

Carnegie Mellon University
香港中文大學 The Chinese University of Hong Kong
ByteDance
UCLA
Microsoft

LLM Training Is Expensive

$3.9 Million

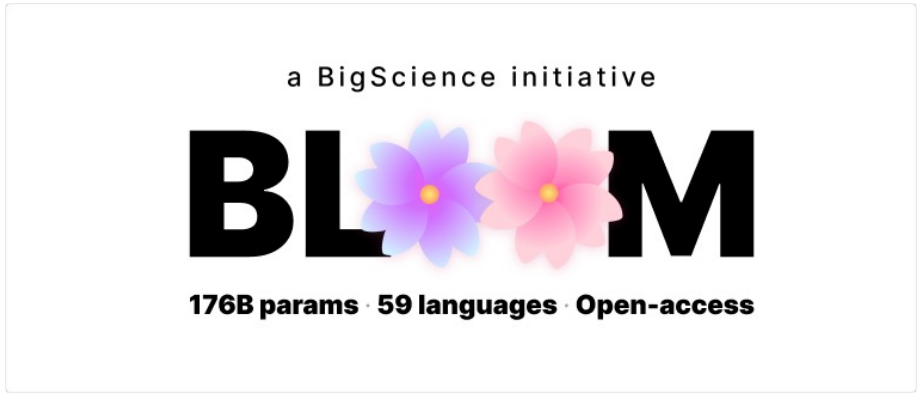

$2-5 Million

GPT-3

$4.6 Million

Spot Instances Reduce Training Cost

up to 10× Savings

up to 9× Savings

up to 3× Savings

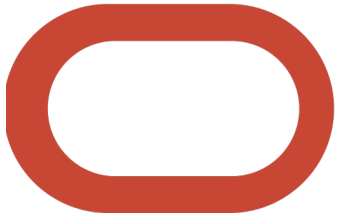

2× Savings

Spot Instances Reduce Training Cost, But Are Preemptible

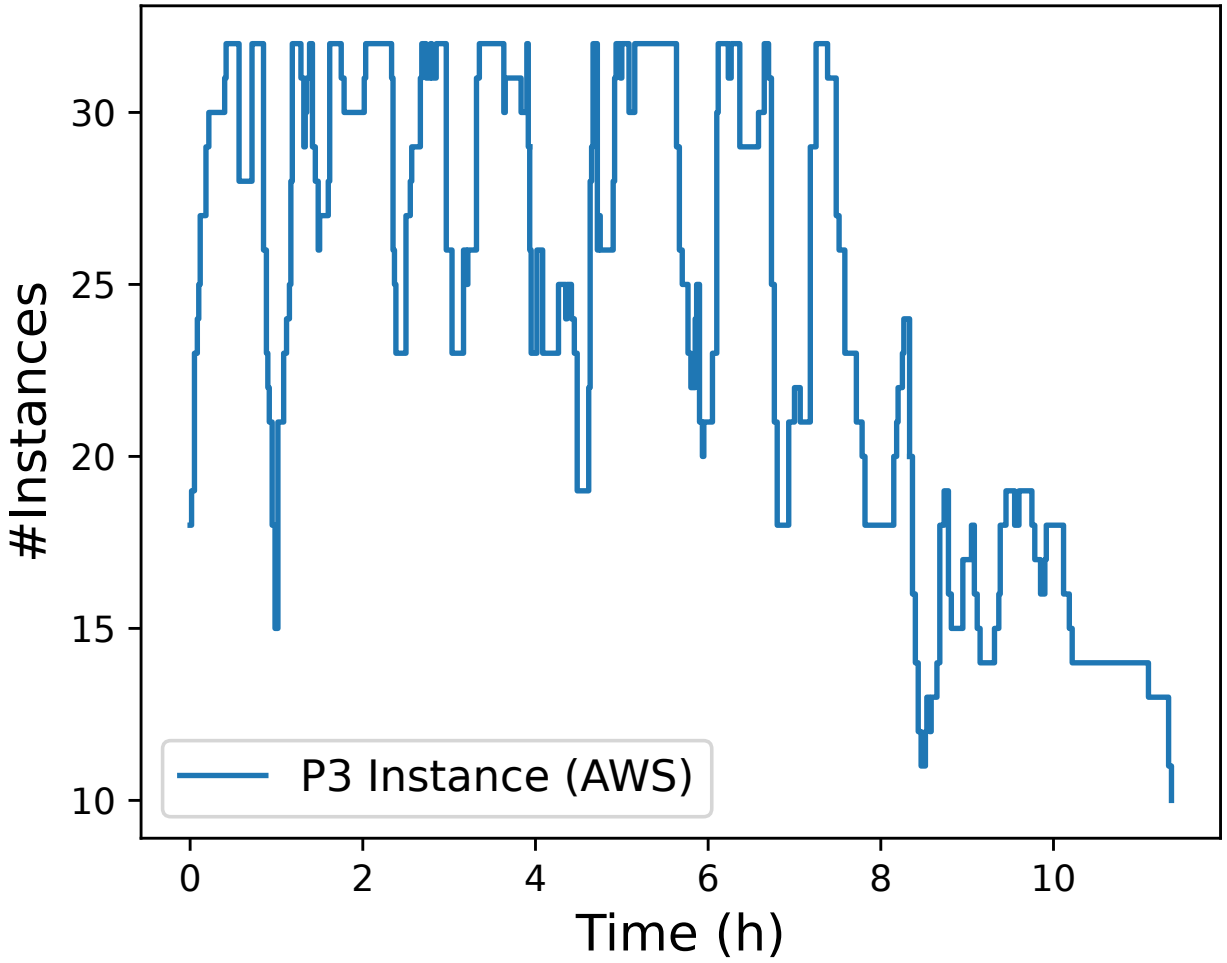

Preemption once every
5 minutes,
worst case 1 minute

Background: Distributed DNN Training

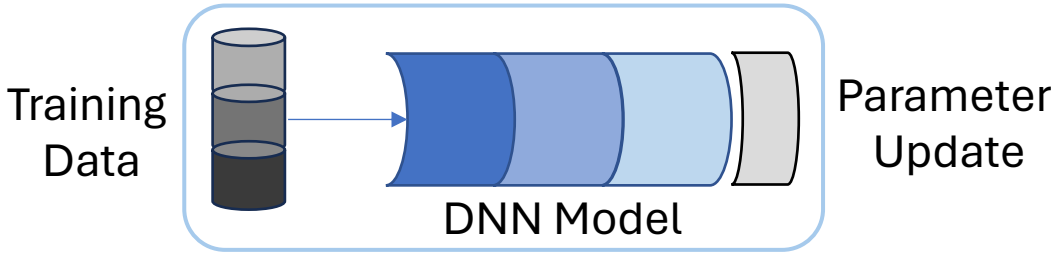

Single GPU Training

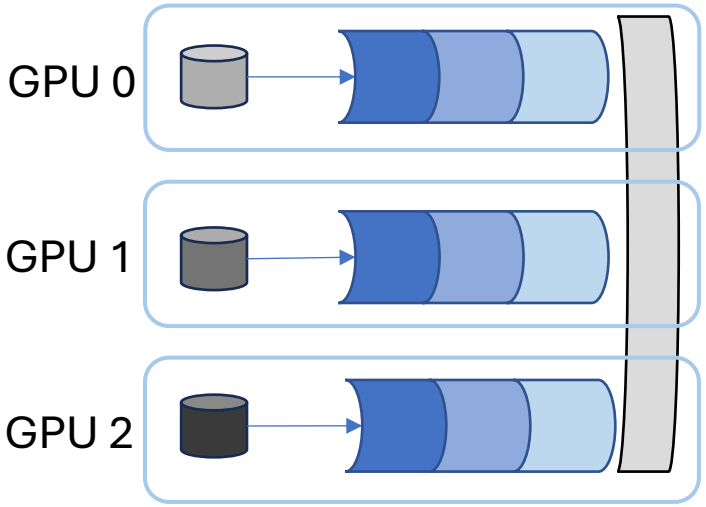

Data Parallelism

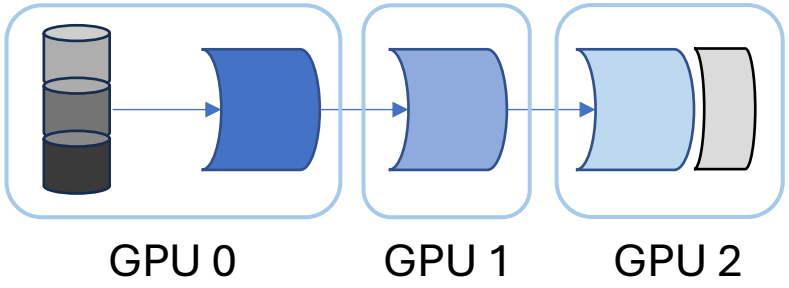

Pipeline Parallelism

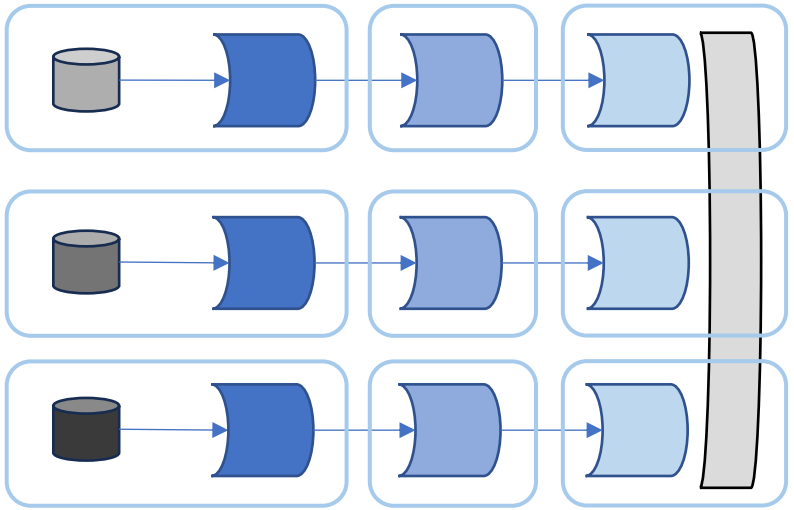

Hybrid Parallelism

Checkpoint-based: Varuna (EuroSys '22)

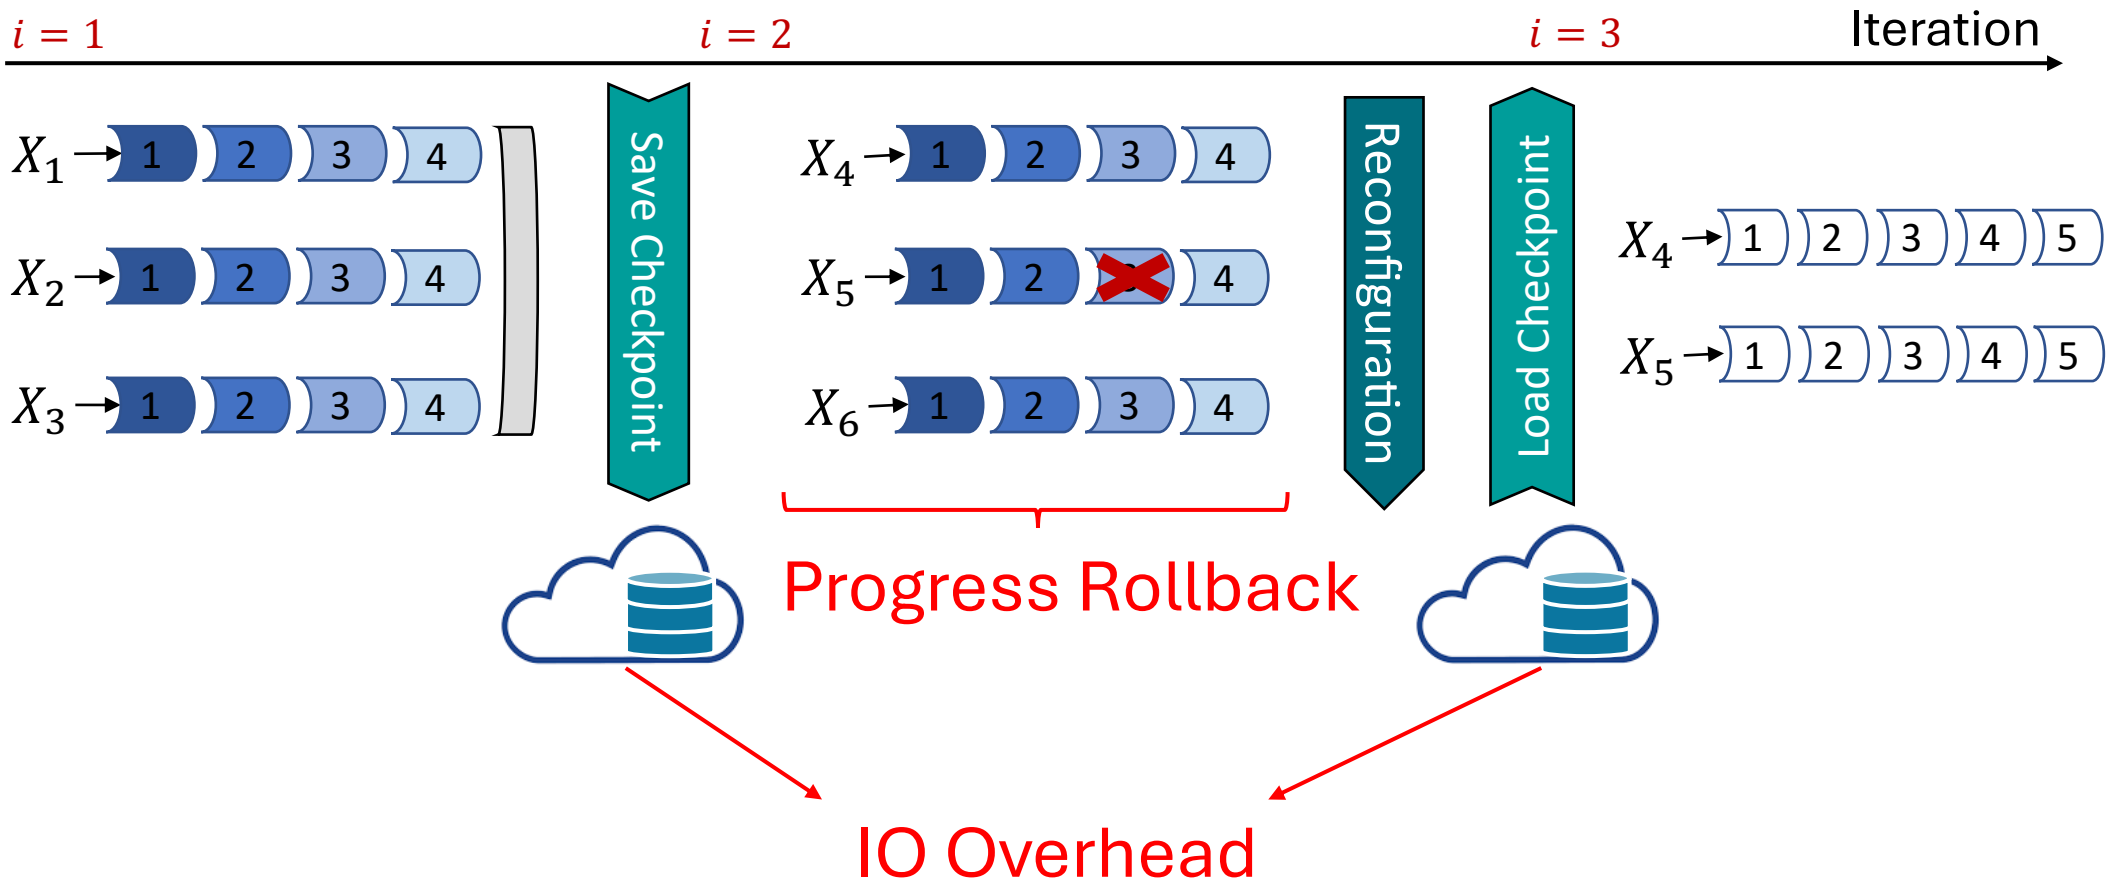

Redundancy-based: Bamboo (NSDI '23)

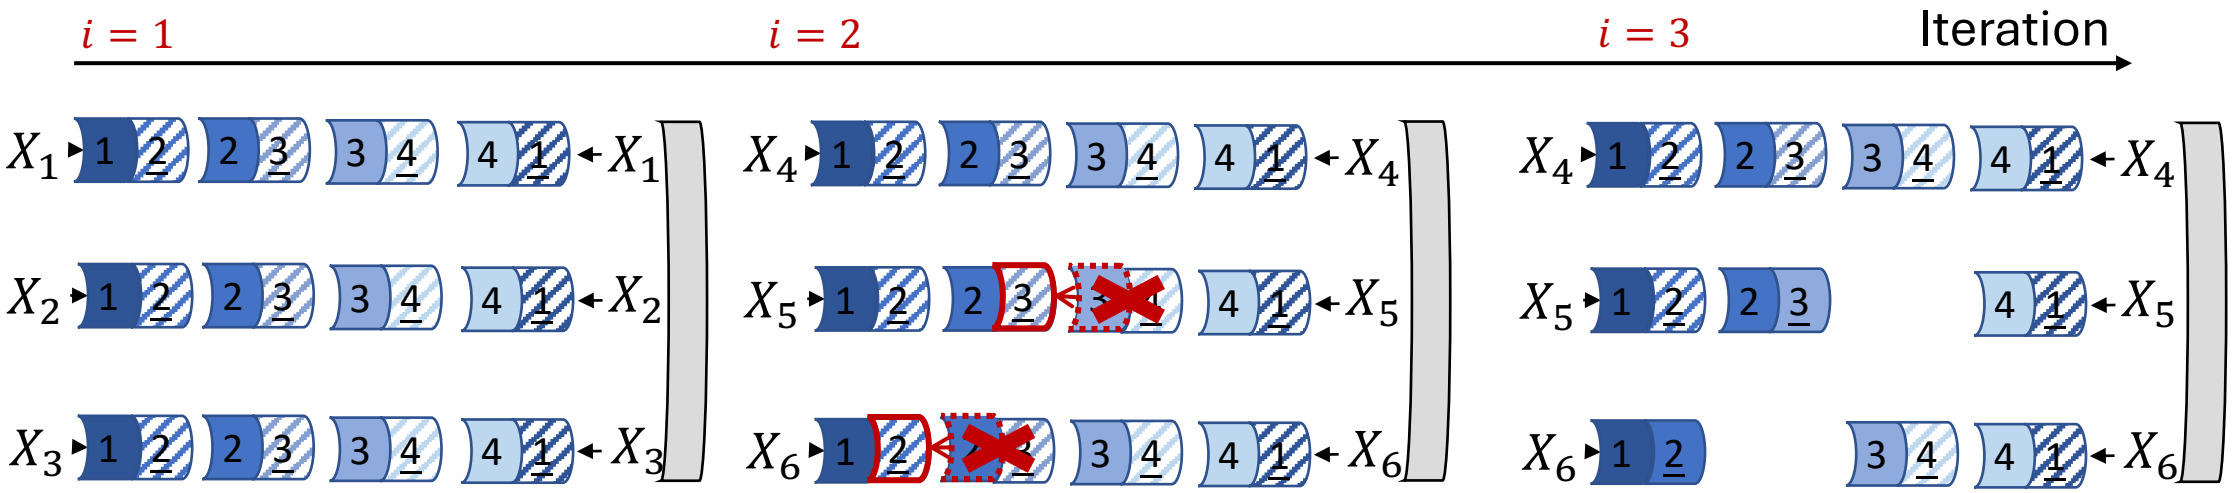

Redundant Computation and Memory Overhead

Parcae: Proactive, Liveput-Optimized

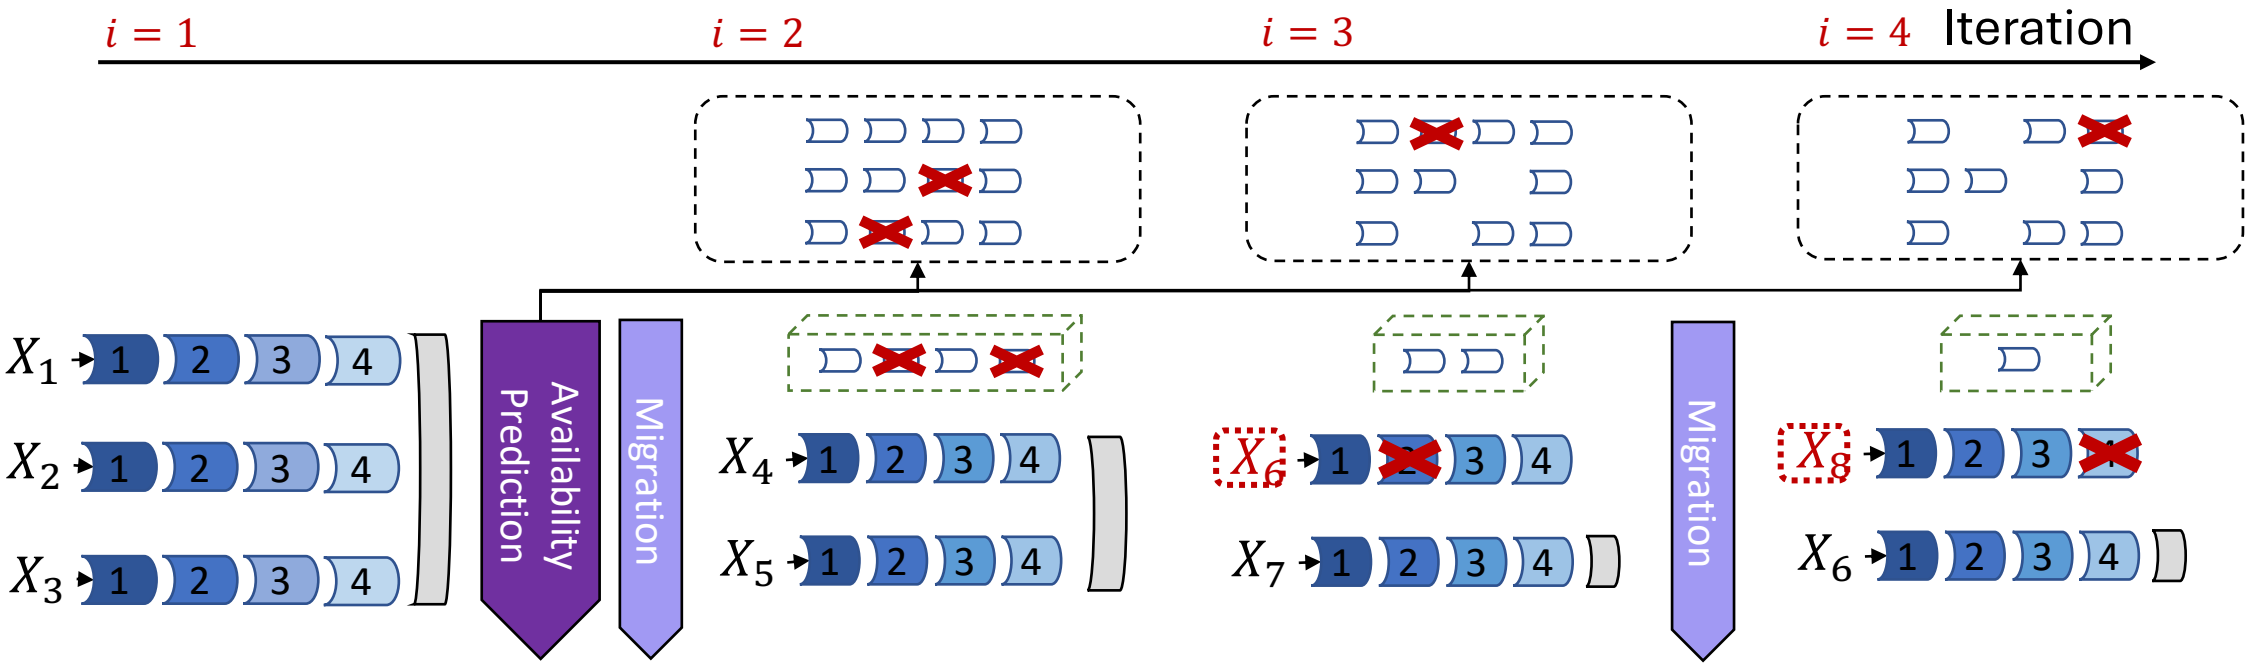

Proactively plan for future preemptions to maximize *preemption-aware* throughput (i.e., Liveput)

Up To **10**× Improvement, Near Optimal

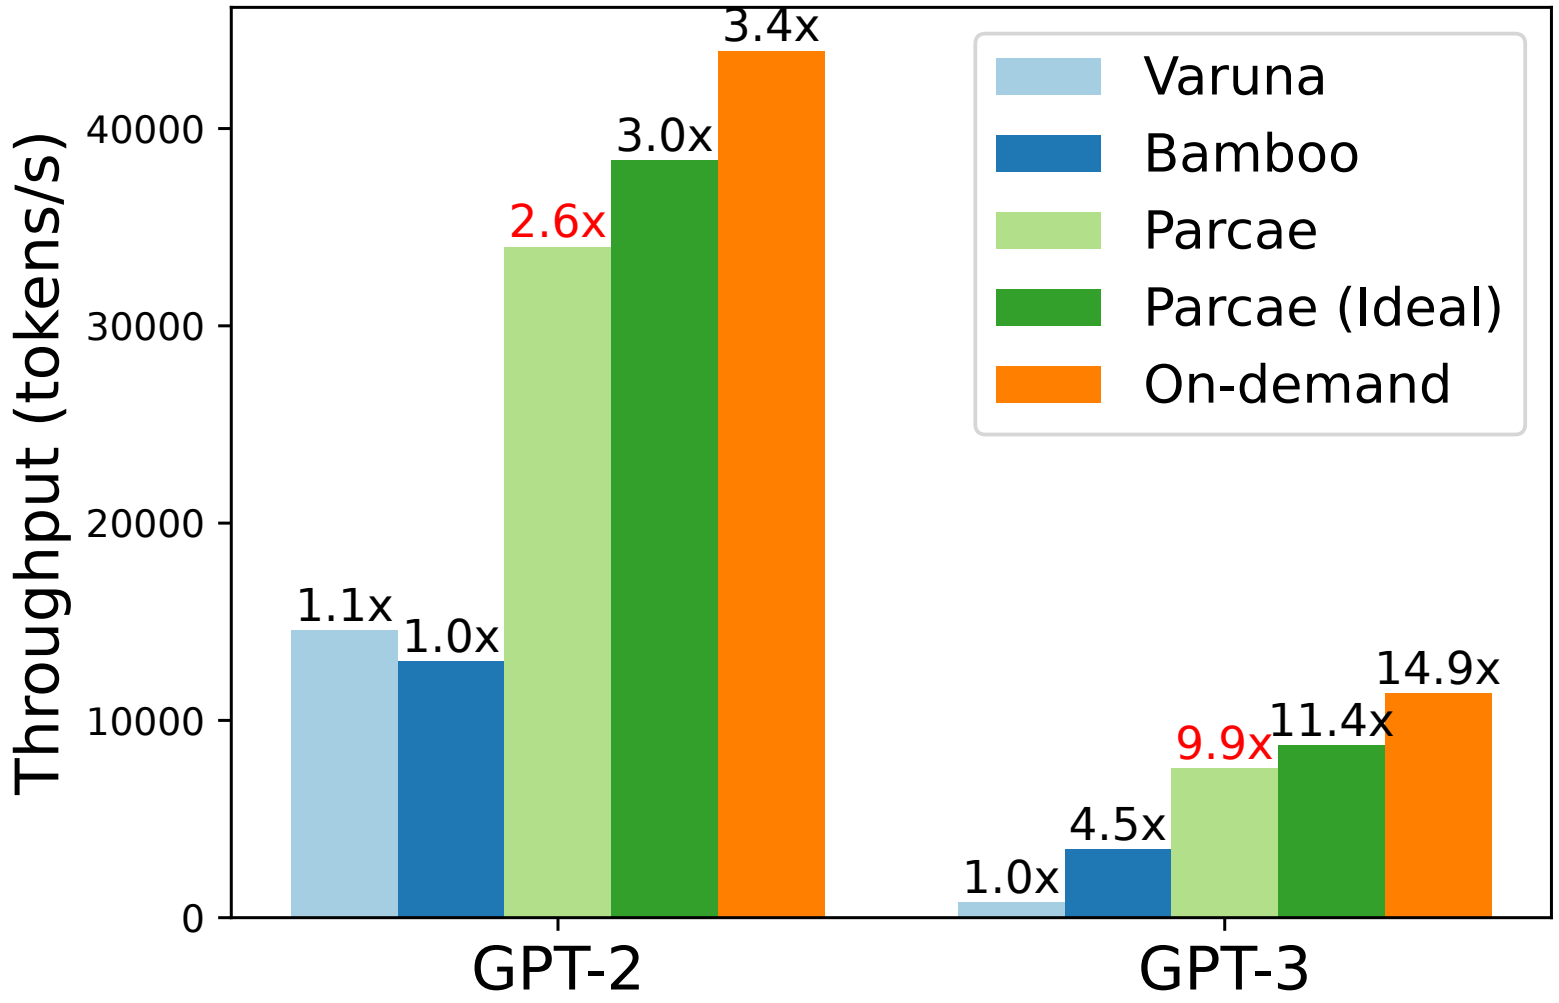

Liveput

$$\text{LIVEPUT}(D, P, \mathcal{V}) = \mathbb{E}_{\underbrace{\vec{v} \sim \mathcal{V}}}[\text{THROUGHPUT}(\underbrace{D_{\vec{v}}, P_{\vec{v}}})]$$

Preemption Scenario

Parallelization Strategy after Preemption

Liveput Example

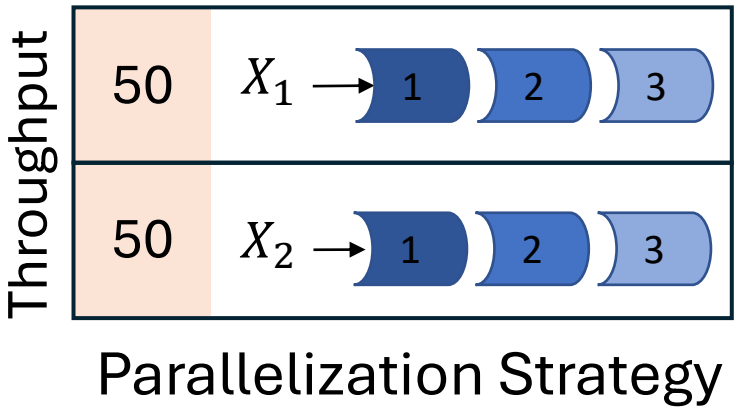

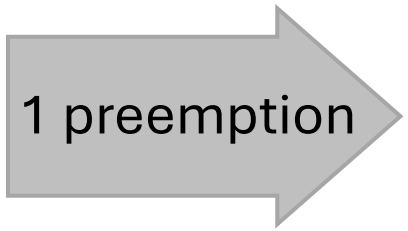

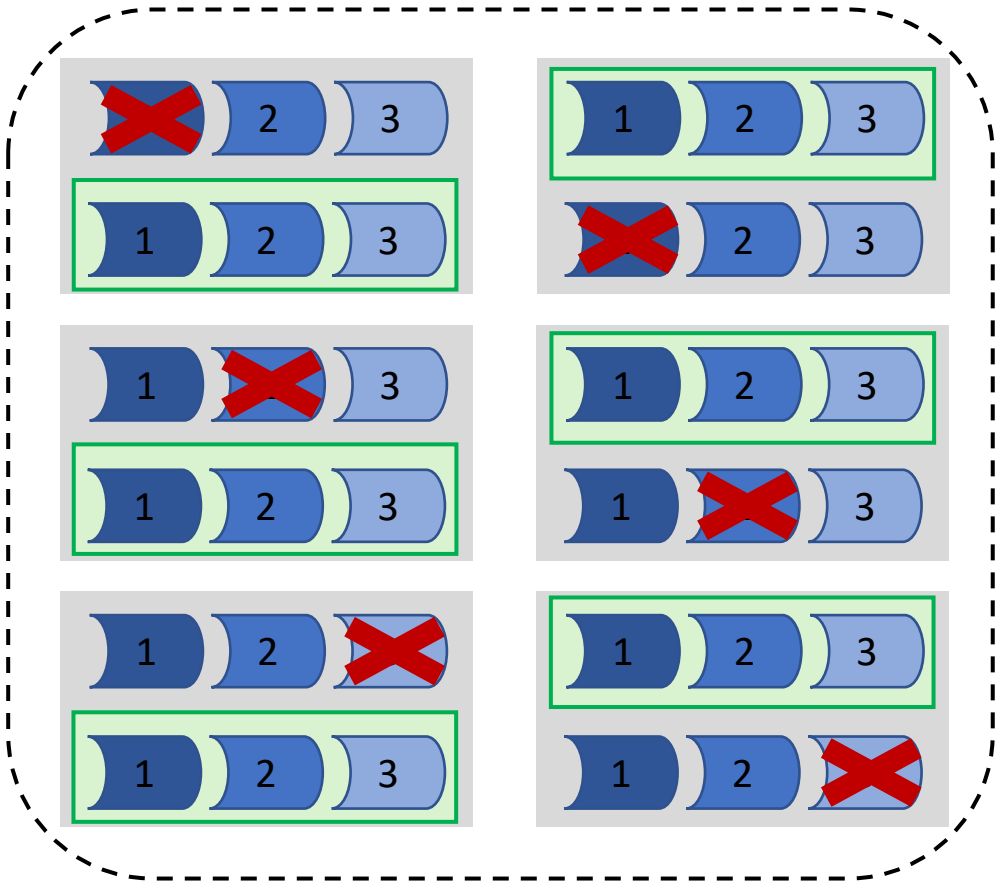

LIVEPUT (1 preemption) = $100\% \times 50 = 50$

100% remain 1 pipeline

Maximizing Throughput Is Sub-Optimal

Configurations	#preemptions	Distribution After Preemption	Throughput	LIVEPUT
50 X_1 → 1 2 3	0	100% → 1 2 3 / 1 2 3	100 👍	$100\% \times 100 = 100$ 👍
50 X_2 → 1 2 3	1	100% → 1 2 3		$100\% \times 50 = 50$
30 X_1 → 1 2	0	100% → 1 2 / 1 2 / 1 2	90	$100\% \times 90 = 90$
30 X_2 → 1 2				
30 X_3 → 1 2	1	100% → 1 2 / 1 2		$100\% \times 60 = 60$ 👍

Challenges

1. Predict Liveput
2. Maximize Liveput
3. Handle Preemption

Our Contributions

1. Availability Predictor
2. Liveput Optimizer
3. Live Migration

Instance-wise Preemption Is Unpredictable

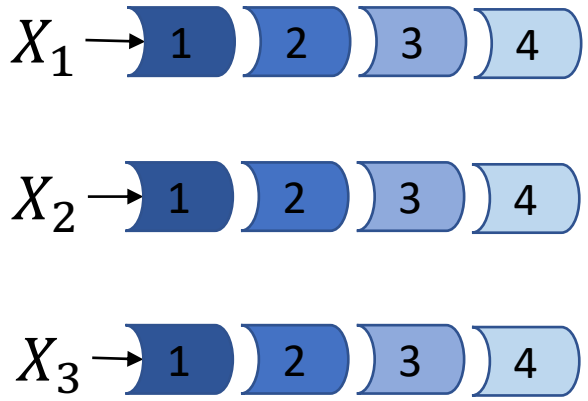

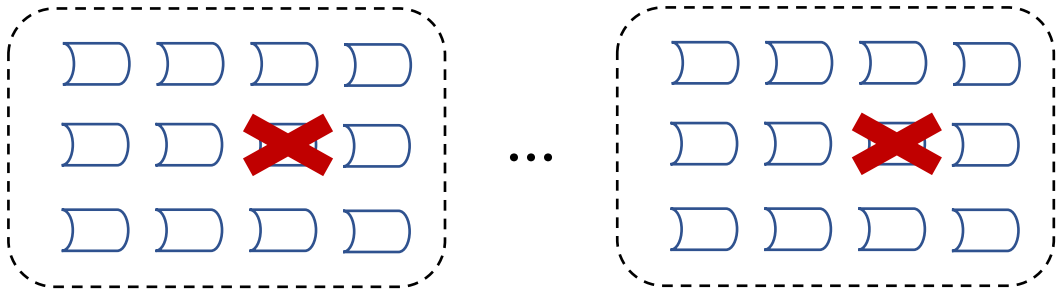

Parallelization Strategy

Preemption Mapping ☹

LIVEPUT

Short-Term Instance Availability Is Predictable

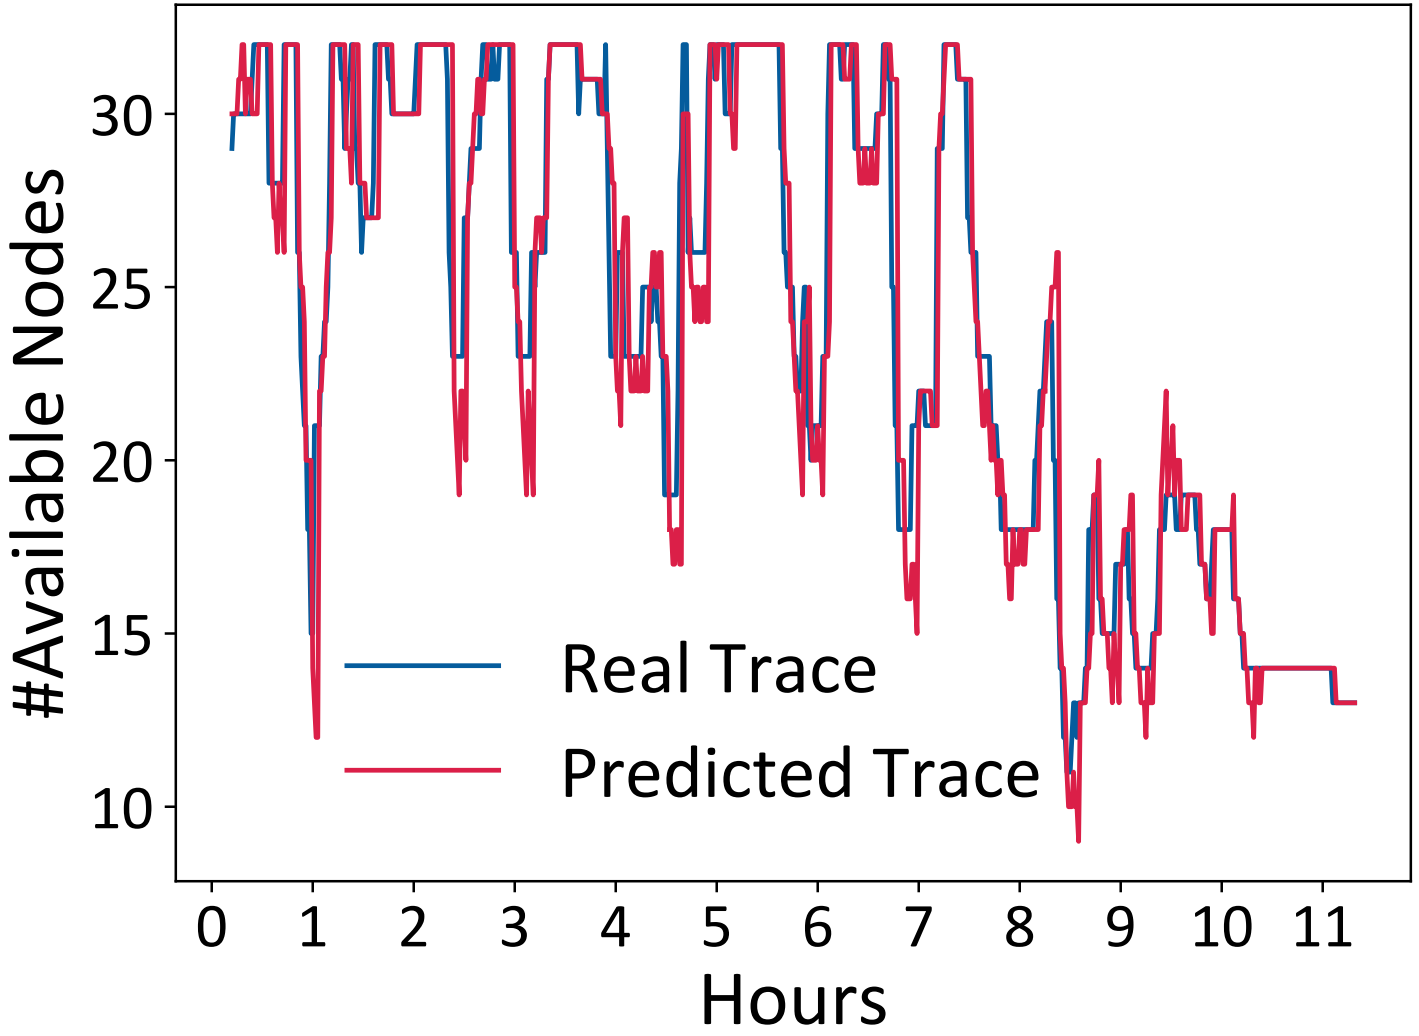

Availability Predictor

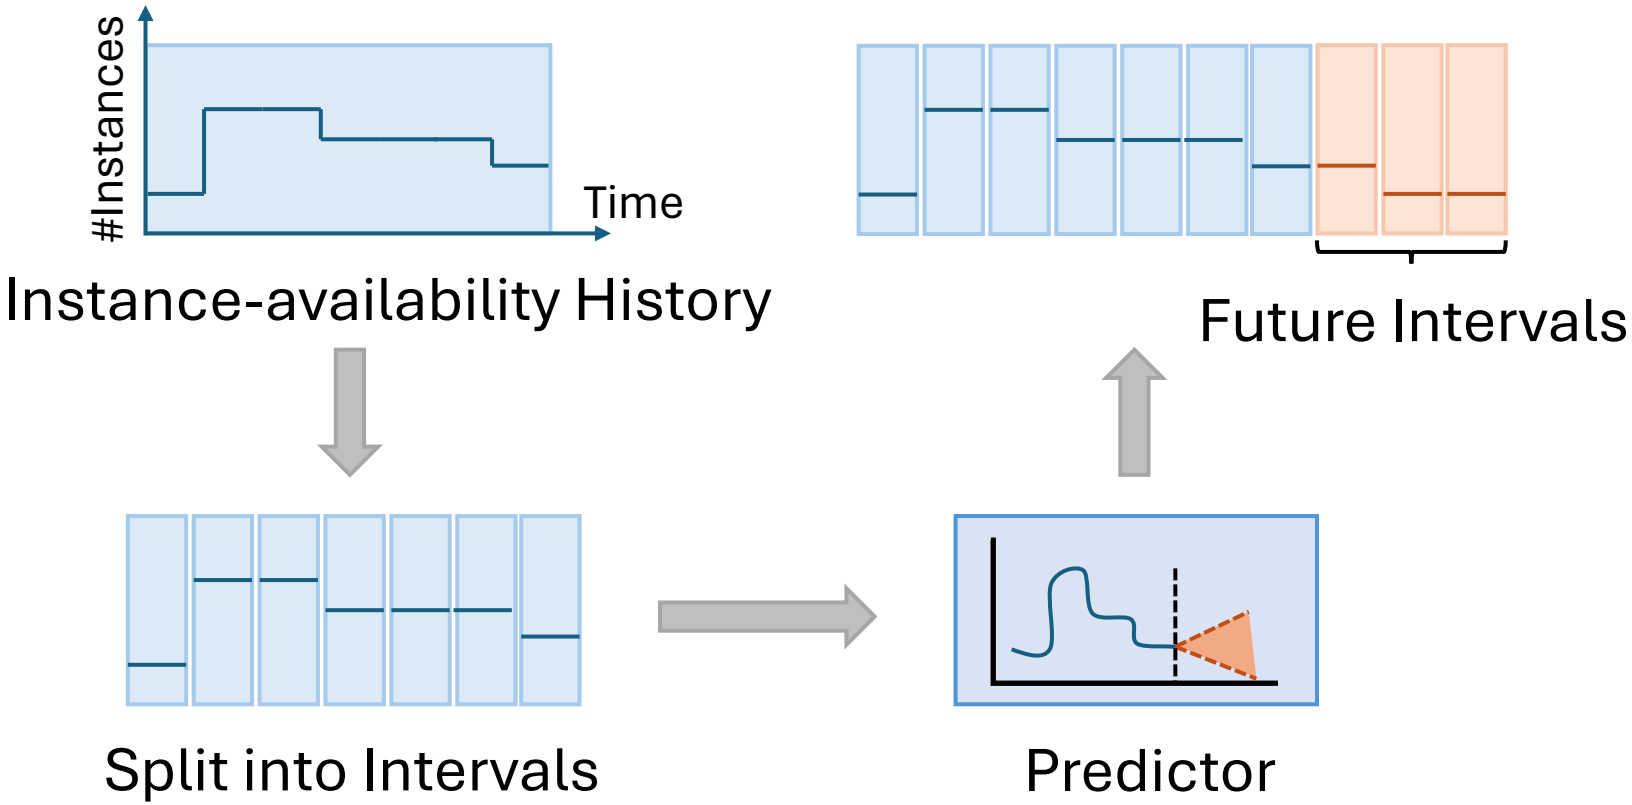

Time-series forecasting problem!

Preemption Mapping Sampling

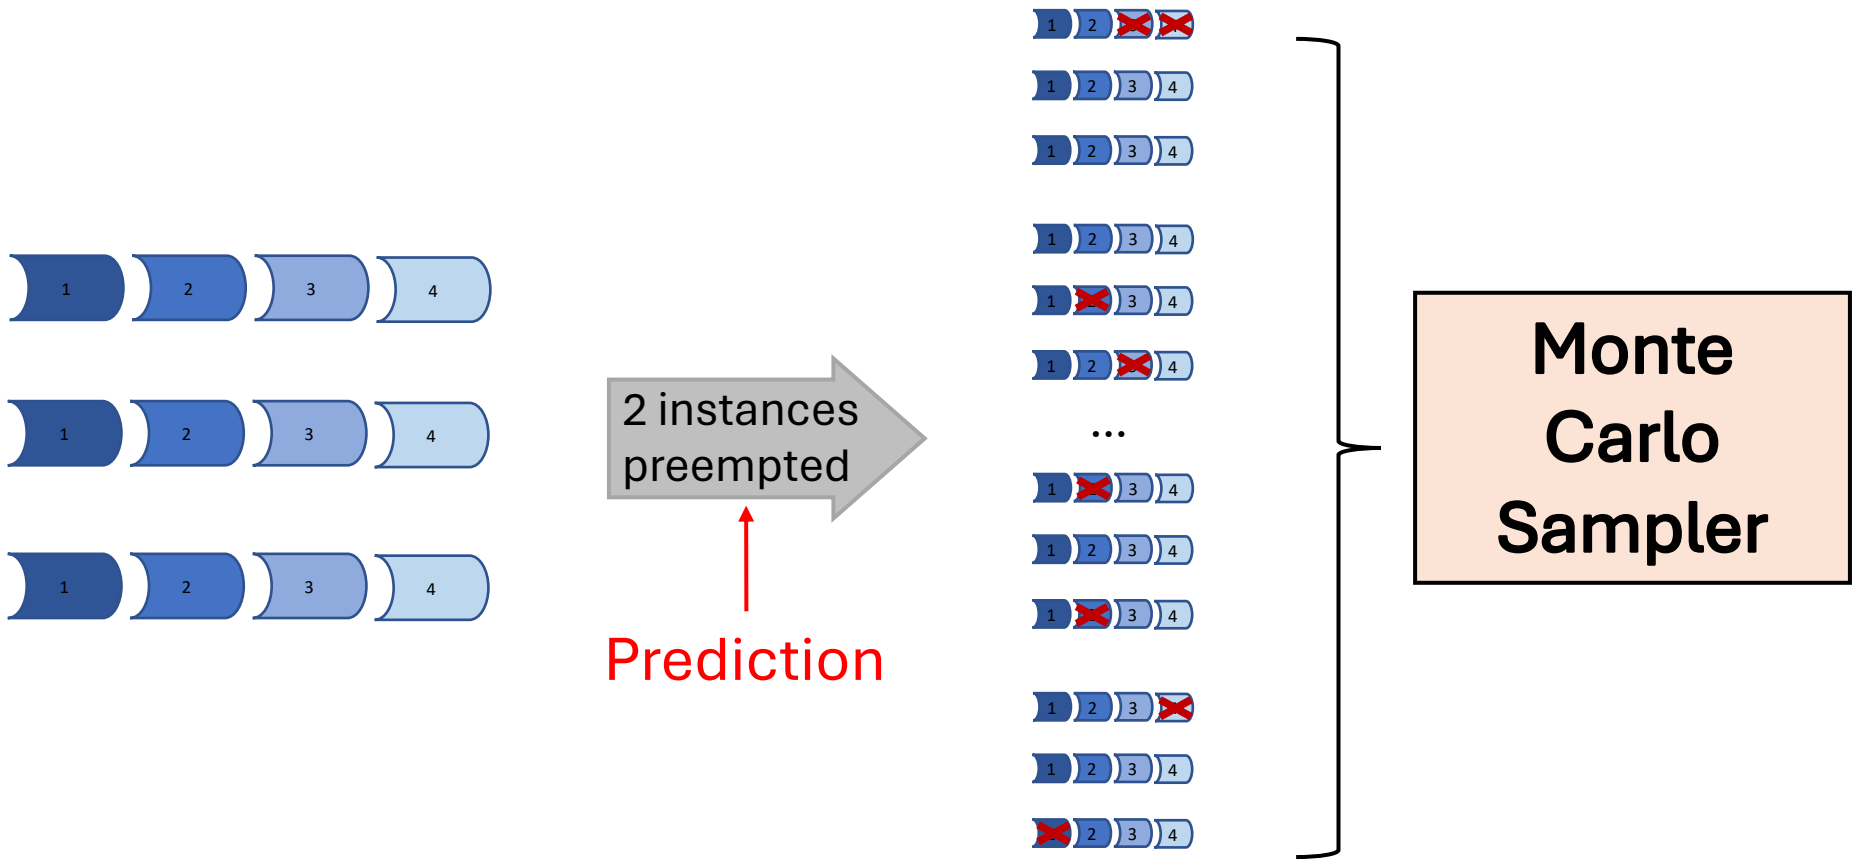

Large preemption mapping space

Liveput Optimizer

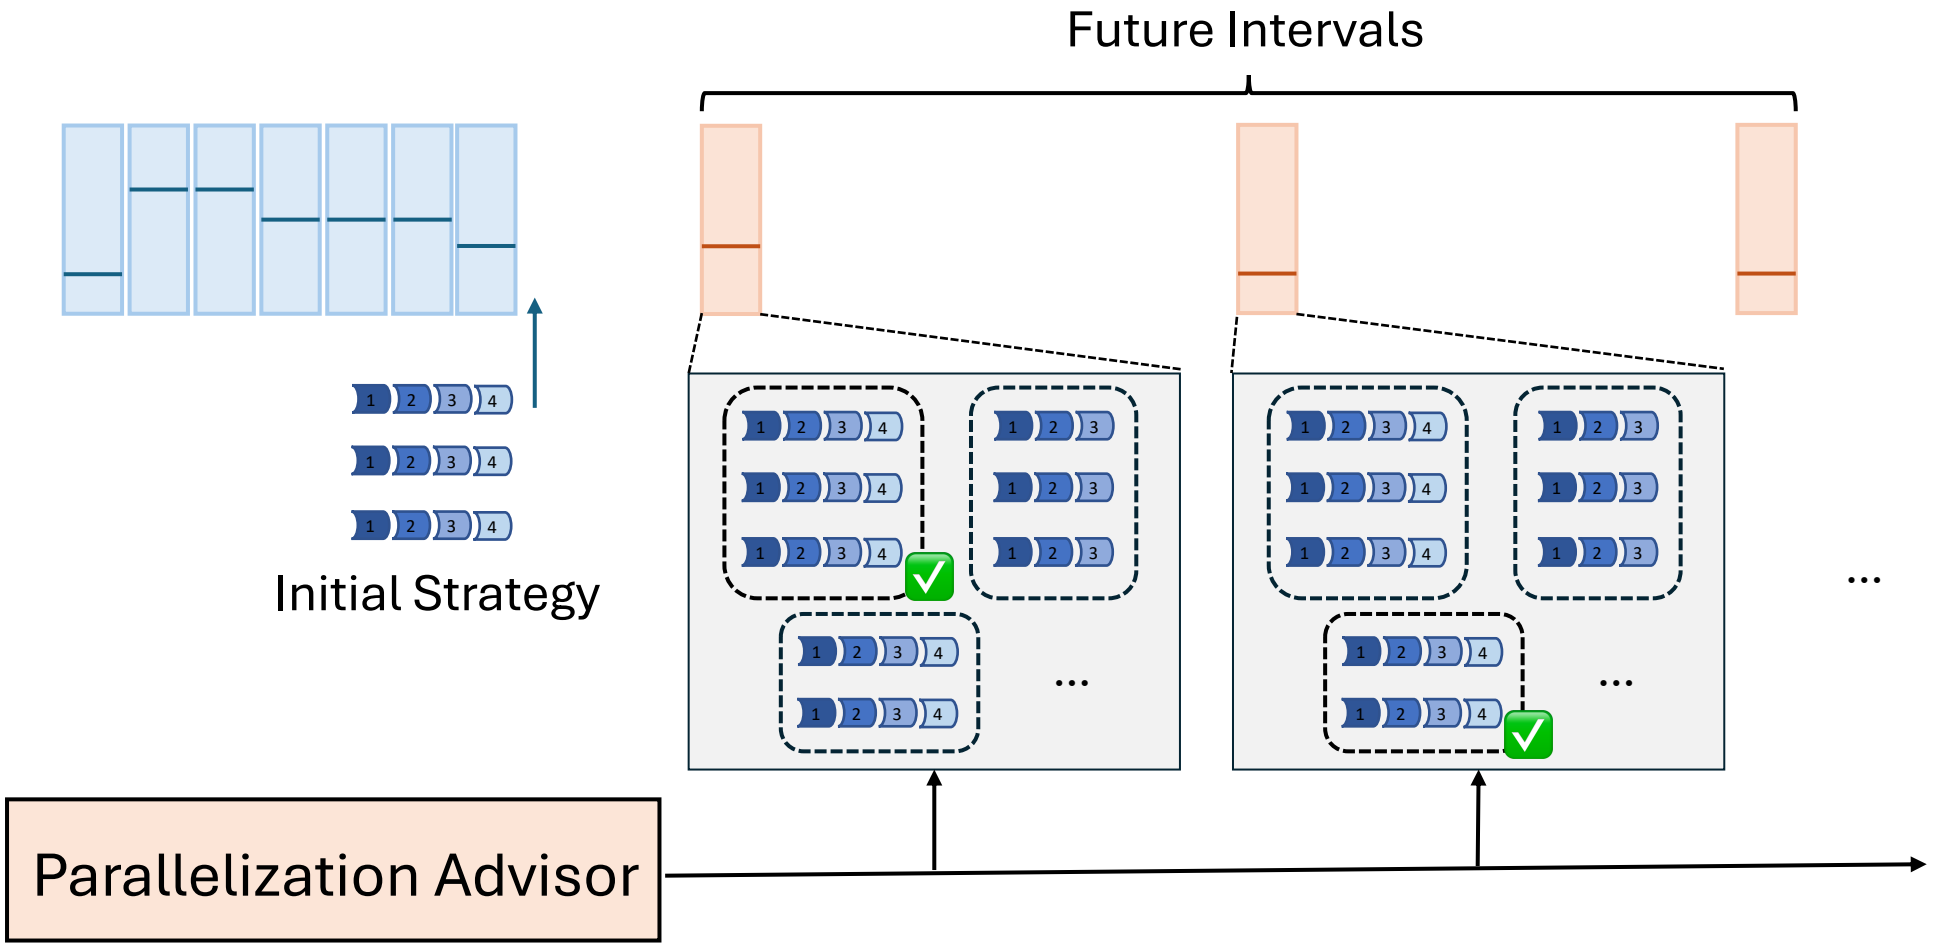

Combinatorial optimization space!

Fine-grained Live Migration

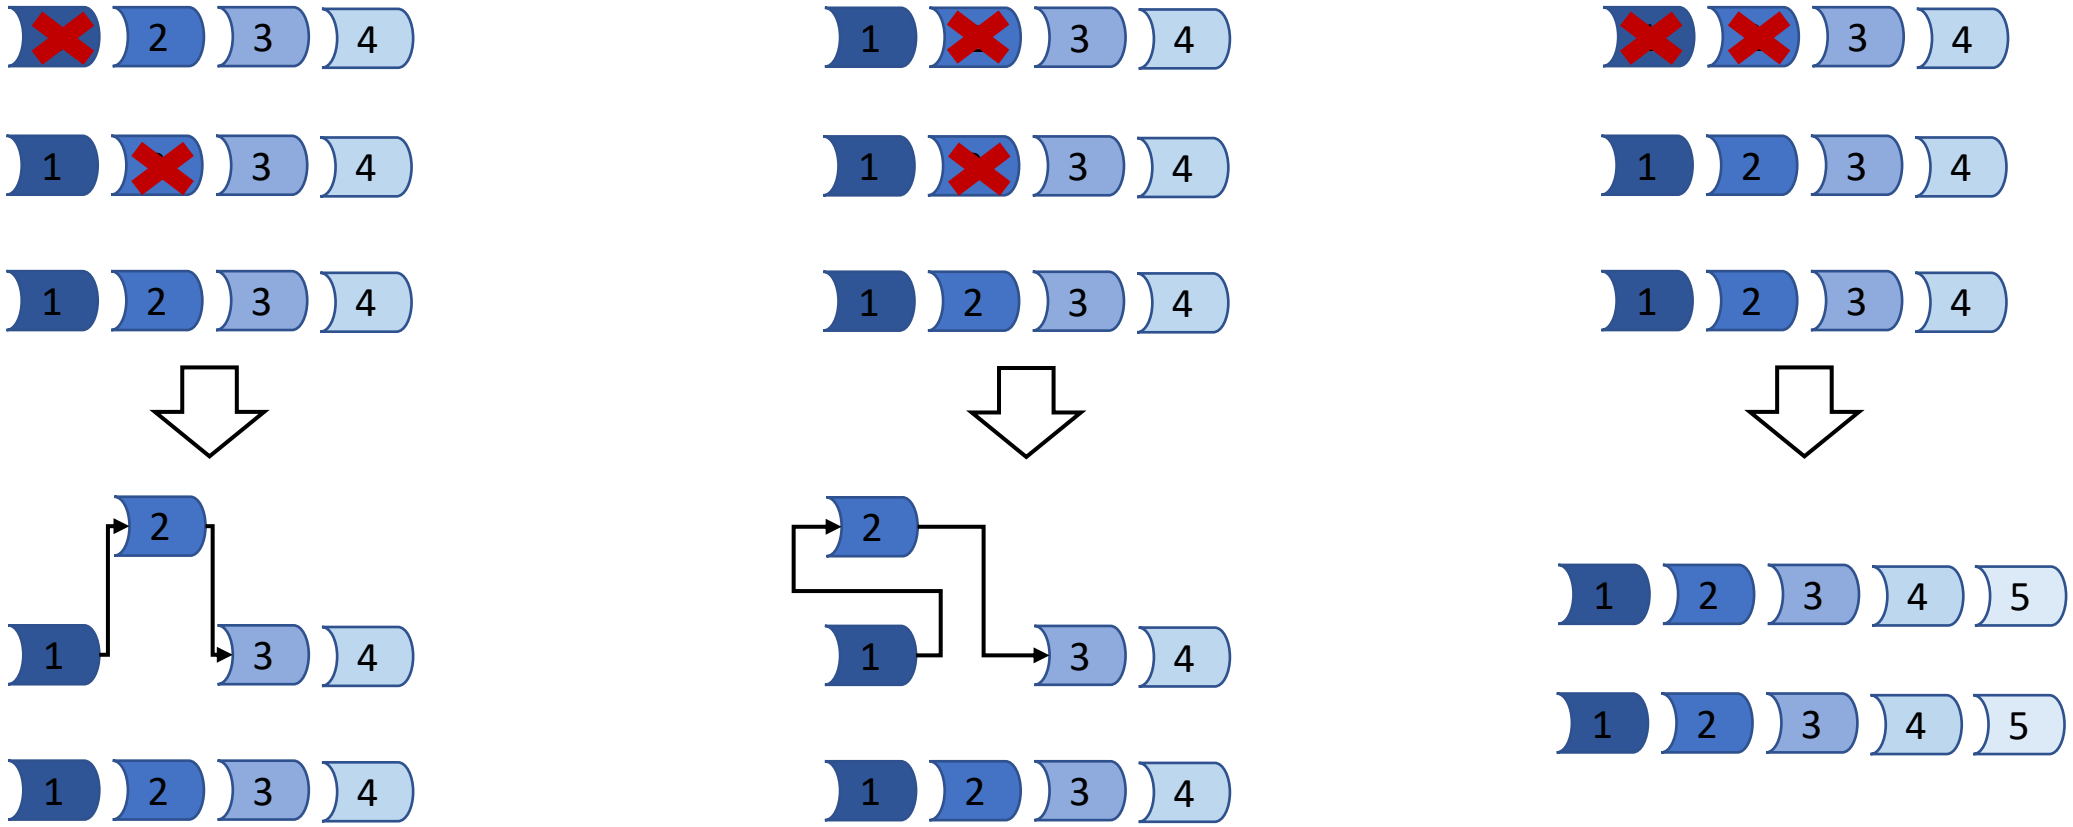

Intra-stage Migration

Inter-stage Migration

Pipeline Migration

GPT-3 (6B) Training on 32 Instances with Deterministic Replay

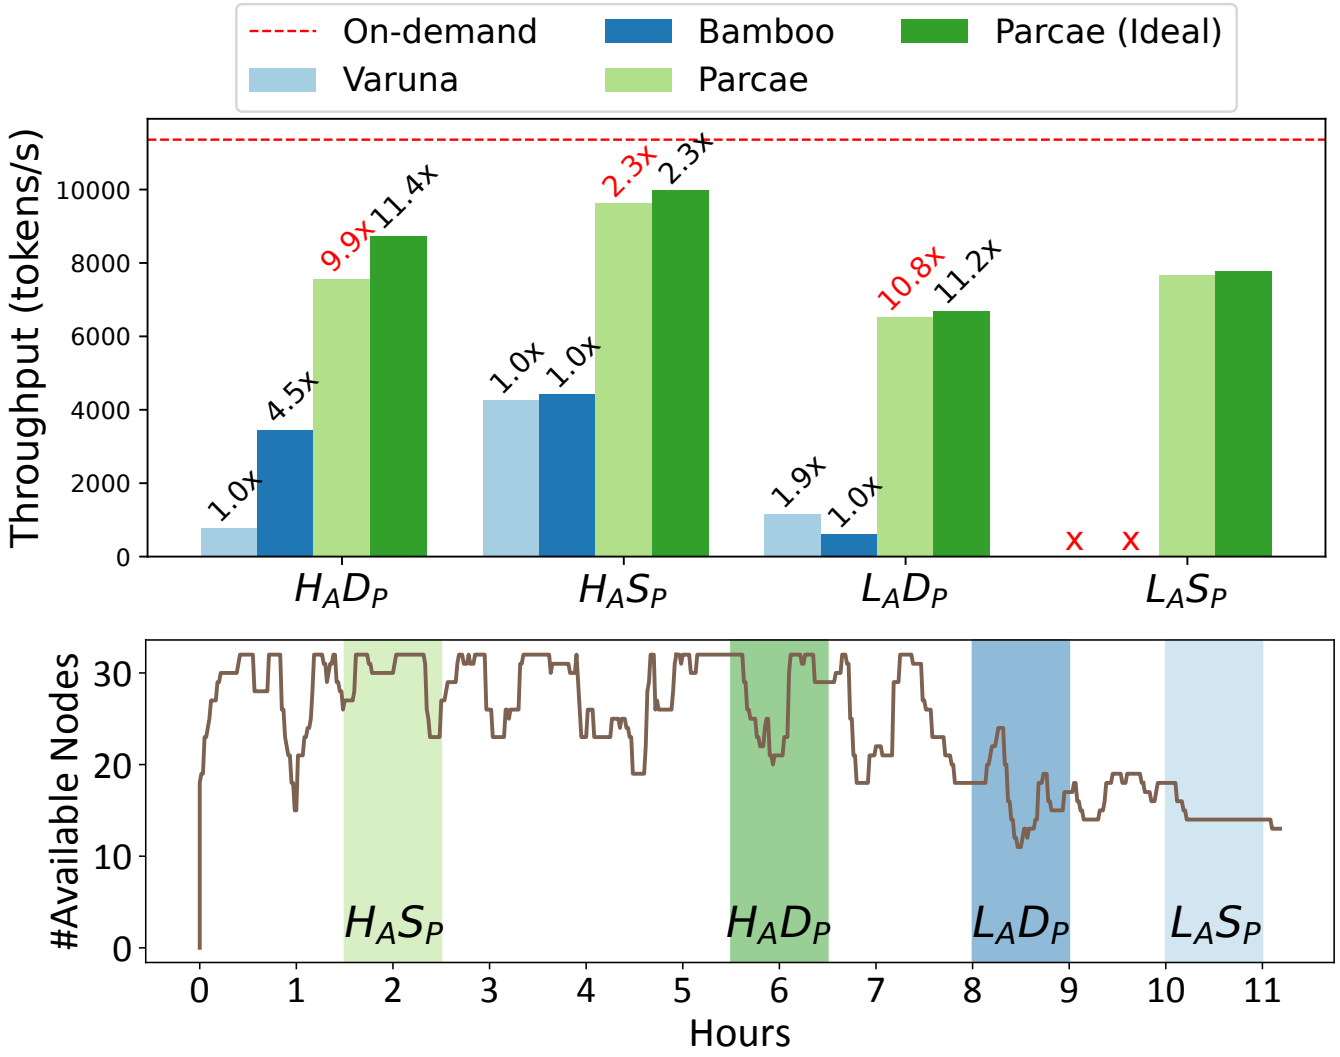

- Up to 10× higher throughput
- Better scalability
- Near optimal performance

Parcae

- Proactive, Liveput-Optimized DNN training on spot instances
- Up to 10× higher throughput than checkpoint- and redundancy-based systems
- Code: https://github.com/JF-D/Parcae