AReaL: A Large-Scale Asynchronous Reinforcement Learning System for Language Reasoning

Wei Fu (me), Jiaxuan Gao , Xujie Shen , Chen Zhu , Zhiyu Mei , Chuyi He , Shusheng Xu, Guo Wei , Jun Mei , Jiashu Wang, Tongkai Yang, Binhang Yuan, Yi Wu

Contact & Collaboration: fuwth17@gmail.com & jxwuyi@gmail.com

Challenges of Sync. RL Systems

Challenge 1: Significant ***GPU idle times during*** inference/generation when lengths vary

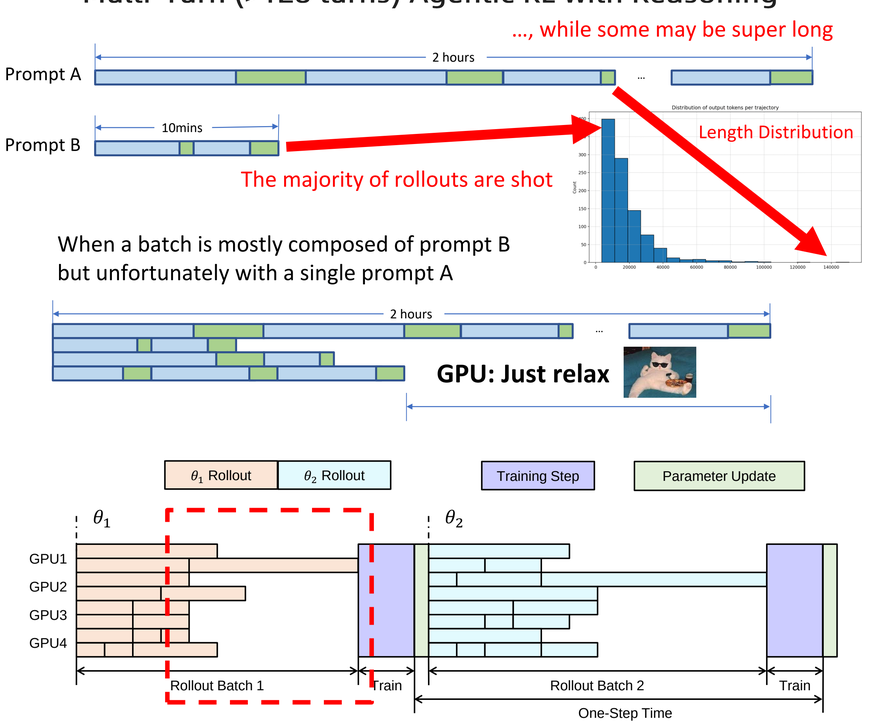

Challenge 2: ***Hard to scale up***

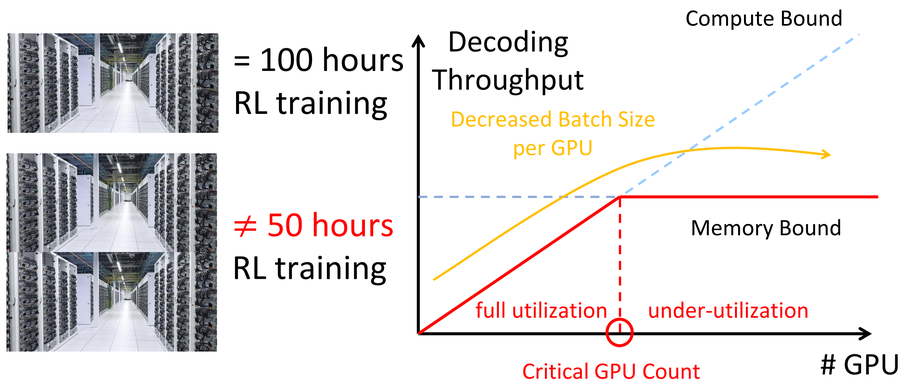

Fully Asynchronous RL in AReaL

Both training and inference achieve full GPU utilization

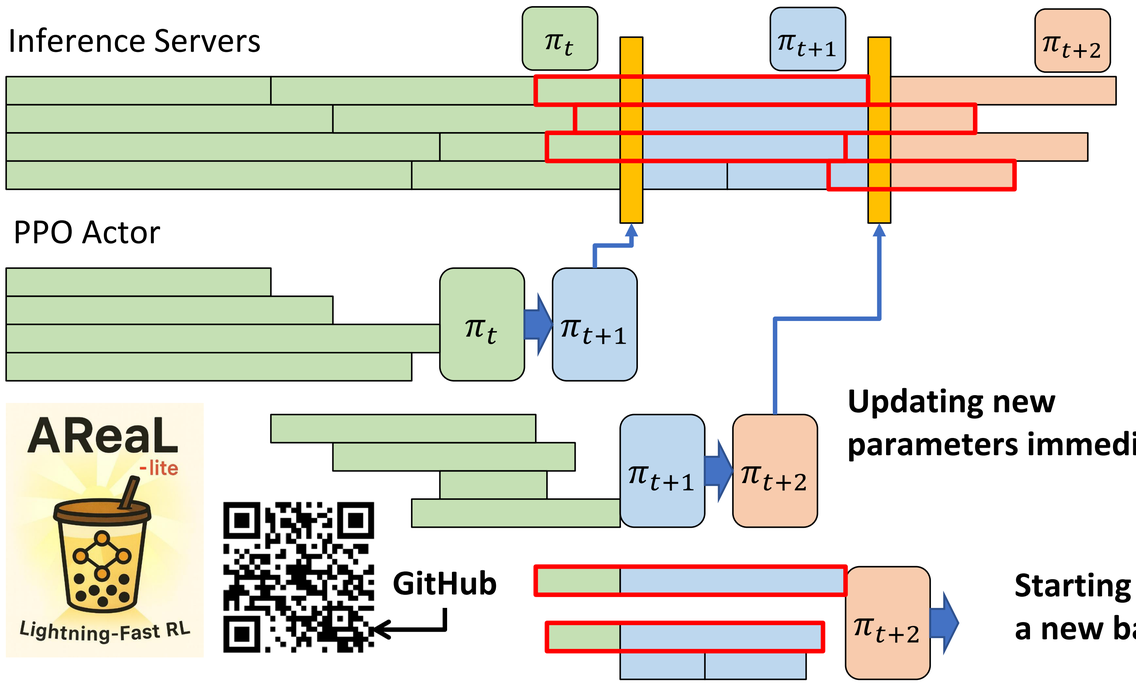

Keys to the challenges:

1. Continuous Rollout Batching with Interruption

2. Disaggregated and Overlapped Training & Inference

Rollout Interruption & Parameter Update

Full Trajectory

Cross-Version Trajectories

Algorithmic Challenges

•Data Staleness

Data generated from old policies impedes efficient learning.

•Inconsistent Policy Versions

Interrupted trajectories involve segments generated from different policy versions, which violates the PPO objective.

•Staleness-Aware Training

Control the maximum staleness with algorithm-system co-design.

•Decoupled PPO objective

$$J(\theta) = \mathbb{E}_{q\sim\mathcal{D}, a_t\sim\pi_{\text{behav}}}\left[\sum_{t=1}^{H} \min\Big(\underbrace{\frac{\pi_\theta}{\pi_{\text{behav}}}}_{\text{Importance Ratio}}\hat{A}_t,\ \overbrace{\frac{\pi_{\text{prox}}}{\pi_{\text{behav}}}\text{clip}\Big(\underbrace{\frac{\pi_\theta}{\pi_{\text{prox}}}}_{\text{Trust Region Center}}, 1-\epsilon, 1+\epsilon\Big)}^{\text{Importance Ratio}}\hat{A}_t\Big)\right]$$

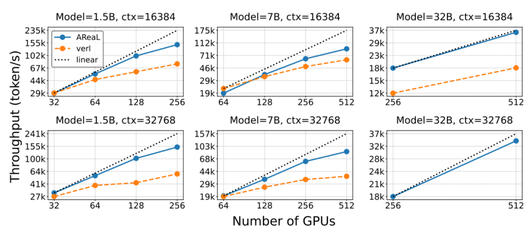

2x Higher **Effective Training Throughput Compared with veRL**

Model	LiveCodeBench ↑	Training Hours ↓
14B basemodel	53.4	-
w/ VeRL	57.9*	44.4
w/ Sync.AReaL	56.7	48.8
w/ AReaL (ours)	**58.1**	**21.9**
32B basemodel	57.4	-
w/ VeRL	-	46.4
w/ Sync.AReaL	**61.2**	51.1
w/ AReaL (ours)	61.0	**31.1**

2x Convergence Speed **in Competitive Programming Tasks**

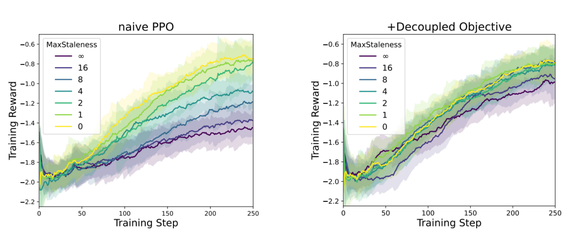

Both algorithm modifications are imperative **to async. RL performance**