

Highly Parallelized Reinforcement Learning Training with Relaxed Assignment Dependencies

Zhouyu He^{1,2*}, Peng Qiao^{1,2*}, Rongchun Li^{1,2†}, Yong Dou^{1,2}, Yusong Tan^{1,2‡}

¹College of Computer Science and Technology, National University of Defense Technology

²National Key Laboratory of Parallel and Distributed Computing, National University of Defense Technology
{rongchunli, ystan}@nudt.edu.cn

Abstract

As the demands for superior agents grow, the training complexity of Deep Reinforcement Learning (DRL) becomes higher. Thus, accelerating training of DRL has become a major research focus. Dividing the DRL training process into subtasks and using parallel computation can effectively reduce training costs. However, current DRL training systems lack sufficient parallelization due to data assignment between subtask components. This assignment issue has been ignored, but addressing it can further boost training efficiency. Therefore, we propose a high-throughput distributed RL training system called TianJi. It relaxes assignment dependencies between subtask components and enables event-driven asynchronous communication. Meanwhile, TianJi maintains clear boundaries between subtask components. To address convergence uncertainty from relaxed assignment dependencies, TianJi proposes a distributed strategy based on the balance of sample production and consumption. The strategy controls the staleness of samples to correct their quality, ensuring convergence. We conducted extensive experiments. TianJi achieves a convergence time acceleration ratio of up to 4.37 compared to related comparison systems. When scaled to eight computational nodes, TianJi shows a convergence time speedup of 1.6 and a throughput speedup of 7.13 relative to XingTian, demonstrating its capability to accelerate training and scalability. In data transmission efficiency experiments, TianJi significantly outperforms other systems, approaching hardware limits. TianJi also shows effectiveness in on-policy algorithms, achieving convergence time acceleration ratios of 4.36 and 2.95 compared to RLlib and XingTian.

Introduction

Reinforcement learning (RL) is a powerful approach for addressing decision-making problems. An agent is trained to interact with the environment based on its policy to maximize long-term rewards (Li 2017; Chadi and Mousannif 2023). To address complex spatial problems, value function approximation methods using neural networks are commonly employed (Sutton et al. 1999; Grondman et al. 2012; Schulman et al. 2015; Wilcox et al. 2022). Deep Reinforcement Learning (DRL) combines the strong expressive

power of deep learning with the decision-making capabilities of RL, offering a versatile system for flexible decision-making and control (Mnih et al. 2013; Lillicrap et al. 2015; Schulman et al. 2017). Research and applications of DRL are growing more complex (Brittain, Alvarez, and Bredden 2024; Wang et al. 2023; Weihs et al. 2020; Szot et al. 2023; Chen et al. 2023; Jiang et al. 2024). In complex scenarios, the scale of the state and action spaces is vast, and computational complexity grows exponentially (Kakade 2003; Neumann and Gros 2022; Li 2023; Liu et al. 2022; Shen et al. 2023b). Modern RL methods require substantial computational resources (Silver et al. 2016; Vinyals et al. 2019; Sun et al. 2023; Liu et al. 2024). For example, DQN (Mnih et al. 2013) needed continuous training for 12 to 14 days on a single game, involving millions of interactions. OpenAI Five (Berner et al. 2019) utilized 256 P100 GPUs and trained for 10 months. Due to these immense computational demands, RL research has increasingly focused on accelerating training.

Parallelization accelerates training by decomposing computational tasks and executing them simultaneously across multiple computing devices or even nodes (Isard et al. 2007; Dean et al. 2012; Liu et al. 2021; Qiao et al. 2023). The typical RL training process can be grouped into three subtasks, as ①②③ shown in Algorithm 1. These subtasks can be parallelized either individually or in combination. However, these parallel subtasks must be executed in a sequential order considering the assignment dependencies between them, as I-III shown in Algorithm 1. When these subtasks are assigned into different abstraction which we will discuss in the following, the assignment dependencies impair the parallel performance. Reasonably relaxing the limitations caused by the above assignment dependencies is the key to accelerating training.

Therefore, the previous works of accelerating the DRL training can be categorized into four groups. In Figure 1, we show the abstraction, data dependencies, communication, execution order, and resource utilization of various RL training systems. The Gorila-style architecture (Nair et al. 2015; Liang et al. 2018; Pan et al. 2022; Zhu et al. 2023; Mei et al. 2023) abstracts sampling-related tasks ① and ② as actors and task ③ as a learner, enabling parallelization by component. Gorila-style achieves parallelization of equivalent serial execution logic while preserving data dependen-

*These authors contributed equally.

†Corresponding author.

‡Corresponding author.

Algorithm 1: Pseudo-code for Function Approximation-based Temporal Difference(0)

Input: exploration rate ϵ , discount factor γ , network update rate α , value function V
Output: network parameters θ

- 1: Initialize environment, initialize state s .
- 2: **for** each episode **do**
- 3: **while** s is not terminal **do**
- 4: Get state s and weights θ .
 {I: Assignment dependence from ③ to ①.}
- 5: Choose action a with ϵ -greedy policy based on $V(s, \theta)$.
 {①: Policy inference.}
- 6: Get action a .
 {II: Assignment dependence from ① to ②.}
- 7: Take action a , get reward r and next state s' , $s = s'$.
 {②: Simulation execution.}
- 8: Get trajectory (s, a, r, s') .
 {III: Assignment dependence from ② to ③.}
- 9: Compute the TD error:
 $\delta = r + \gamma \times V(s'; \theta) - V(s; \theta)$;
 Update θ using SGD:
 $\theta = \theta + \alpha \times \delta \times \nabla V(s; \theta)$.
 {③: Policy update.}
- 10: **end while**
- 11: **end for**
- 12: **return** θ

cies (I, II, III). These dependencies result in process waiting, as evidenced by the numerous idle blocks in Figure 1(b). A detailed analysis of other architectures is provided in Section 2.2. Existing systems exhibit significant inter-component dependencies, which severely hinder performance.

Accelerating RL training is a rapidly evolving field. However, current systems often struggle with limited parallelization. They typically decompose large computational tasks into smaller subtasks for parallel processing, which leads to various inter-component assignment dependencies. Our analysis of parallelizability in RL shows that relaxing these dependencies is crucial for speeding up training. Existing work either overlooks or fails to achieve these relaxations (I, II, III). Therefore, we propose TianJi, a distributed RL training system designed to relax assignment dependencies. Our main contributions are as follows:

- We identified the relaxing of assignment dependencies I-III is crucial for accelerating the DRL training. Therefore, we propose TianJi, a distributed reinforcement learning training system that relaxes these dependencies.
- We propose a decentralized, data-driven training model that transforms inter-component assignment dependencies into asynchronous, loosely-coupled processes. TianJi defines clear component boundaries with internal self-loop computations and event-driven asynchronous communication between components.
- We introduce a distributed training strategy that balances the training sample production and consumption. Performance analysis manages sample staleness and adjusts sample distribution, ensuring convergence.
- We conducted extensive experiments in which TianJi demonstrated up to a 4.37-fold speedup in convergence

time compared to related systems. When scaled to eight computing nodes, TianJi achieved a 1.6-fold improvement in convergence time and a 7.13-fold increase in throughput, highlighting its training acceleration capabilities and scalability.

Related Work

Parallelization of Reinforcement Learning

Most RL methods are variants of the Temporal Difference (TD) (Sutton 1988). TD is a commonly used method in RL that updates the value function based on each step's trajectory. The TD(θ), shown in Algorithm 1, represents the general RL training process. The agent interacts with the environment through multiple discrete time steps t . At each time step t , the agent observes a state s and selects an action a from a set of possible actions according to the policy θ (denoted as ①). The agent then interacts with the environment, which advances the simulator to yield the next state s' and a scalar reward r (denoted as ②). Steps ① and ② are repeated until the agent reaches a terminal state or the specified time step, which is simplified to one step in Algorithm 1. During training, the TD error is calculated from the collected trajectories to learn the state value function under the policy from a series of incomplete episodes (denoted as ③).

Computation task ①, ②, and ③ can be parallelized either individually or in combination. Abstracting these tasks as separate or combined components leads to varying levels of data isolation due to different degrees of computational isolation. The dependencies between these tasks are as follows: ① requires the state s from ② and the latest network parameters θ from ③ (denoted as I); ② needs the action a from ① (denoted as II); and ③ requires the trajectories collected from ② (denoted as III). When parallelizing, it is crucial to consider the constraints imposed by these dependencies. Therefore, relaxing these dependencies is key to achieving high levels of parallelism. Shen (Shen et al. 2023a) indicated that the parallelism and asynchrony of A3C accelerate convergence in theory. GALA further relaxes the model dependency among workers, resulting in faster convergence than the synchronized A2C. Inspired by previous theoretical and experimental works, we relax assignment dependencies to better utilize computational resources.

Parallelization of Existing DRL Training Systems

The Gorila-style architecture (Nair et al. 2015; Liang et al. 2018; Pan et al. 2022; Zhu et al. 2023; Mei et al. 2023) abstracts sampling-related ① and ② as actors, and ③ as learners, enabling parallel execution. This approach isolates sample production from consumption but requires components to exchange samples and parameters, increasing inter-component communication. Gorila contributed to increasing sampling parallelism by decoupling sampling from model training, allowing each to be parallelized independently. Subsequent developments in Gorila-style architectures include XingTian, which addressed communication bottlenecks with an asynchronous communication channel; SRL, which proposed a data flow-based abstraction to improve resource efficiency. Although these improvements optimize

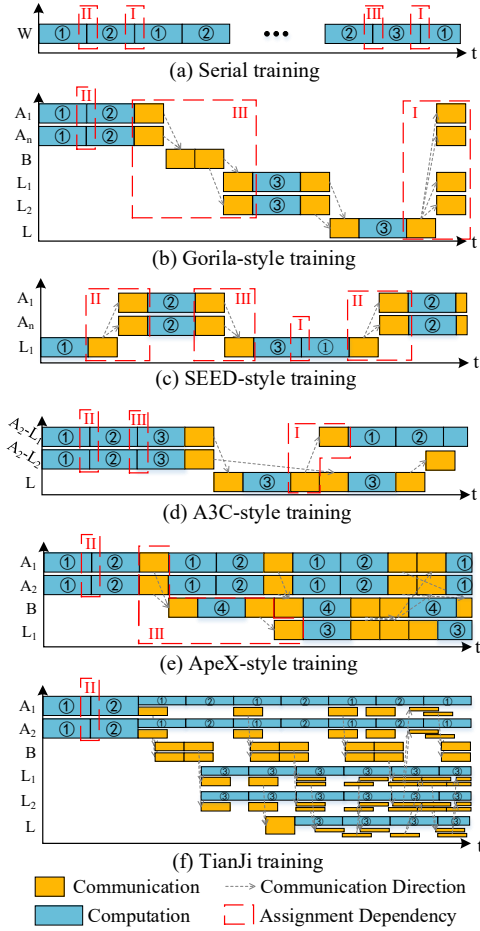


Figure 1: The spatiotemporal diagram of typical DRL systems during training. The annotations match those in Algorithm 1, with ④ representing the computation of Prioritized Experience Replay (PER). W denotes a worker, A an actor, B a buffer, and L a learner. The A2-L1 represents a component comprising two actors and one learner. This diagram illustrates the abstraction, assignment dependencies, communication patterns, execution sequence, and resource utilization in systems.

sample quality, communication, and task-resource mapping, they do not contribute to enhancing training parallelism. As shown in Figure 1, Gorila-style architectures are constrained by component assignment dependencies(I, II, III), resulting in significant idle periods.

The SEED-style architecture (Espoholt et al. 2019; Petrenko et al. 2020; Zhu et al. 2023; Mei et al. 2023) identifies inefficiencies in resource utilization present in the Gorila-style architecture. In Gorila, actors alternate between two dissimilar tasks, ① and ②. In SEED-style, model-related tasks ① and ③ are abstracted into a learner, while actors handle ②. At each environment step, the state is sent to the learner, which infers the action and then returns it to the actor. This introduces a new issue: latency. To address this issue, SEED implements a high-performance gRPC library, and SampleFactory designs dual-buffer sampling method.

Similar to Gorila-style, SEED-style architecture also employs equivalent serial execution logic, constrained by component assignment dependencies(I, II, III).

The A3C-style architecture (Mnih et al. 2016; Assran et al. 2019) abstracts tasks ①, ②, and ③ into a single worker, extending computation across multiple workers. Each worker typically uses a Gorila-like architecture internally, with isolated model parameters between workers. A3C employs an Hogwild!-like(Recht et al. 2011) asynchronous update method, while GALA uses a ring-based asynchronous update method. A3C-style reduces parameter synchronization dependencies among workers, so dependency III is only related to the global worker or neighboring workers, significantly increasing training throughput. However, assignment dependencies within each worker(I and II) still exist.

The ApeX-style architecture(Horgan et al. 2018; Espoholt et al. 2018; Assran et al. 2019) is conceptually similar to Gorila-style but theoretically relaxes dependencies I and III. However, this relaxation introduces convergence uncertainty. At any given time, the policy parameters θ among multiple actors are inconsistent, as are the policies of the actors and the learner. To address this, ApeX designs Prioritized Experience Replay(PER)(Schaul et al. 2015) to correct staleness. However, PER requires data consistency. The global priority segment tree maintenance and updates prevent asynchronous data transfer, effectively moving dependency III to the Buffer, which limits the acceleration of training throughput.

This Work

We found that assignment dependencies significantly restrict parallelization. To address this, we proposed TianJi, a distributed reinforcement learning training system that relaxes these dependencies. TianJi abstracts ① and ② into a single actor, thereby mitigating the negative impact of II. TianJi introduces decentralized data-driven training, transforming assignment dependencies I and III between components into asynchronous data exchanges (see Section 3.1). Decentralized computing leads to data isolation, resulting in inter-process data transfers. Efficient data transfer is the base of achieving high-throughput training. TianJi implements asynchronous communication with hidden overheads (see Section 3.2). Although relaxing assignment dependencies can introduce data staleness, TianJi addresses this by employing a distributed strategy based on production-consumption balance, ensuring training convergence through sample distribution adjustment. This strategy also facilitates scaling, enabling the system to overcome performance bottlenecks (see Section 3.3).

Decentralized Data-driven Training

Existing systems typically use a global core component to allocate and coordinate multiple instances, ensuring equivalence between the parallel and serial implementations. These implementations maintain assignment dependencies either fully or partially, resulting in significant idle time during training, as illustrated in Figure 1. Therefore, TianJi reorganizes DRL computations to enhance parallelism and introduces a decentralized data-driven training approach. In

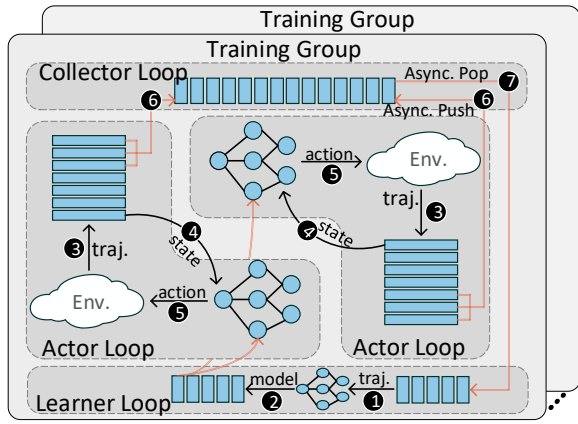


Figure 2: Data-driven training flowchart with relaxed assignment dependencies.

this approach, assignment dependency II is reduced to intracomponent dependency, while dependencies I and III are converted to asynchronous data exchanges.

TianJi abstracts sampling-related computation tasks (①, ②) as "Actor" and model updates (③) as "Learner," parallelizing these at the component level. Figure 2 illustrates the training process of TianJi. Once computation begins, each role independently performs looped computations and triggers data exchanges between components at appropriate times. In the Learner Loop, steps (①, ②) are executed repeatedly. Step ① involves loading trajectories from local storage, using an asynchronous communication trigger to determine if data reception is needed. Step ② involves learning and updating the model, which is then stored locally, with a trigger to decide whether sending model parameters is necessary. In the Actor Loop, steps (③-⑤) are executed in sequence. Step ③ involves the agent interacting with the environment to obtain trajectories or states, which are stored locally, with a trigger to determine if these trajectories should be sent to the buffer. Step ④ involves model inference, with a trigger to decide whether a new model should be received. Step ⑤ involves generating actions through model inference and sending these actions to the environment. In the Buffer Loop, steps (⑥, ⑦) are executed repeatedly to manage asynchronous data reception and sending. TianJi supports scaling the computation across multiple training groups.

Event-driven Asynchronous Communication

Distributed components create data isolation, necessitating communication between components. Experiments have shown that communication can sometimes take more time than computation (Pan et al. 2022; Zhao et al. 2023). We observed that learners do not require samples to be uniformly distributed across actors. Therefore, common practices such as group communication or uniformly requesting data from each actor can be both impractical and costly. TianJi employs an actively pushing asynchronous communication mode, where communication requests are initiated by the sender. Once the data at the sender is ready, it immediately triggers the communication, actively pushing the data to the intended recipient. The receiver uses a prob-

ing mechanism. The sender does not wait for an immediate request and continues its computations. Communication is confirmed during subsequent data transmissions, allowing computation and communication to overlap.

We will analyze the overlap between communication and computation in TianJi by examining critical path transitions during the sample collection phase. In this phase, the Actor collects data and sends it to the Buffer, where the samples are received. Figure 3 illustrates the critical path transitions in both "Single Actor" and "Multiple Actors" scenarios. In the "Single Actor" scenario, the sampling time exceeds the reception time. Causing computation overshadows communication, making the Actor's sampling the critical determinant of the path. Conversely, in the "Multiple Actors" scenario, increasing the number of actors effectively reduces the sampling time. When the sampling time becomes shorter than the communication time, communication can no longer be entirely concealed, shifting the key path to the Buffer. Although computation can obscure communication, this concealment has its limitations.

The theoretical collect time T_c can be calculated. When the communication hidden limit is not reached, the communication time can be masked by the computation time, placing the critical path within the "Actor". When the communication hidden limit is reached, the communication time can no longer be masked by the computation time, causing the critical path to shift to the Buffer. From the Buffer's perspective, increasing the number of actors is equivalent to linearly accelerating the sample collection process. Based on the critical path transition rules, we can derive the following formula:

$$T_c = \begin{cases} (T_{sp} + T_{sd}) \times N_s, & \text{if } \frac{(T_s + T_{sd})}{N_A} > T_{rv} \\ T_{rv} \times N_s, & \text{others} \end{cases} \quad (1)$$

where, N_s is the number of collected trajectories, T_{sp} is the time of a single sampling. T_{sd} and T_{rv} are the time of sending and receiving, respectively. N_A is the number of actors.

Event-driven asynchronous communication eliminates redundant waiting and enables the overlap of computation and communication, thereby facilitating more efficient and rational data exchange between components.

Distribution Strategy

Distributed strategies handle the configuration of components and their allocation to computing resources, impacting the location of performance bottlenecks and resource utilization. When a performance bottleneck arises, efforts outside this bottleneck do not contribute to convergence. Extensive experiments have demonstrated that performance bottlenecks vary across different algorithms, applications, and hardware. These studies indicate that some algorithms are more prone to bottlenecks in specific areas, resulting in fixed strategies targeting particular performance issues. However, because of the unpredictability of performance bottlenecks, fixed optimization techniques cannot guarantee overall throughput improvement. Reverb (Cassirer et al. 2021) proposes SPI rate limiting, which controls when items can be inserted into or sampled from a table. However, SPI

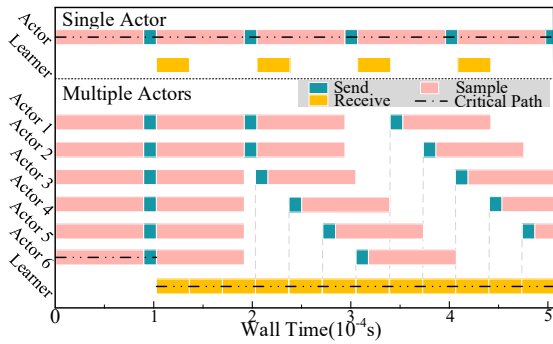


Figure 3: As the number of actors increases, the critical path shifts. Critical path is the sequence of tasks that determines the minimum time required to complete the computation.

reflects the ratio over a period of time, varies significantly across different scales, and involves blocking operations. SPI is both inaccurate and inefficient. Therefore, TianJi proposes a distributed strategy based on performance analysis and production-consumption balance to enable scalable training that addresses performance bottlenecks. This strategy also maintains sample freshness by ensuring consistency between sample distribution and serial execution, which is essential for training convergence.

Given the algorithm, application, and computing resources, the goal of a distributed strategy is to maximize computational throughput and ensure training convergence. TianJi allows the assignment of specific computing resources to each component, which can be easily managed using a configuration file. For example, with a only CPUs setup and a total of M cores, performance analysis tools can measure single-step training to obtain the consumption throughput TR_L and production throughput TR_A .

In serial execution, once the sample pool is filled, the sample staleness P stabilizes. The formula is:

$$P = \frac{B_s}{N_s} \quad (2)$$

where, B_s is the size of a batch trajectories, N_s is the size of buffer.

The optimization objective under the distributed strategy is to maximize training throughput while controlling the sample obsolescence ratio, as follows:

$$\text{Maximize } \min(TR_L, TR_A) \quad (3)$$

$$\text{Subject to: } \begin{cases} \frac{N_c}{N_s} = P \\ M_L + M_A = M \\ TR_L \approx TR_A \end{cases} \quad (4)$$

where, N_c is the number of trajectories consumed per unit time, M_A and M_L are the number of cores obtained by the Learner and the Actor, respectively.

Through performance analysis and objective function optimization, TianJi achieves load balancing and optimal mapping of computational tasks and resources.

Evaluation

The evaluation answer the following questions: (1) How does TianJi compare to existing systems in terms of over-

all performance optimization (Section 4.2)? (2) How effective are the key components as demonstrated by the ablation study (Section 4.3)?

Setting

Testbed. We configured two hardware platforms for our experiments. The first platform is a CPU-only Slurm cluster with 8 computing nodes. Each node is equipped with 2 Intel Xeon Gold 6248 processors, providing a total of 40 physical cores per node. Each node has 384GB of memory, and they are interconnected using ConnectX-6 high-speed interconnects. The second platform is a heterogeneous machine, equipped with one A100 GPU, 40 physical cores, and 376GB of memory.

Baseline. The three baselines are: ApeX, a distributed architecture for large-scale deep reinforcement learning that employs prioritized experience replay; RLlib, a widely used DRL training system based on Ray; and XingTian, which leverages communication-computation overlap in DRL to offer a more efficient asynchronous communication channel than other systems.

Algorithms & Environment. We conducted experiments in Atari and OpenAI Gym. The algorithms evaluated are DQN(off-policy) and PPO(on-policy). We used RLlib’s default network architecture and parameters as benchmarks for both Gym and Atari tasks. All comparisons used identical network architectures and hyperparameters across the same games to ensure fairness.

Metrics. The evaluation focuses on learning throughput and performance. learning throughput: This includes sample throughput, receive throughput, and training throughput, which represent the number of frames processed per second. learning performance: This is assessed using episode return and final time, which indicate the reward value and the time taken to achieve a specified reward, respectively. Reward are averaged over a sliding window of 100 episodes.

Overall Performance Comparison

ApeX relaxes certain assignment dependencies and is the most similar work to TianJi. Since ApeX does not support multi-core processing or multiple learners, experiments are conducted on a single GPU machine. Figure 4(a) illustrates the episode return over time, highlighting the learning performance of each system. The current episode return is derived from clipped rewards within a single life. TianJi demonstrates superior learning performance compared to ApeX. TianJi’s distributed strategy, based on production-consumption balance, corrects the sample distribution and ensures. In contrast, ApeX still encounters assignment dependencies due to the data consistency requirements of PER. Conversely, TianJi achieves true relaxation of assignment dependencies, significantly enhancing computational efficiency, as shown in Figure 4(b). ApeX’s open-source implementation is non-scalable, restricting comparisons to a single node. Theoretically, PER’s negative impact on computational efficiency worsens with an increasing number of components.

Additional performance comparisons were carried out on eight CPU machines, as illustrated in Figure 4(c). For on-

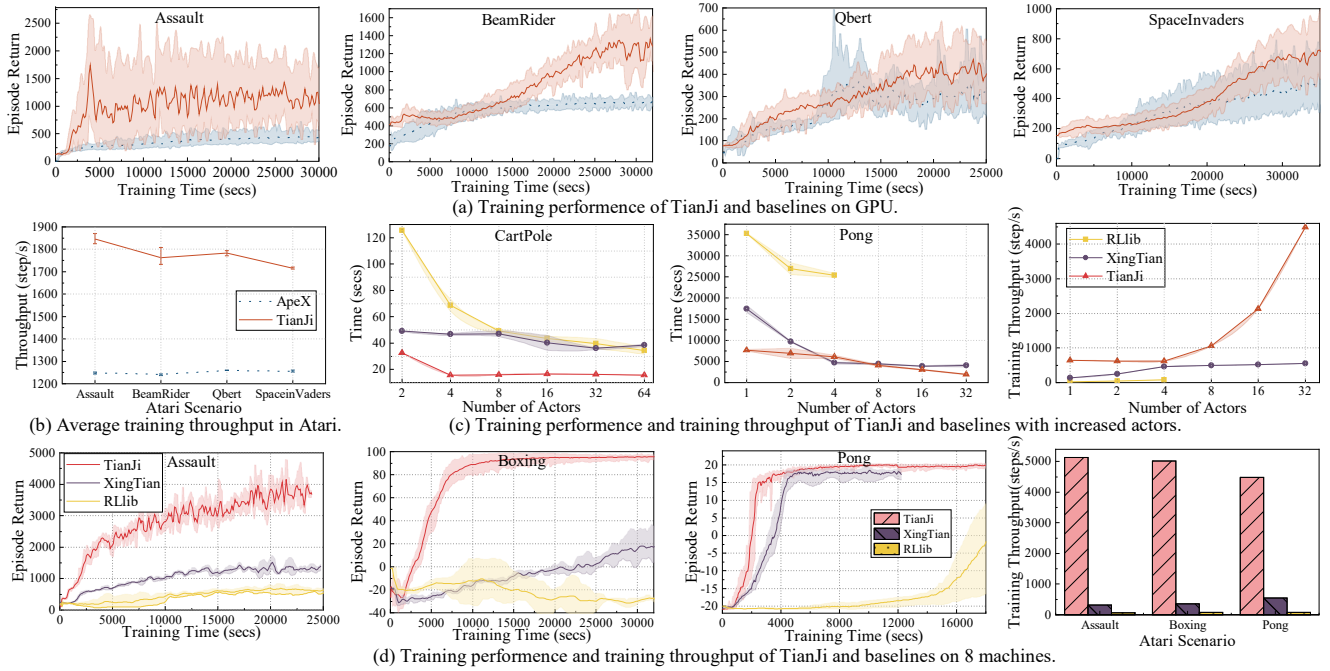


Figure 4: TianJi outperforms the baselines on learning performance and computational efficiency.

policy method learning Cartpole, we compared the final time to reach an average reward of 300. TianJi outperformed the baselines with varying numbers of actors. Compared to RLib, TianJi achieved a final time speedup of up to 4.36. For off-policy method learning Pong, we compared the final time and training throughput. When the number of actors exceeds 4, XingTian’s final time tends to stabilize. This occurs because the performance bottleneck shifts to training, and increasing sampling efficiency does not improve convergence. The training throughput data show that TianJi’s distributed strategy effectively addresses performance bottlenecks. Compared to XingTian, TianJi achieves 1.6-fold faster convergence and 7.13-fold higher computational efficiency. The experimental results demonstrate that TianJi achieves several times higher training throughput and faster convergence compared to the baselines. Figure 4(d) provides a detailed performance comparison across eight computational nodes, where episode returns are calculated from game rewards within a single life. TianJi exhibits markedly superior learning performance compared to XingTian and RLib while achieving significantly higher throughput.

Ablation Study

Data Transmission Efficiency We conducted a virtual communication experiment to evaluate the data transmission efficiency of TianJi and other systems. The experiment, which followed the communication pattern of a standard DistRL algorithm, concluded after collecting 10,000 samples, with collection time recorded. As a baseline, we selected XingTian, which offers an asynchronous communication channel and exhibits superior data transmission efficiency compared to other systems. We evaluated the data transmission efficiency by measuring the time to receive the

10,000 samples(referred to as collection time) and the data transmission throughput(referred to as throughput).

Figure 5(a) shows that TianJi consistently outperforms XingTian in data transmission efficiency for all message sizes under single-node and single-actor conditions. Figure 5(b) depicts how collection time varies with the number of actors while the message size remains fixed at 512 KB. The figure indicates that, for a fixed sample time(e.g., 0.001), communication bottlenecks occur with more than four actors, and increasing the number of actors beyond this does not reduce collection time(see Section 3.2.2 for analysis). With a sample time of 0.01, the communication bottleneck occurs with 16 actors. TianJi consistently exhibits lower sample collection times than XingTian across all conditions. Notably, TianJi has a lower communication bottleneck than XingTian. After reaching the bottleneck, the collection time depends only on receiver operations(e.g., local storage), and TianJi’s reception time remains lower. Figure 5(c) shows how communication efficiency varies as the number of actors increases across multiple computing nodes, with actors evenly distributed among 4 nodes. TianJi consistently has lower collection times than XingTian in all scenarios. Figure 5(d) shows that TianJi achieves throughput close to Ethernet bandwidth. Testing was conducted between two nodes with Ethernet bandwidth of 125 MB/s and InfiniBand(IB) bandwidth of 12.5 GB/s. Under Ethernet conditions, both TianJi and XingTian show similar performance trends: as message size increases, throughput gradually rises and approaches Ethernet bandwidth. In contrast, under IB conditions, TianJi exceeds previous throughput limits, while XingTian does not. Throughput decreases after 2048 KB due to increased message generation time with larger message sizes. TianJi reaches the Ethernet bandwidth limits, demonstrating supe-

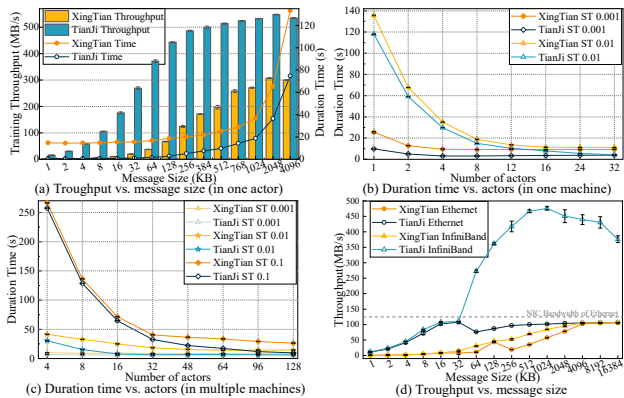


Figure 5: Comparison of data transfer efficiency between TianJi and XingTian. ST is "sample time".

rior communication efficiency compared to XingTian.

The Impact of the Distributed Strategy on Convergence

TianJi proposes a distributed strategy based on production-consumption balance to correct sample distribution and maximize training throughput. This section examines the impact of sample staleness and computational-resource mapping on training. Training was conducted on a single node with PPO. The optimal computational-resource mapping, identified by the distributed strategy, includes 2 learners and 4 actors, using 16 cores, denoted as L2A8-C16. A random computational-resource mapping, labeled L1A14-C16, was also evaluated. After introducing asynchrony, simulations with serial sample distributions were conducted using two sample ratios: 1:1 (denoted as New) and 1:8 (denoted as Staleness). Four sets of control experiments were conducted. The results, shown in Figure 6(d), indicate that optimal computational-resource mapping and production-consumption balance significantly accelerate computational efficiency and enhance learning performance.

Scaling Beyond Performance Bottlenecks Extensive experiments show that once training hits a performance bottleneck, adding resources beyond this point does not improve training performance. Predicting the location of this bottleneck is challenging because various algorithms, applications, and hardware configurations can create bottlenecks in different areas. This section will demonstrate how TianJi uses effective distribution methods to tackle different performance bottlenecks. TianJi achieves accelerated convergence beyond performance bottlenecks and scales training throughput to near-linear levels. Performance bottlenecks can be categorized into three types: sample-intensive (①②), training-intensive (③), and communication-intensive. In the DQN(Cartpole) experiment, Figure 6(a) illustrates how sampling throughput, receiving throughput, and training throughput change as the number of actors increases in the DQN(Cartpole) experiment. The number of learners, resources, and computational load is kept constant, resulting in stable training throughput (blue dashed line). As the number of actors increases, sampling throughput (red solid line) also rises. The receiving throughput (yellow dashed line) aligns with the sampling throughput, indicating that no communi-

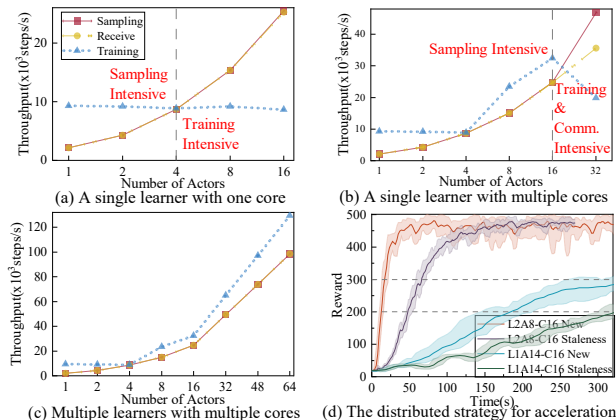


Figure 6: TianJi's distributed strategy achieves continuous scaling beyond performance bottlenecks and accelerate convergence by adjusting the sample distribution.

cation bottleneck is present at this stage. When the number of actors exceeds 4, the learner's training throughput becomes the bottleneck, and adding more actors does not accelerate performance. Figure 6(b) shows that increasing the number of physical cores for the learners and proportionally enlarging the batchsize can overcome the previous training bottleneck, shifting the equilibrium point to 16 actors. Further increasing the number of actors raises production throughput, but receiving and training throughput fall below this level. This indicates that both communication and training bottlenecks have been reached. Replicating the equilibrium configuration shown in Figure 6(b) and expanding it in groups, as depicted in Figure 6(c), shows that continuing to increase the number of actors results in a linear increase in sample receiving and training throughput. TianJi employs performance analysis and a distributed strategy based on production-consumption balance to dynamically allocate computation to resources, allowing for sustained expansion beyond performance bottlenecks.

Conclusion

We found that assignment dependencies are the key to ensuring the equivalence for paralleling DRL training, but also the bottleneck that hinders the performance. Therefore, we introduce TianJi, which reduces assignment dependencies between subtasks using a decentralized, data-driven training approach combined with event-driven asynchronous communication. Additionally, TianJi proposes a distributed strategy based on balancing sample production and consumption to alleviate the convergence issue introduced by relaxing the dependencies. Experimental results show that relaxing assignment dependencies and improving sample quality significantly enhance computational and training efficiency. TianJi achieves a 1.6-fold speedup in convergence time and a 7.13-fold speedup in training throughput compared to XingTian when scaled to eight machines. The ablation study demonstrates the effectiveness of key components. In future work, we will further deepen the theoretical proof of convergence under the relaxed assignment dependencies.

Acknowledgments

This work is sponsored in part by the National Natural Science Foundation of China under Grant No.62421002. We extend our gratitude to Peilin Lu and Ruihan Li for their significant contributions to the experiments.

References

- Assran, M.; Romoff, J.; Ballas, N.; Pineau, J.; and Rabbat, M. 2019. Gossip-based actor-learner architectures for deep reinforcement learning. *Advances in Neural Information Processing Systems*, 32.
- Berner, C.; Brockman, G.; Chan, B.; Cheung, V.; Debiak, P.; Dennison, C.; Farhi, D.; Fischer, Q.; Hashme, S.; Hesse, C.; et al. 2019. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*.
- Brittain, M. W.; Alvarez, L. E.; and Breeden, K. 2024. Improving autonomous separation assurance through distributed reinforcement learning with attention networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 22857–22863.
- Cassirer, A.; Barth-Maron, G.; Brevdo, E.; Ramos, S.; Boyd, T.; Sottiaux, T.; and Kroiss, M. 2021. Reverb: A Framework For Experience Replay. *arXiv:2102.04736*.
- Chadi, M.-A.; and Mousannif, H. 2023. Understanding reinforcement learning algorithms: The progress from basic Q-learning to proximal policy optimization. *arXiv preprint arXiv:2304.00026*.
- Chen, L.; Zhu, T.; Liu, J.; Liang, J.; and Xiao, Y. 2023. End-to-end entity linking with hierarchical reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 4173–4181.
- Dean, J.; Corrado, G.; Monga, R.; Chen, K.; Devin, M.; Mao, M.; Ranzato, M.; Senior, A.; Tucker, P.; Yang, K.; et al. 2012. Large scale distributed deep networks. *Advances in neural information processing systems*, 25.
- Espeholt, L.; Marinier, R.; Stanczyk, P.; Wang, K.; and Michalski, M. 2019. Seed rl: Scalable and efficient deep-rl with accelerated central inference. *arXiv preprint arXiv:1910.06591*.
- Espeholt, L.; Soyer, H.; Munos, R.; Simonyan, K.; Mnih, V.; Ward, T.; Doron, Y.; Firoyiu, V.; Harley, T.; Dunning, I.; et al. 2018. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures. In *International conference on machine learning*, 1407–1416. PMLR.
- Grondman, I.; Busoniu, L.; Lopes, G. A.; and Babuska, R. 2012. A survey of actor-critic reinforcement learning: Standard and natural policy gradients. *IEEE Transactions on Systems, Man, and Cybernetics, part C (applications and reviews)*, 42(6): 1291–1307.
- Horgan, D.; Quan, J.; Budden, D.; Barth-Maron, G.; Hessel, M.; Van Hasselt, H.; and Silver, D. 2018. Distributed prioritized experience replay. *arXiv preprint arXiv:1803.00933*.
- Isard, M.; Budi, M.; Yu, Y.; Birrell, A.; and Fetterly, D. 2007. Dryad: distributed data-parallel programs from sequential building blocks. In *Proceedings of the 2nd ACM SIGOPS/EuroSys European conference on computer systems 2007*, 59–72.
- Jiang, Y.; Yang, Y.; Lan, Z.; Zhan, G.; Li, S. E.; Sun, Q.; Ma, J.; Yu, T.; and Zhang, C. 2024. Rocket Landing Control with Random Annealing Jump Start Reinforcement Learning. *arXiv preprint arXiv:2407.15083*.
- Kakade, S. M. 2003. *On the sample complexity of reinforcement learning*. University of London, University College London (United Kingdom).
- Li, S. E. 2023. Deep reinforcement learning. In *Reinforcement learning for sequential decision and optimal control*, 365–402. Springer.
- Li, Y. 2017. Deep reinforcement learning: An overview. *arXiv preprint arXiv:1701.07274*.
- Liang, E.; Liaw, R.; Nishihara, R.; Moritz, P.; Fox, R.; Goldberg, K.; Gonzalez, J.; Jordan, M.; and Stoica, I. 2018. RLlib: Abstractions for distributed reinforcement learning. In *International conference on machine learning*, 3053–3062. PMLR.
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2015. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*.
- Liu, S.; Qiao, P.; Dou, Y.; and Li, R. 2021. Ddper: Decentralized Distributed Prioritized Experience Replay. In *2021 IEEE International Conference on Multimedia and Expo (ICME)*, 1–6.
- Liu, Y.; Li, Y.; Xu, X.; Dou, Y.; and Liu, D. 2022. Heterogeneous Skill Learning for Multi-agent Tasks. In *Neural Information Processing Systems*.
- Liu, Z.; Xu, X.; Qiao, P.; and Li, D. 2024. Acceleration for Deep Reinforcement Learning using Parallel and Distributed Computing: A Survey. *ACM Comput. Surv.*, 57(4).
- Mei, Z.; Fu, W.; Gao, J.; Wang, G.; Zhang, H.; and Wu, Y. 2023. SRL: Scaling Distributed Reinforcement Learning to Over Ten Thousand Cores. *arXiv preprint arXiv:2306.16688*.
- Mnih, V.; Badia, A. P.; Mirza, M.; Graves, A.; Lillicrap, T.; Harley, T.; Silver, D.; and Kavukcuoglu, K. 2016. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, 1928–1937. PMLR.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; and Riedmiller, M. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*.
- Nair, A.; Srinivasan, P.; Blackwell, S.; Alcicek, C.; Fearon, R.; De Maria, A.; Panneershelvam, V.; Suleyman, M.; Beattie, C.; Petersen, S.; et al. 2015. Massively parallel methods for deep reinforcement learning. *arXiv preprint arXiv:1507.04296*.
- Neumann, O.; and Gros, C. 2022. Scaling laws for a multi-agent reinforcement learning model. *arXiv preprint arXiv:2210.00849*.
- Pan, L.; Qian, J.; Xia, W.; Mao, H.; Yao, J.; Li, P.; and Xiao, Z. 2022. Optimizing communication in deep reinforcement learning with xingtian. In *Proceedings of the 23rd ACM/I-FIP International Middleware Conference*, 255–268.

- Petrenko, A.; Huang, Z.; Kumar, T.; Sukhatme, G.; and Koltun, V. 2020. Sample factory: Egocentric 3d control from pixels at 100000 fps with asynchronous reinforcement learning. In *International Conference on Machine Learning*, 7652–7662. PMLR.
- Qiao, P.; He, Z.; Li, R.; Jiang, J.; Dou, Y.; and Li, D. 2023. MLPs: Efficient Training of MiniGo on Large-scale Heterogeneous Computing System. In *2022 IEEE 28th International Conference on Parallel and Distributed Systems (ICPADS)*, 475–482.
- Recht, B.; Re, C.; Wright, S.; and Niu, F. 2011. Hogwild!: A lock-free approach to parallelizing stochastic gradient descent. *Advances in neural information processing systems*, 24.
- Schaul, T.; Quan, J.; Antonoglou, I.; and Silver, D. 2015. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*.
- Schulman, J.; Levine, S.; Abbeel, P.; Jordan, M.; and Moritz, P. 2015. Trust region policy optimization. In *International conference on machine learning*, 1889–1897. PMLR.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Shen, H.; Zhang, K.; Hong, M.; and Chen, T. 2023a. Towards Understanding Asynchronous Advantage Actor-Critic: Convergence and Linear Speedup. *IEEE Transactions on Signal Processing*, 71: 2579–2594.
- Shen, S.; Ma, C.; Li, C.; Liu, W.; Fu, Y.; Mei, S.; Liu, X.; and Wang, C. 2023b. RiskQ: Risk-sensitive Multi-Agent Reinforcement Learning Value Factorization. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems*, volume 36, 34791–34825. Curran Associates, Inc.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. 2016. Mastering the game of Go with deep neural networks and tree search. *nature*, 529(7587): 484–489.
- Sun, Z.; Qiao, P.; Dou, Y.; Li, Q.; and Li, R. 2023. PALA: Parallel Actor-Learner Architecture for Distributed Deep Reinforcement Learning. *Chinese Journal of Computers*, 46(2): 229–243.
- Sutton, R. S. 1988. Learning to predict by the methods of temporal differences. *Machine learning*, 3: 9–44.
- Sutton, R. S.; McAllester, D.; Singh, S.; and Mansour, Y. 1999. Policy gradient methods for reinforcement learning with function approximation. *Advances in neural information processing systems*, 12.
- Szot, A.; Schwarzer, M.; Agrawal, H.; Mazouze, B.; Metcalf, R.; Talbott, W.; Mackraz, N.; Hjelm, R. D.; and Toshev, A. T. 2023. Large language models as generalizable policies for embodied tasks. In *The Twelfth International Conference on Learning Representations*.
- Vinyals, O.; Babuschkin, I.; Czarnecki, W. M.; Mathieu, M.; Dudzik, A.; Chung, J.; Choi, D. H.; Powell, R.; Ewalds, T.; Georgiev, P.; et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *nature*, 575(7782): 350–354.
- Wang, C.; Durugkar, I.; Liebman, E.; and Stone, P. 2023. DM²: Decentralized Multi-Agent Reinforcement Learning via Distribution Matching. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 11699–11707.
- Weihs, L.; Salvador, J.; Kotar, K.; Jain, U.; Zeng, K.-H.; Mottaghi, R.; and Kembhavi, A. 2020. Allenact: A framework for embodied ai research. *arXiv preprint arXiv:2008.12760*.
- Wilcox, A.; Balakrishna, A.; Dedieu, J.; Benslimane, W.; Brown, D.; and Goldberg, K. 2022. Monte carlo augmented actor-critic for sparse reward deep reinforcement learning from suboptimal demonstrations. *Advances in neural information processing systems*, 35: 2254–2267.
- Zhao, L.; Dai, X.; Zhao, Z.; Xin, Y.; Hu, Y.; Qian, J.; Yao, J.; and Li, K. 2023. High-throughput Sampling, Communicating and Training for Reinforcement Learning Systems. In *2023 IEEE/ACM 31st International Symposium on Quality of Service (IWQoS)*, 1–10. IEEE.
- Zhu, H.; Zhao, B.; Chen, G.; Chen, W.; Chen, Y.; Shi, L.; Yang, Y.; Pietzuch, P.; and Chen, L. 2023. {MSRL}: Distributed Reinforcement Learning with Dataflow Fragments. In *2023 USENIX Annual Technical Conference (USENIX ATC 23)*, 977–993.