

MALib: A Parallel Framework for Population-based Multi-agent Reinforcement Learning

Ming Zhou^{1*}; Ziyu Wan¹, Hanjing Wang¹, Muning Wen¹, Runzhe Wu¹,
Ying Wen^{1†}, Yaodong Yang², Weinan Zhang¹ and Jun Wang²

¹Shanghai Jiao Tong University, ²University College London

Abstract

Population-based multi-agent reinforcement learning (PB-MARL) refers to the series of methods nested with reinforcement learning (RL) algorithms, which produces a self-generated sequence of tasks arising from the coupled population dynamics. By leveraging auto-curricula to induce a population of distinct emergent strategies, PB-MARL has achieved impressive success in tackling multi-agent tasks. Despite remarkable prior arts of distributed RL frameworks, PB-MARL poses new challenges for parallelizing the training frameworks due to the additional complexity of multiple nested workloads between sampling, training and evaluation involved with heterogeneous policy interactions. To solve these problems, we present *MALib*, a scalable and efficient computing framework for PB-MARL. Our framework is comprised of three key components: (1) a centralized task dispatching model, which supports the self-generated tasks and scalable training with heterogeneous policy combinations; (2) a programming architecture named *Actor-Evaluator-Learner*, which achieves high parallelism for both training and sampling, and meets the evaluation requirement of auto-curriculum learning; (3) a higher-level abstraction of MARL training paradigms, which enables efficient code reuse and flexible deployments on different distributed computing paradigms. Experiments on a series of complex tasks such as multi-agent Atari Games show that *MALib* achieves throughput higher than 40K FPS on a single machine with 32 CPU cores; 5× speedup than RLlib and at least 3× speedup than OpenSpiel in multi-agent training tasks. *MALib* is publicly available at <https://github.com/sjtu-marl/malib>.

1 Introduction

Training intelligent agents that can adapt to a diverse set of complex environments and agents has been a long-standing challenge. A feasible way to handle these tasks is multi-agent reinforcement learning (MARL) [56, 2], which has shown great potentials to solve multi-agent tasks such as real-time strategy games [51, 37], autonomous driving [61], traffic light control [53] and ride-hailing [60, 21]. In particular, the PB-MARL algorithms combine deep reinforcement learning (DRL) and dynamical population selection methodologies (e.g., game theory [10], evolutionary strategies [40]) to generate *auto-curricula* [54, 19]. In such a way, PB-MARL continually generates advanced intelligence and has achieved impressive successes in some non-trivial tasks, like Dota2 [35], StarCraftII [50] and Leduc Poker [26]. However, due to the intrinsic dynamics arising from multi-agent and population interactions, these algorithms have intricately nested structure and are extremely data-thirsty, requiring a flexible and scalable training framework to ground their effectiveness.

The deployment of PB-MARL shares some common procedures with conventional (MA)RL, but many challenges still remain, including the auto-curricula learning process and exponential compositional sampling. PB-MARL inherently comprises heterogeneous tasks, including simulation, policy inference, policy training and policy support expansion. All of these tasks are coupled to the underlying mutable policy combinations, which further complicates the PB-MARL.

*The first three authors are core developers. Ming contributed to the system architecture design and development; Ziyu contributed to the algorithm implementations; Hanjing contributed to the data server decoupling.

†Correspondence to: Ying Wen <ying.wen@sjtu.edu.cn>.

Figure 1 presents a typical case in PB-MARL in which the learning process of Double Oracle [27, 5] or Policy Space Response Oracle (PSRO) methods [31, 33] can be incorporated. For example, at each iteration of PSRO, the algorithm will generate some policy combinations and executes independent learning for each agent. Such a nested learning process comprises rollout, training, evaluation in sequence, and works circularly until the algorithm finds the estimated Nash Equilibrium. We note that these sub-tasks perform highly heterogeneous, i.e., the underlying policy combination is different from agent to agent. Furthermore, the evaluation stage involves tremendous simulations that cross fast-growing joint policy sets. Thus, we believe these requirements make distributed computing unavoidable for achieving high efficiency in executions.

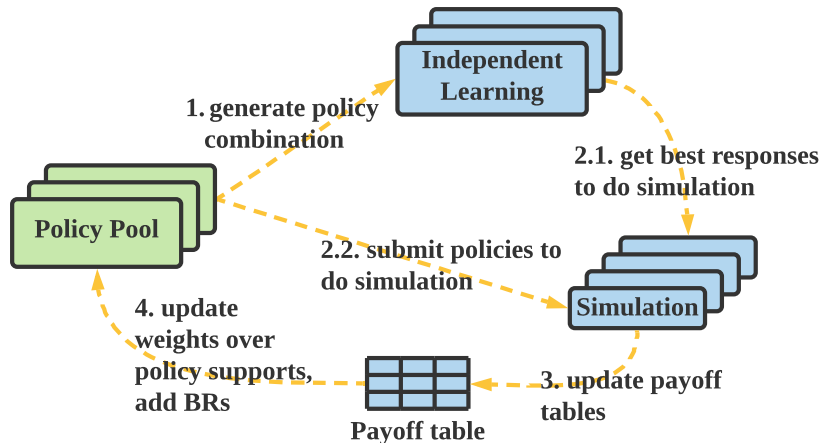


Figure 1: The diagram of PB-MARL learning process (e.g., Double Oracle/PSRO methods).

Most of existing training frameworks are originally designed for single-agent RL, and few attempts have been made to miscellaneous types of training schema in MARL, especially PB-MARL. In single-agent DRL, to meet the requirements of distributed computing, there have been tremendous distributed frameworks proposed to solve the data processing problem, e.g., RLlib [23], SEED RL [6], Sample-Factory [38] and Acme [11]. Despite that single-agent RL methods can also be applied in multi-agent settings [35] by acting as independent learners [47], more effective training schemas like centralized training & decentralized execution [25, 39], self-play [9, 20, 31], population-based learning, etc. require the training performs in a non-independent style. Obviously, it requires extensive effort for the single-agent learning mode to handle such interactive requirements, and poses new challenges to the existing distributed reinforcement learning frameworks. Though works like OpenSpiel [17] have attempted to support self-play and PSRO, they merely consider the good PB-MARL abstractions and scalability. Therefore, we claim that existing frameworks cannot fully satisfy the new requirements brought by PB-MARL.

In this paper, we stress the necessity and unfulfilled requirements for deploying PB-MARL and provide our systematic solution. We summarize that the essentials of a MARL distributed framework should (1) provide a highly efficient controlling mechanism to support self-supervised auto-curricula training and heterogeneous policy interactions in PB-MARL; (2) provide a high-level abstraction of MARL training schema to simplify the implementation and a unified and scalable training criteria. Based on the above analysis and comparisons, we present *MALib* as shown in Figure 2, a parallel framework that meets these requirements, to solve PB-MARL tasks in a high-parallelism style. We evaluate the performance of *MALib* from two distinct perspectives, i.e., system and algorithm. System performance is illustrated by the sample efficiency and throughput with increasing number of workers, while the algorithmic-side performance comparison is conducted over the reproduced algorithms with baselines, including typical PB-MARL algorithms and conventional MARL algorithms.

The main contributions of this work are summarized as follows:

1. We propose a centralized task dispatching model for PB-MARL, which achieves high

flexibility in training tasks over auto-curricula policy combinations.

2. We propose an independent programming model, *Actor-Evaluator-Learner*, to improve the efficiency of data processing and asynchronous execution.
3. To improve the code reuse and compatibility of heterogeneous MARL algorithms, we abstract the MARL training schema, which also bridges the gap between single-point and multi-point optimization.
4. We also provide mainstream MARL algorithm implementations and verify system performance and algorithm performance on *MALib* with different system settings, including single machine and clusters.

2 Related Work

A fundamental challenge in MARL is that the agents tend to overfit other players [18], making it hard for the algorithms to achieve robust performance. To solve this problem, interacting with heterogeneous agents or diverse policies of co-players is unavoidable. PB-MARL is a feasible approach to solve this problem, and prior works include population-based training [3, 14], fictitious self-play [50, 9] and meta-game analysis [31, 33].

Another difficulty is the data processing, the same as all DRL tasks. For deep reinforcement learning, high throughput allows algorithms to achieve a faster convergence rate and high data efficiency. There are many distributed reinforcement learning algorithms/frameworks proposed in recent years [32, 28, 45]. Among them, a standard implementation is to design training and rollout workers in fully distributed control, i.e., the Actor-Learner model. Also, some of them try to mitigate the communication loads between CPU and GPU [1] to improve the single-machine performance. Despite the impressive successes in distributed RL, most of them require users to do extra parallel programming to fit their custom requirements. RLlib [23] solved this problem by building a DRL framework on the top of Ray [30], work in a logically centralized control manner. Furthermore, the solution to MARL tasks among these frameworks is to model the MARL tasks as single-agent tasks, which decreases the computing efficiency in more general MARL settings since it requires heterogeneous agent/policy interaction in the training process.

Apart from the efforts in DRL, there are also tremendous works that integrate distributed computing techniques into deep learning architectures. Frameworks like Pytorch [36] and Horovod [44] implemented their distributed training logic over MPI [52]. General distributed tools like Ray [30] and Fiber [59] provide a universal API for building distributed applications, which relieve users from parallel programming such as MPI and OpenMP [4]. There are also some works that focus on MARL implementation and abstraction [42, 49], but most of them are too narrow to fit general MARL tasks, focusing on a specific domain [41]. As we claimed in the aforementioned content, the distributed computing support for MARL is necessary to the wider access of this exciting area.

Table 1: Comparison between *MALib* and existing distributed reinforcement learning frameworks from three dimensions.

Framework	Single-Agent	Multi-Agent	Population Management
RLlib [23]	✓	✓	×
SeedRL [6]	✓	×	×
Sample-Factory [38]	✓	✓	×
<i>MALib</i>	✓	✓	✓

To meet the distributed computing requirements of PB-MARL, we built our *MALib* on top of Ray and provided an efficient training framework. Table 1 presents the comparison between *MALib* and exiting distributed reinforcement learning framework from three dimensions, i.e., single-agent RL support, multi-agent RL support and population management. Despite some of them support multi-agent RL algorithms, they are essentially independent learning, or require

users’ extra efforts to implement algorithms in other training paradigms. Furthermore, a key dimension for PB-MARL is the population management, i.e., maintaining a policy pool for each agent, support policy expansion, update policy distribution in auto-curriculum learning, etc. *MALib* considered these requirements and gives corresponding implementations.

3 Parallel Programming Abstractions for PB-MARL

In this section, we will give an introduction to our framework from three components: the *Centralized Task Dispatching* model, the *Actor-Evaluator-Learner* model and the abstractions of MARL training paradigms, as shown in Figure 2. With these key implementations, we tense *MALib* serve for PB-MARL in auto-curriculum learning task schedule, execution in high performance and implementation with high code reuse.

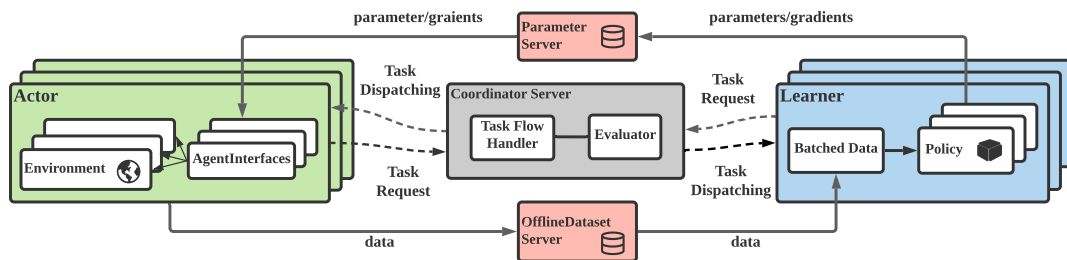


Figure 2: Overview of the *MALib* architecture. The *Coordinator Server* schedules the learning tasks; workers like *Actor* and *Learner* work in parallel and data dependencies are decoupled by *Parameter* and *OfflineDataset* servers. Actor is responsible for rollout/simulation tasks with k environments each, and Learner is responsible for the optimization of a policy pool. The collected experiences are processed before being sent to the *OfflineDataset* server. After Learner/Actor completes tasks, it will send a task request to *Coordinator* server for evaluation or promote the generation of next learning stage.

3.1 Centralized Task Dispatching Model

As introduced in Section 2, parallelizations for RL in previous work can be roughly classified into the Fully Distributed Control (FDC) [7, 6, 38] and the Hierarchical Parallel Task (HPT) model [23] fixed training task flow and policy interaction manners. Though these frameworks have abstractions for RL tasks, the extraordinary types of MARL training schema limit their performance, so users have to make extra efforts for customization. Furthermore, the PB-MARL algorithms like PSRO [31] and AlphaRank [34] require mutation in policy combination and policy space expansion in auto-curricula, which are ignored in previous frameworks.

We propose the *Centralized Task Dispatching (CTD)* model to meet these requirements in PB-MARL. Figure 3 presents the comparisons between previous parallel control model and our CTD model. This figure borrows the flowcharts from RLlib to better explain our design in parallel task control. The same as RLlib, we implemented the CTD model on top of Ray [30], which allows Python-implemented tasks to be naturally distributed over a large cluster.

The CTD model considers both of the advantages from FDC and HPT models. Specifically, the CTD has a centralized controller **D** to update the description of underlying policy combinations iteratively and generate new learning tasks, then dispatches them to working processes (**A**, **B** and **C**). The working processes in CTD work independently but do not coordinate with each other like in FTD. Furthermore, the working processes also work in a semi-passive manner, i.e., they will send task requests to **D** after completing tasks, which differs from the HPT model where the working process is fully passive. In fact, the semi-passive execution can be highly performant since the working processes will not handle the centralized controller all the time, so that **D** can work in highly parallelism to process more tasks to make sure the system run in high efficiency,




Figure 3: We abstract parallel processes as $\mathbf{A}\sim\mathbf{D}$ in this figure. Processes perform autonomous control in (a) *Fully Distributed Control Model (FDC)* and centralized control in (b) the *Hierarchical Parallel Task Model (HPT)* and (c) our *Centralized Task Dispatching Model (CTD)*. \mathbf{D} represents the centralized control process which is responsible for task dispatching. The working processes ($\mathbf{A}\sim\mathbf{C}$) in CTD execute in semi-passive manner, performing higher parallelism than the fully-passive execution in HPT.

especially for Python, which has a global interpreter lock. In our implementation, we modeled \mathbf{D} as the *Coordinator Server*, and working processes \mathbf{A} , \mathbf{B} and \mathbf{C} could be Actors, Learners and decoupled data servers.

Defining the Task Graph. The execution logic of the CTD model can be formulated as a closed-loop task graph for PB-MARL. In the beginning, we initialize a set of policy pools $\mathcal{P} = \{\mathcal{P}_i \mid i = 1, \dots, M\}$ for each agent, while some of them can share the same policy pool. Then, each agent $\{a_j \mid j = 1, \dots, N\}$ from the environment will choose one policy π_{a_j} from its belonging policy pool \mathcal{P}_i to form policy combinations. We formulate a policy combination at intermediate training stage t as $\text{Comb}_t = \prod_j^N \pi_{a_j}$. Based on the generated Comb_t , the coordinator will dispatch rollout tasks and training tasks to working processes. In general, the amount of tasks is determined by specific algorithms, e.g., PSRO generates N training tasks for N agents if the nested RL algorithm performs independent learning, or $m \leq N$ tasks for some centralized learning algorithms. Let $X(\theta_k)$ represent the collected data from rollout workers with policy parameters θ_k . Then for each rollout iteration, we have:

$$X(\theta_k) \sim P(s, a, s' \mid a \sim \text{Comb}_t(\theta_k)),$$

where θ_k represents the policy parameters at iteration k . For a given policy combination, a batch of collected data will be stored as $\mathcal{D} = \{X(\theta_k) \mid k = 1, \dots, h\}$, and two parallel evaluation tasks for rollout and training will be executed periodically as

$$\text{Eval}_{\text{rollout}} = f(X(\theta_k \leftarrow \theta')), \text{ and } \theta' = \underset{\theta}{\text{argmax}} \text{Eval}_{\text{train}} = f(X' \sim \mathcal{D} \mid \theta).$$

Until the global evaluator from the coordinator server reports staged stopping based on either one of them, then a new policy combination Comb_{t+1} will be produced with $\text{Comb}_{t+1} \sim G(\text{Eval}_{\text{rollout}}, \text{Eval}_{\text{train}})$, where G could be a specific evaluation function from PB-MARL algorithm like PSRO. Figure 4 shows the execution of circled task graph.

Decoupling the Task-Data Flow. The data flow mentioned here denotes the flow from data collecting to data sampling, also parameters push & pull. In general, a data flow is private to a policy combination, and the corresponding working processes perform in high efficiency as one-to-many (one learner to multiple actors). Although such mode has shown many advantages in prior distributed frameworks [7, 6, 38, 12], it limits the parallelism in the PB-MARL case since the data dependencies could be many-to-many here, i.e., each rollout task is corresponding to policies from multiple learning processes whose learning paces differ to each one. We decouple the data flow from task execution using *Parameter Server* [22] and *OfflineDataset Server*, i.e., *Parameter Server* for parameters synchronous between Actors and Learners, *OfflineDataset Server* for data saving and sampling, as shown in Figure 2.

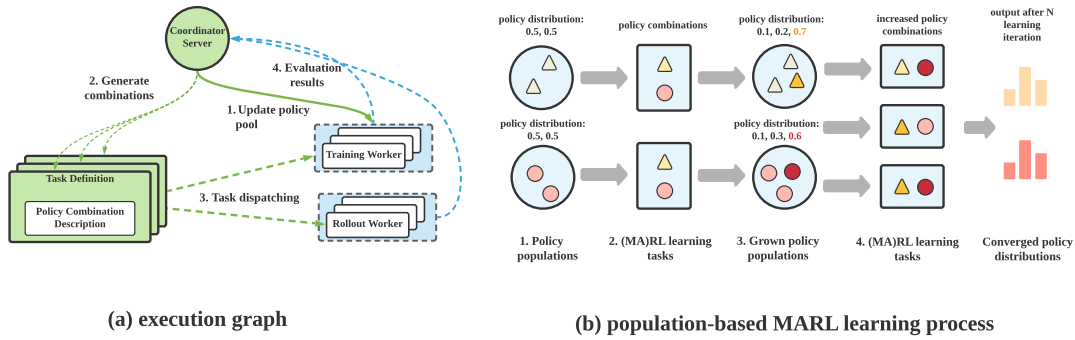


Figure 4: (a) The task execution graph of PB-MARL in *MALib*; (b) The auto-curriculum learning paradigm of PB-MARL. In general, a PB-MARL algorithm will start from an initialized policy set and the distribution over the policies (**step 1**), then find new policies (**step 2**) via (MA)RL algorithms. After that, the policy set will be updated with new policies, also the policy distribution (**step 3**), and follow with (MA)RL tasks over expanded policy combinations (**step 4**). This process will be cyclically executed until meeting the convergence such as an estimated Nash Equilibrium, then output the final policy distribution.

3.2 Actor-Evaluator-Learner Model

In single RL, it is common to decouple the training and rollout tasks [6, 7, 46, 38] using the *Actor-Learner* model, where the Learner operates policy training and the Actor operates data collecting. Furthermore, the evaluation program interleaves the Learners and Actors. However, in the case of PB-MARL, such a design appears to be inadequate since there is no centralized evaluator to integrate evaluations of multiple learning tasks which is required by PB-MARL tasks. Thus, we propose the *Actor-Evaluator-Learner* to meet this requirement, as shown in Figure 5. The *Evaluator* is nested with a payoff table to record the evaluation results of each policy combination. In general, the evaluation of a PB-MARL algorithm is built on top of the table, e.g., PSRO generates policy distribution over existing policies by estimating a Nash Equilibrium with this table.

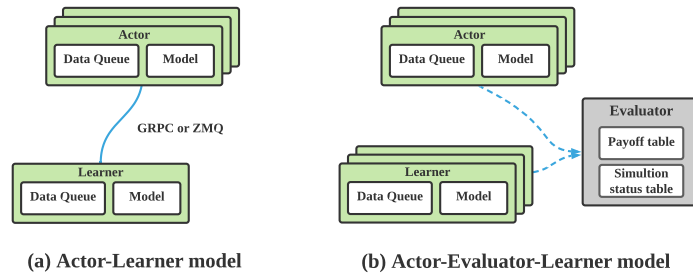


Figure 5: Comparison between Actor-Learner and Actor-Evaluator-Learner model.

3.3 Abstractions for MARL Training Paradigms

PB-MARL is nested with (MA)RL algorithms. In this section, we introduce a set of learners abstracted for different MARL paradigms. Our initial implementation offers five learners to meet the basic requirements, which also considers the asynchronous and synchronous training styles and some implementations of their respective MARL algorithms. Most previous works have tried to apply modular design to (MA)RL algorithm implementation and training. Still, they focus on the algorithm types that originate from RL, not the training paradigms, i.e., value-based

or policy-gradient-based learners [12]. Though such a modular design presents a completed implementation logic to users, it has low reuse because of the nested training logic in algorithm implementation. We list four typical training paradigms here and present the corresponding abstractions.

Table 2: *MALib*'s training abstraction supports various of (MA)RL and PB-MARL algorithms.

Algorithm Family	Independent	Centralized	Asynchronous
Value Based	DQN [29]	QMIX [39]	Gorrila [32]
Actor Critic	A2C [16]	MAAC [13]	IMPALA [7]
Population-based Training	PSRO [31]	×	Pipeline-PSRO [26]

Independent Learning. It is basically equivalent to single-agent deep reinforcement learning, with little or no need to change the original framework. The independent learner serves for independent learning algorithms like DQN [29] and PPO [43]. In this learning paradigm, one learner for one policy, policy learning is independent to other agents (policies).

Centralized Learning. In MARL, centralized learning means the learning requires shared information from other agents. In MADDPG [25], the shared information indicates the state-action pairs from all agents, and each agent has its individual critic, which accepts the shared information to perform policy optimization. Another variants are QMIX [39] and QDPP [57], which differ from MADDPG in that all agents share a global critic to do simultaneous policy optimization. Gradient shared algorithms like Networked-agent learning [58] could also be framed as centralized learning. Though its authors claimed that their method is fully decentralized, it is explained from the view of optimization.

Asynchronous Learning. Methods like IMPALA [7] use multiple Learners to update policy networks in asynchronous manner. We provide an implementation to support this paradigm. Furthermore, this learning interface can be coped with the former two as a nested implementation.

Synchronous Learning. Synchronous learning interface makes the Actors and Learner work in sequence, like the conventional *Actor-Learner* model, i.e., support multiple Actors work for one Learner. Furthermore, this learning interface can be coped with the Independent and Centralized learning interfaces as a nested implementation.

We demonstrate the exiting (MA)RL algorithm families within our abstractions in Table 2. The synchronous learning dimension is eliminated from the table since it is the default training manner for conventional (MA)RL algorithms. Furthermore, we show the aforementioned training paradigms in Figure 6.

4 Evaluation

Special efforts are paid to optimize the learning performance of multi-agent tasks in the design of *MALib*. In this section, metrics including data throughput, sampling efficiency and training time are reported for a comprehensive understanding of *MALib* system performance. All the experiment results listed are obtained with one of the following hardware settings: (1) *System# 1*: A 32-core computing node with dual graphics cards. (2) *System# 2*: A two-node cluster with each node owns 128-core and a single graphics card. All the GPUs mentioned are of the same model (NVIDIA RTX3090). We present the primary results here, and readers can find more details in Appendix B.




Figure 6: Depiction of scopes of independent/centralized and async/sync training diagrams. Independent and centralized paradigms have differences in data flow dependency: whether to share information among agents. And async and sync paradigms have differences in the task flow dependency: whether to strictly interleaving rollout and training.

4.1 Throughput Comparisons

We conduct the throughput evaluation on *MALib* and compare the results with some of the existing SOTA distributed RL frameworks, i.e. RLlib [23], Sample-Factory [38] and SEED RL [6]. Specially, Sample-Factory and SEED RL are highly tailored for TPU and GPU instances. Therefore, we repeat the group of experiments, with only CPU instances and GPU acceleration activated correspondingly.

Environment. As the environment for throughput comparison, we adopt the multi-agent version of Atari games (MA-Atari) from PettingZoo [48], which is a collection of 2D video games with more than one agent in each game. In our experiments, we use the two-players Pong game, with RGB-image frames in $12 \times 12 \times 3$ resolution.

Baselines and Settings. We choose IMPALA for SEED RL and RLlib, APPO for Sample-Factor and *MALib*. The evaluation of training throughput was conducted in different worker configurations over three minutes of continuous training, considering performance fluctuations caused by environment reset, data concatenation and other factors like threading lock. For each worker, we fixed the number of environments as 100. The number of workers ranges from 1 to 128 to compare the upper bound and bottleneck in parallelism performance of different frameworks.

Figure 7 shows the results of comparison on System# 1. Note that due to resource limitation, with only CPU cores, RLlib failed to launch with more than 32 workers while that threshold for GPU-accelerated RLlib is 8 workers on the same node. Despite of the extra abstraction layer introduced for tackling PB-MARL problems, the CPU version of *MALib* outperforms other frameworks in the conventional MA-Atari environment for the scalability to different size of worker groups. And *MALib* achieves comparable performance with Sample-Factory in the GPU-accelerated settings, which is a framework specially tailored for training conventional RL algorithms on single GPU node.

4.2 Algorithm Implementation and Performance

We have implemented a series of algorithms in *MALib* as listed in Appendix A, including independent learning algorithms like DQN, PPO and SAC, along with MARL algorithms like MADDPG and QMIX. Furthermore, all of these algorithms can be applied to Population-based Training (PBT), self-play, which have been supported by *MALib*. For the algorithmic performance evaluation, we focus on the convergence rate, which is derived from sample efficiency and training time consumption. As for the evaluated algorithms, we use PSRO training for PB-MARL, MADDPG and QMIX for conventional MARL algorithms. With the consideration of fairness, we make the algorithm implementation from different frameworks consistent in network settings. Appendix B presents the details.




Figure 7: Throughput comparison among the existing RL frameworks and MALib. Due to resource limitation(32 cores, 256G RAM), RLib fails under heavy loads(CPU case: #workers>32, GPU case: >8). MALib outperforms other frameworks with only CPU and achieves comparable performance with the highly tailored framework Sample-Factory with GPU despite higher abstraction introduced. To better illustrate the scalability of *MALib*, we show the MA-Atari and SC2 throughput on system#2 under different worker settings, the 512-workers group on SC2 fails due to resource limitation.




Figure 8: Comparisons of PSRO between *MALib* and OpenSpiel. (a) indicates that *MALib* achieves the same performance on exploitability as OpenSpiel; (b) shows that the convergence rate of *MALib* is $3\times$ faster than OpenSpiel; (c) shows that *MALib* achieves a higher execution efficiency than OpenSpiel, since it requires less time consumption to iterate the same learning steps, which means *MALib* has the potential to scale up in more complex tasks that need run for much more steps.

Table 3: *MALib*'s implementation result of population-based methods.

Algorithm	Time(sec)	Population Size
Fictitious Self-Play	7562	60
PSRO(fictitious play)	4862	40
PSRO(α -rank)	1776	21

Population-based training for Leduc Poker. We compared the PSRO algorithm with OpenSpiel's [17] implementation on Leduc Poker, a common benchmark in Poker AI. In order to stress how the evaluation effects the running time, we change the meta-solver to fictitious play, an approximation of Nash Equilibrium with tractable solving time on a large payoff matrix. To get a relatively accurate empirical payoff, the number of simulations was sat to be 2000 for each policy combination and the learning iteration was limited to 100 steps, i.e., the maximum of population size is 100.

We evaluate the convergence of PSRO with exploitability [26]. As shown in Figure 8b, we cut 70% execution time while maintaining exploitability with a similar quality as OpenSpiel(Figure 8). Moreover, it indicates that *MALib* is more capable of complicated games since the executing time of OpenSpiel grows much faster than *MALib*. Other methods like Fictitious Self-Play [9] (FSP) and Self-Play [] (SP) were implemented and evaluated in the same environment settings. In Table 3, we compare the execution time required and population size expanded when the exploitability goes to 0.5. The results show that PSRO outperforms the other two methods.

Specifically, SP fails to converge since this Poker game is not a purely transitive game, and FSP is more time-consumption than PSRO. It might be that PSRO considers the interactions and meta-game between different policies in populations, and solves it to approximate the Nash Equilibrium of the underlying game, which results in faster convergence rate. Furthermore, when an exact meta-game solver such as the LP-solver or the α -rank solver [34, 55], PSRO will converge in shorter time and smaller population size.

MADDPG for MPE. Multi-agent Particle Environments [25] (MPE) is a typical benchmarking environment for MARL algorithms. It offers plenty of scenarios covering cooperative and competitive tasks. We compared MADDPG with RLib implementation on seven scenarios under different worker settings, covered cooperative, competitive, and mixed cooperative-competitive tasks.

The experiments were conducted under different worker settings to compare the convergence performance. Figure 9 shows the results on *simple adversary*, it indicates that *MALib*'s implementation performs more steadily than RLib's, especially when the worker number increases, RLib implementation show high variance and fail to converge. We point out that the difference in implementations between *MALib* and RLib is that the former executes learning in a fully asynchronous and the RLib's implementation executes learning in sequential. Despite the data sampling executes in parallel, RLib requires extra efforts on tuning to solve the data starvation in the training stage.

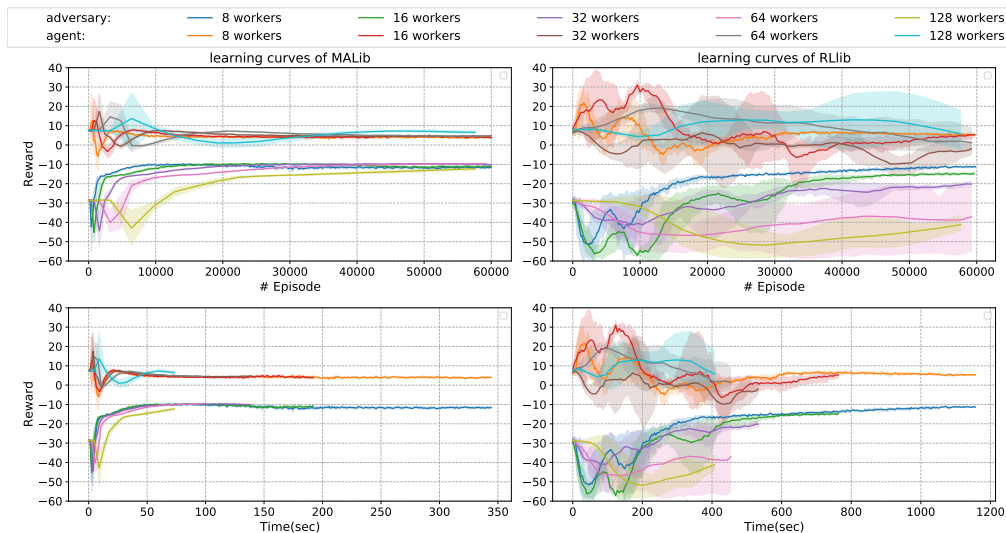


Figure 9: Comparisons of MADDPG in *simple adversary* under different rollout worker settings. Figures in the top row depict each agent's episode reward w.r.t. the number of sampled episodes, which indicate that *MALib* converges faster than RLib with equal sampled episodes. Figures in the bottom row show the average time and average episode reward at the same number of sampled episodes, which indicates that *MALib* achieves 5 \times speedup than RLib.

5 Conclusion

In this paper, we introduced *MALib* for PB-MARL to boost the research from high efficient training and code implementation, also compared with novel distributed RL frameworks in system and algorithm performance. Despite that *MALib* is currently built upon some of the low-level stacks from Ray [30] and lacks further optimization for GPU, it achieved higher throughput than previous work on multi-agent Atari tasks, especially in the CPU-only setting. For the algorithm performance, *MALib* shows at least 3x speedup with limited simulation parallelism on PB-MARL tasks compared to the existing library and achieved state-of-the-art scores on MPE tasks 5x faster than the parallel implementation in RLib. System development is a long-term work, we plan to

further improve MALib by optimizing the usage of the heterogeneous computing resources and testing the performance on much larger scale clusters in the future.

Acknowledgments

The SJTU team is supported by “New Generation of AI 2030” Major Project (2018AAA0100900), Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102), National Natural Science Foundation of China (62076161, 61632017) and Shanghai Sailing Program (21YF1421900). We also thank Zhicheng Zhang, Hangyu Wang, Ruiwen Zhou, Weizhe Chen, Minghuan Liu, Yunfeng Lin, Xihuai Wang, Derrick Goh and Linghui Meng for many helpful discussions, suggestions and comments on the project, and Ms. Yi Qu for her help on the design work.

References

- [1] Mohammad Babaeizadeh, Iuri Frosio, Stephen Tyree, Jason Clemons, and Jan Kautz. Reinforcement learning through asynchronous advantage actor-critic on a gpu. 2017.
- [2] Lucian Buşoniu, Robert Babuška, and Bart De Schutter. Multi-agent reinforcement learning: An overview. In *Innovations in multi-agent systems and applications-1*, pages 183–221. Springer, 2010.
- [3] Micah Carroll, Rohin Shah, Mark K Ho, Thomas L Griffiths, Sanjit A Seshia, Pieter Abbeel, and Anca Dragan. On the utility of learning about humans for human-ai coordination. 2019.
- [4] Barbara Chapman, Gabriele Jost, and Ruud Van Der Pas. Using openmp, 2007.
- [5] Le Cong Dinh, Yaodong Yang, Zheng Tian, Nicolas Perez Nieves, Oliver Slumbers, David Henry Mguni, Haitham Bou Ammar, and Jun Wang. Online double oracle. *arXiv preprint arXiv:2103.07780*, 2021.
- [6] Lasse Espeholt, Raphaël Marinier, Piotr Stanczyk, Ke Wang, and Marcin Michalski. Seed rl: Scalable and efficient deep-rl with accelerated central inference. 2020.
- [7] Lasse Espeholt, Hubert Soyer, Remi Munos, Karen Simonyan, Volodymyr Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, et al. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2018.
- [8] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning*, pages 1861–1870. PMLR, 2018.
- [9] Johannes Heinrich, Marc Lanctot, and David Silver. Fictitious self-play in extensive-form games. In *International conference on machine learning*, pages 805–813. PMLR, 2015.
- [10] Johannes Heinrich and David Silver. Deep reinforcement learning from self-play in imperfect-information games. *arXiv preprint arXiv:1603.01121*, 2016.
- [11] Matt Hoffman, Bobak Shahriari, John Aslanides, Gabriel Barth-Maron, Feryal Behbahani, Tamara Norman, Abbas Abdolmaleki, Albin Cassirer, Fan Yang, Kate Baumli, et al. Acme: A research framework for distributed reinforcement learning. *arXiv preprint arXiv:2006.00979*, 2020.
- [12] Matt Hoffman, Bobak Shahriari, John Aslanides, Gabriel Barth-Maron, Feryal Behbahani, Tamara Norman, Abbas Abdolmaleki, Albin Cassirer, Fan Yang, Kate Baumli, et al. Acme: A research framework for distributed reinforcement learning. *arXiv preprint arXiv:2006.00979*, 2020.

- [13] Shariq Iqbal and Fei Sha. Actor-attention-critic for multi-agent reinforcement learning. In *International Conference on Machine Learning*, pages 2961–2970. PMLR, 2019.
- [14] Max Jaderberg, Wojciech M Czarnecki, Iain Dunning, Luke Marris, Guy Lever, Antonio Garcia Castaneda, Charles Beattie, Neil C Rabinowitz, Ari S Morcos, Avraham Ruderman, et al. Human-level performance in 3d multiplayer games with population-based reinforcement learning. *Science*, 364(6443):859–865, 2019.
- [15] Max Jaderberg, Valentin Dalibard, Simon Osindero, Wojciech M Czarnecki, Jeff Donahue, Ali Razavi, Oriol Vinyals, Tim Green, Iain Dunning, Karen Simonyan, et al. Population based training of neural networks. *arXiv preprint arXiv:1711.09846*, 2017.
- [16] Vijay R Konda and John N Tsitsiklis. Actor-critic algorithms. In *Advances in neural information processing systems*, pages 1008–1014. Citeseer, 2000.
- [17] Marc Lanctot, Edward Lockhart, Jean-Baptiste Lespiau, Vinicius Zambaldi, Satyaki Upadhyay, Julien Pérolat, Sriram Srinivasan, Finbarr Timbers, Karl Tuyls, Shayegan Omidshafiei, et al. Openspiel: A framework for reinforcement learning in games. *arXiv preprint arXiv:1908.09453*, 2019.
- [18] Marc Lanctot, Vinicius Zambaldi, Audrunas Gruslys, Angeliki Lazaridou, Karl Tuyls, Julien Pérolat, David Silver, and Thore Graepel. A unified game-theoretic approach to multiagent reinforcement learning. 2017.
- [19] Joel Z Leibo, Edward Hughes, Marc Lanctot, and Thore Graepel. Autocurricula and the emergence of innovation from social interaction: A manifesto for multi-agent intelligence research. *arXiv preprint arXiv:1903.00742*, 2019.
- [20] David S Leslie and Edmund J Collins. Generalised weakened fictitious play. *Games and Economic Behavior*, 56(2):285–298, 2006.
- [21] Minne Li, Zhiwei Qin, Yan Jiao, Yaodong Yang, Jun Wang, Chenxi Wang, Guobin Wu, and Jieping Ye. Efficient ridesharing order dispatching with mean field multi-agent reinforcement learning. In *The World Wide Web Conference*, pages 983–994, 2019.
- [22] Mu Li, Li Zhou, Zichao Yang, Aaron Li, Fei Xia, David G Andersen, and Alexander Smola. Parameter server for distributed machine learning. In *Big Learning NIPS Workshop*, volume 6, page 2, 2013.
- [23] Eric Liang, Richard Liaw, Philipp Moritz, Robert Nishihara, Roy Fox, Ken Goldberg, Joseph E Gonzalez, Michael I Jordan, and Ion Stoica. Rllib: Abstractions for distributed reinforcement learning. 80:3059–3068, 2018.
- [24] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In Yoshua Bengio and Yann LeCun, editors, *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.
- [25] Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Advances in Neural Information Processing Systems*, pages 6379–6390, 2017.
- [26] Stephen McAleer, John Lanier, Roy Fox, and Pierre Baldi. Pipeline psro: A scalable approach for finding approximate nash equilibria in large games. 2020.
- [27] H Brendan McMahan, Geoffrey J Gordon, and Avrim Blum. Planning in the presence of cost functions controlled by an adversary. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pages 536–543, 2003.

- [28] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937. PMLR, 2016.
- [29] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- [30] Philipp Moritz, Robert Nishihara, Stephanie Wang, Alexey Tumanov, Richard Liaw, Eric Liang, Melih Elibol, Zongheng Yang, William Paul, Michael I Jordan, et al. Ray: A distributed framework for emerging {AI} applications. In *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*, pages 561–577, 2018.
- [31] Paul Muller, Shayegan Omidshafiei, Mark Rowland, Karl Tuyls, Julien Pérolat, Siqi Liu, Daniel Hennes, Luke Marris, Marc Lanctot, Edward Hughes, Zhe Wang, Guy Lever, Nicolas Heess, Thore Graepel, and Rémi Munos. A generalized training approach for multiagent learning. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.
- [32] Arun Nair, Praveen Srinivasan, Sam Blackwell, Cagdas Alcicek, Rory Fearon, Alessandro De Maria, Vedavyas Panneershelvam, Mustafa Suleyman, Charles Beattie, Stig Petersen, et al. Massively parallel methods for deep reinforcement learning. *arXiv preprint arXiv:1507.04296*, 2015.
- [33] Nicolas Perez Nieves, Yaodong Yang, Oliver Slumbers, David Henry Mguni, and Jun Wang. Modelling behavioural diversity for learning in open-ended games. *arXiv preprint arXiv:2103.07927*, 2021.
- [34] Shayegan Omidshafiei, Christos Papadimitriou, Georgios Piliouras, Karl Tuyls, Mark Rowland, Jean-Baptiste Lespiau, Wojciech M Czarnecki, Marc Lanctot, Julien Perolat, and Remi Munos. α -rank: Multi-agent evaluation by evolution. *Scientific reports*, 9(1):1–29, 2019.
- [35] OpenAI. Openai five. <https://blog.openai.com/openai-five/>, 2018.
- [36] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. pages 8024–8035, 2019.
- [37] Peng Peng, Ying Wen, Yaodong Yang, Quan Yuan, Zhenkun Tang, Haitao Long, and Jun Wang. Multiagent bidirectionally-coordinated nets: Emergence of human-level coordination in learning to play starcraft combat games. *arXiv preprint arXiv:1703.10069*, 2017.
- [38] Aleksei Petrenko, Zhehui Huang, Tushar Kumar, Gaurav Sukhatme, and Vladlen Koltun. Sample factory: Egocentric 3d control from pixels at 100000 fps with asynchronous reinforcement learning. In *International Conference on Machine Learning*, pages 7652–7662. PMLR, 2020.
- [39] Tabish Rashid, Mikayel Samvelyan, Christian Schröder de Witt, Gregory Farquhar, Jakob N. Foerster, and Shimon Whiteson. QMIX: monotonic value function factorisation for deep multi-agent reinforcement learning. In Jennifer G. Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 4292–4301. PMLR, 2018.
- [40] Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864*, 2017.
- [41] Mikayel Samvelyan, Tabish Rashid, Christian Schroeder de Witt, Gregory Farquhar, Nantas Nardelli, Tim G. J. Rudner, Chia-Man Hung, Philip H. S. Torr, Jakob Foerster, and Shimon Whiteson. The StarCraft Multi-Agent Challenge. pages 2186–2188, 2019.

- [42] Mikayel Samvelyan, Tabish Rashid, Christian Schroeder De Witt, Gregory Farquhar, Nantas Nardelli, Tim GJ Rudner, Chia-Man Hung, Philip HS Torr, Jakob Foerster, and Shimon Whiteson. The starcraft multi-agent challenge. pages 2186–2188, 2019.
- [43] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *CoRR*, abs/1707.06347, 2017.
- [44] Alexander Sergeev and Mike Del Balso. Horovod: fast and easy distributed deep learning in TensorFlow. *arXiv preprint arXiv:1802.05799*, 2018.
- [45] Adam Stooke and Pieter Abbeel. Accelerated methods for deep reinforcement learning. *arXiv preprint arXiv:1803.02811*, 2018.
- [46] Peng Sun, Jiechao Xiong, Lei Han, Xinghai Sun, Shuxing Li, Jiawei Xu, Meng Fang, and Zhengyou Zhang. Tleague: A framework for competitive self-play based distributed multi-agent reinforcement learning. *arXiv preprint arXiv:2011.12895*, 2020.
- [47] Ming Tan. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *Proceedings of the tenth international conference on machine learning*, pages 330–337, 1993.
- [48] Justin K Terry, Benjamin Black, Mario Jayakumar, Ananth Hari, Luis Santos, Clemens Diefendahl, Niall L Williams, Yashas Lokesh, Ryan Sullivan, Caroline Horsch, and Praveen Ravi. Pettingzoo: Gym for multi-agent reinforcement learning. *arXiv preprint arXiv:2009.14471*, 2020.
- [49] Yuandong Tian, Qucheng Gong, Wenling Shang, Yuxin Wu, and C Lawrence Zitnick. Elf: An extensive, lightweight and flexible research platform for real-time strategy games. pages 2659–2669, 2017.
- [50] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.
- [51] Oriol Vinyals, Timo Ewalds, Sergey Bartunov, Petko Georgiev, Alexander Sasha Vezhnevets, Michelle Yeo, Alireza Makhzani, Heinrich Küttler, John Agapiou, Julian Schrittwieser, et al. Starcraft ii: A new challenge for reinforcement learning. *arXiv preprint arXiv:1708.04782*, 2017.
- [52] David W Walker and Jack J Dongarra. Mpi: a standard message passing interface. *Supercomputer*, 12:56–68, 1996.
- [53] Cathy Wu, Aboudy Kreidieh, Eugene Vinitsky, and Alexandre M Bayen. Emergent behaviors in mixed-autonomy traffic. In *Conference on Robot Learning*, pages 398–407. PMLR, 2017.
- [54] Yaodong Yang, Jun Luo, Ying Wen, Oliver Slumbers, Daniel Graves, Haitham Bou Ammar, Jun Wang, and Matthew E Taylor. Diverse auto-curriculum is critical for successful real-world multiagent learning systems. *arXiv preprint arXiv:2102.07659*, 2021.
- [55] Yaodong Yang, Rasul Tutunov, Phu Sakulwongtana, and Haitham Bou Ammar. α -rank: Practically scaling α -rank through stochastic optimisation. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1575–1583, 2020.
- [56] Yaodong Yang and Jun Wang. An overview of multi-agent reinforcement learning from game theoretical perspective. *arXiv preprint arXiv:2011.00583*, 2020.
- [57] Yaodong Yang, Ying Wen, Jun Wang, Liheng Chen, Kun Shao, David Mguni, and Weinan Zhang. Multi-agent determinantal q-learning. In *International Conference on Machine Learning*, pages 10757–10766. PMLR, 2020.

- [58] Kaiqing Zhang, Zhuoran Yang, and Tamer Başar. Decentralized multi-agent reinforcement learning with networked agents: Recent advances. *arXiv preprint arXiv:1912.03821*, 2019.
- [59] Jiale Zhi, Rui Wang, Jeff Clune, and Kenneth O Stanley. Fiber: A platform for efficient development and distributed training for reinforcement learning and population-based methods. *arXiv preprint arXiv:2003.11164*, 2020.
- [60] Ming Zhou, Jiarui Jin, Weinan Zhang, Zhiwei Qin, Yan Jiao, Chenxi Wang, Guobin Wu, Yong Yu, and Jieping Ye. Multi-agent reinforcement learning for order-dispatching via order-vehicle distribution matching. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pages 2645–2653, 2019.
- [61] Ming Zhou, Jun Luo, Julian Villela, Yaodong Yang, David Rusu, Jiayu Miao, Weinan Zhang, Montgomery Alban, Iman Fadakar, Zheng Chen, et al. Smarts: Scalable multi-agent reinforcement learning training school for autonomous driving. *arXiv preprint arXiv:2010.09776*, 2020.

A Algorithm Library

We have integrated a set of popular (MA)RL algorithms. Table 4 gives an overview of these algorithms and tags them according to 1) training interface introduced in Section 3.3, 2) execution mode, and 3) the supported PB-MARL algorithms. The training interfaces could be **Independent** or **Centralized** which are corresponding to independent learning and centralized learning respectively. The execution mode could be Async (asynchronous) or Sync (synchronous). In the initial implementation, we provided three PB-MARL algorithms support, they are Policy Space Response Oracle [31] (PSRO), Fictitious Self-play [9] (FSP), Self-play [10] (SP) and Population-based Training [15] (PBT).

Table 4: Implemented algorithms in *MALib*.

Algorithm	Training Interface	Execution Mode	PB-MARL Support
DQN [29]	Independent	Async/Sync	PSRO/FSP/SP
Gorilla [32]	Independent	Async	PSRO/FSP/SP
A2C [16]	Independent	Sync	PSRO/FSP/SP
A3C [28]	Independent	Async	PSRO/FSP/SP
SAC [8]	Independent	Async/Sync	PSRO/FSP/SP
DDPG [24]	Independent	Async/Sync	PSRO/FSP/SP
PPO [43]	Independent	Sync	PSRO/FSP/SP
APPO	Independent	Async	PSRO/FSP/SP
MADDPG [25]	Centralized	Async/Sync	PBT
QMIX [39]	Centralized	Async/Sync	PBT
MAAC [13]	Centralized	Async/Sync	PBT

B Additional Results

B.1 MADDPG for MPE

We implemented MADDPG in *MALib* with the same configuration as RLLib, i.e., both of the actor and critic uses three layers of 64-units fully-connect network. The experiments were conducted in seven scenarios introduced in PettingZoo [48], with different worker settings (ranges from 1 to 128), as listed below.




Figure 10: Comparisons of MADDPG in *simple adversary* under different rollout worker settings.

Simple Adversary. There are 1 adversary, 2 good agents and 2 landmarks in this scenario. All agents can observe landmarks and other agents. One landmark is tagged as the ‘target landmark’. Good agents are rewarded based on the distance to the target landmark, i.e., closer to the target landmark higher reward, but also receive negative reward based on how close the adversary is to the target landmark. For the adversary, it is rewarded based on distance to the target, but it doesn’t know which landmark is the target landmark. In this scenario, good agents have to learn to ‘split up’ and cover all landmarks to deceive the adversary. Figure 10 shows the comparison from the converged reward and time-consumption.

Simple Crypto. There are 2 good agents (Alice and Bob) and 1 adversary (Eve) in this scenario. Alice must send a private 1 bit message to Bob over a public channel. Alice and Bob are rewarded +2 if Bob reconstructs the message, but are rewarded -2 if Eve reconstructs the message (that adds to 0 if both teams reconstruct the bit). Eve is rewarded -2 based if it cannot reconstruct the signal, zero if it can. Alice and Bob have a private key (randomly generated at beginning of each episode) which they must learn to use to encrypt the message. Figure 11 shows the comparison from the converged reward and time-consumption.

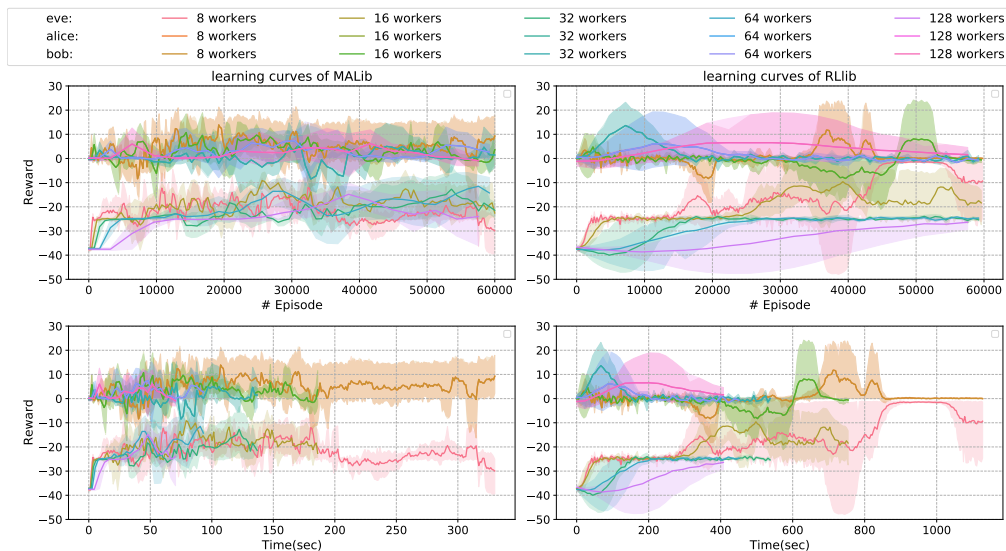


Figure 11: Comparisons between learning curves of MALib and RLLib when running MADDPG in *simple_crypto* environment under different rollout worker settings.

Simple Push. There are 1 good agent, 1 adversary, and 1 landmark in this scenario. The good agent is rewarded based on the distance to the landmark. The adversary is rewarded if it is close to the landmark, and if the agent is far from the landmark (the difference of the distances). Thus the adversary must learn to push the good agent away from the landmark. Figure 12 shows the comparison from the converged reward and time-consumption.

Simple Reference. There are 2 agents and 3 landmarks of different colors in this scenario. Each agent wants to get closer to their target landmark, which is known only by the other agents. Both agents are simultaneous speakers and listeners. Locally, the agents are rewarded by their distance to their target landmark. Globally, all agents are rewarded by the average distance of all the agents to their respective landmarks. The relative weight of these rewards is controlled by the `local_ratio` parameter. Figure 13 shows the comparison from the converged reward and time-consumption.

Simple Speaker Listener. This scenario is similar to `simple_reference`, except that one agent is the ‘speaker’ and can speak but cannot move, while the other agent is the listener (cannot speak,

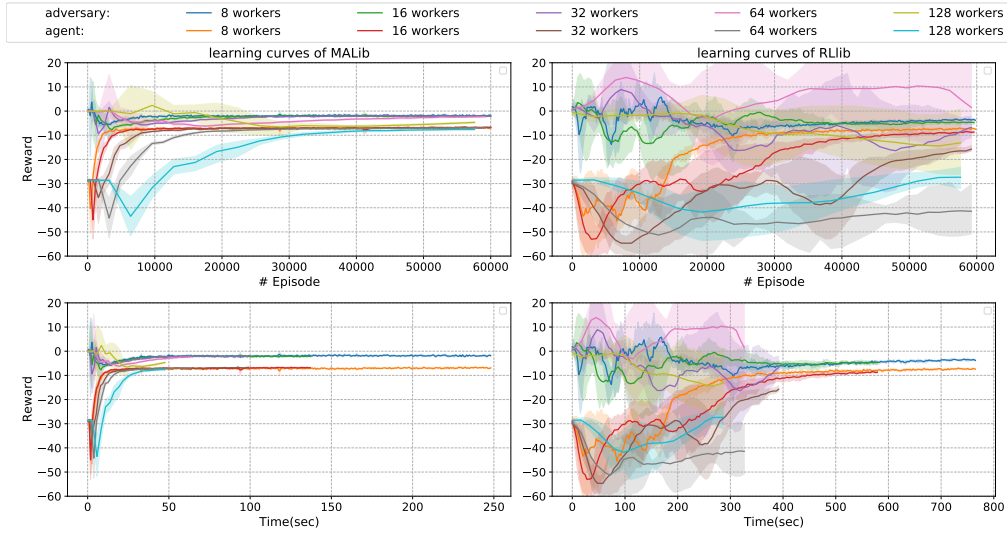


Figure 12: Comparisons between learning curves of *MALib* and *RLib* when running MADDPG in *simple push* under different rollout worker settings.

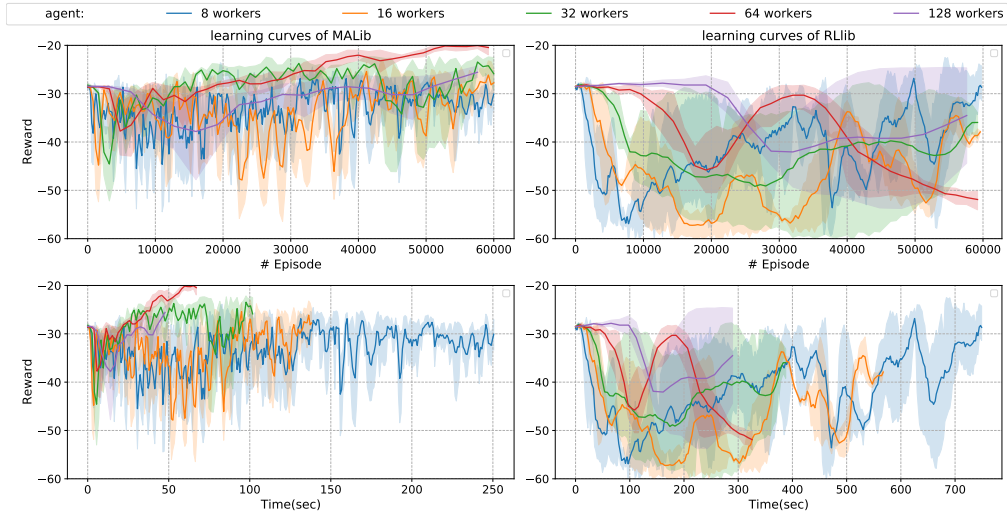


Figure 13: Comparisons between learning curves of *MALib* and *RLib* when running MADDPG in *simple reference* under different rollout worker settings.

but must navigate to correct landmark). Figure 14 shows the comparison from the converged reward and time-consumption.

Simple Spread. There are 3 agents, 3 landmarks in this scenario. Agents must learn to cover all the landmarks while avoiding collisions. More specifically, all agents are globally rewarded based on how far the closest agent is to each landmark (sum of the minimum distances). Locally, the agents are penalized if they collide with other agents (-1 for each collision). The relative weights of these rewards can be controlled with the `local_ratio` parameter. Figure 15 shows the comparison from the converged reward and time-consumption.

Simple Tag. There are 1 good agent, 3 adversaries and 2 obstacles in this scenario. Good agent is faster and receive a negative reward for being hit by adversaries (-10 for each collision). Adversaries are slower and are rewarded for hitting good agents (+10 for each collision). Obstacles block the way. Figure 16 shows the comparison from the converged reward and time-consumption.




Figure 14: Comparisons between learning curves of MALib and RLib when running MADDPG in *simple speaker listener* under different rollout worker settings.




Figure 15: Comparisons between learning curves of MALib and RLib when running MADDPG in *simple spread* under different rollout worker settings.

In summary, *MALib*'s implementation performs more steadily than RLib in different worker settings, and achieved faster convergence rate.

B.2 QMIX for SMAC

StarCraft Multi-Agent Challenge [51] (SMAC) is a multi-agent benchmark that provides elements of partial observability, challenging dynamics, and high-dimensional observation spaces. It is built on top of a real-time strategic game, StarCraftII, creating a testbed for research in cooperative MARL tasks. We compared the *MALib*'s and PyMARL's implementations of QMIX. Table 5 shows the scenarios we tested, and the time consumption (seconds) when win rate achieves 80%.

For the scenario *3s5z*, however, both of *MALib* and PyMARL cannot reach 80% win rate.

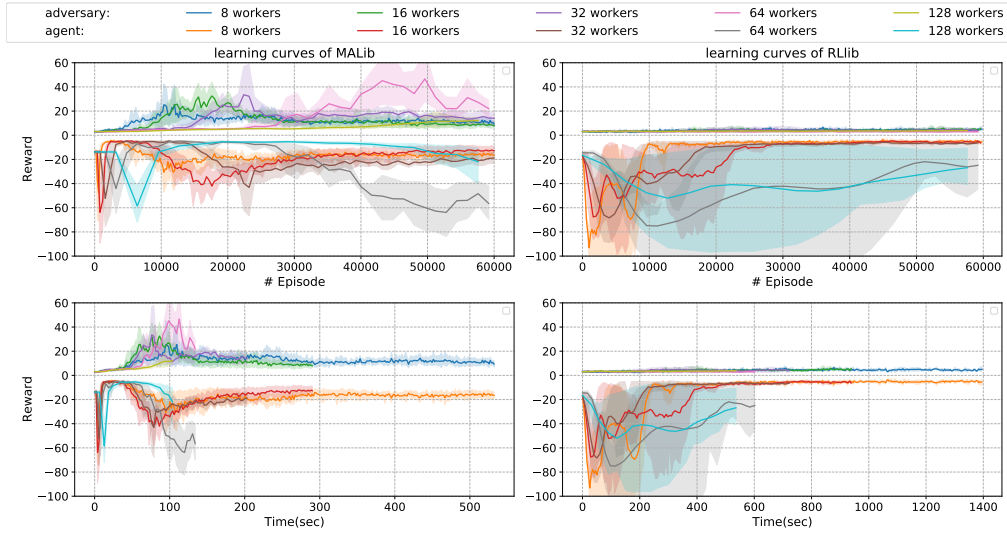


Figure 16: Comparisons between learning curves of MALib and RLib when running MADDPG in *simple tag* under different rollout worker settings.

Table 5: Tested SMAC scenarios and time consumption when Win Rate = 80%

Scenario	Ally Units	Enemy Units	MALib	PyMARL
3m	3 Marines	3 Marines	300	5625
8m	8 Marines	8 Marines	500	12375
2s3z	2 Stalkers & 3 Zealots	2 Stalkers & 3 Zealots	375	7920
3s5z	3 Stalkers & 5 Zealots	3 Stalkers & 5 Zealots	-	-




Figure 17: Additional comparisons between MALib and other distributed RL training frameworks. **(Left)**: System #3 cluster throughput of MALib in 2-player MA-Atari and 3-player SC2. **(Middle)**: 4-player MA-Atari throughput comparison on System #1 without GPU. **(Right)**: 4-player MA-Atari throughput comparison on System #1 with GPU.

B.3 Throughput and Comparisons

In addition to the system #1, #3 descriptions included in the main body, the data throughput of MALib on system #2 is depicted in Figure 17 (Left) to show the performance of different system presets. Each group of experiments are repeated for five times and the corresponding throughput results are reported. As is shown, the effects of random seeds and fluctuations among different trails of throughput tests are insignificant, and therefore error bars are omitted in the following figures for simplicity.

As illustrated in the Figure 17 (Middle, Right), we further evaluated and compared the training throughput with three other distributed RL frameworks in a unequally configured multi-agent MA-Atari game (i.e. 4-player Pong-v1).

Note that in the main text, we reported the results in a 2-player MA-Atari environment, which is less challenging than this case. And the FPS data obtained here has been re-aligned

with timestamps for a fair comparison among all the frameworks, causing a degradation in the reported performance for Sample-Factory.

C User Examples

C.1 Parallel Training Customization

MALib provides four basic training interfaces to support MARL training in different paradigms and distributed computing settings. We present an use case of centralized training customization to show the convenience of training customization. Specifically, the only thing we need to do is to customize a Trainer class which inherits from a standard training interface `Trainer`. Other training customization is similar the this case.

```
# ==== Centralized training customization ====
from malib.agents import CenntralizedAgent
from malib.algorithms.common import Trainer
from malib import runner

# ==== Customized trainer ====
class CentralizedTrainer(Trainer):
    # centralized trainer implemented the logic
    # of centralized training for arbitrary
    # independent RL algorithm like DDPG, A2C or PPO.
    def optimize(self, batch):
        pass

    def preprocess(self, batch, **kwargs):
        # merge sample batch or implement communication
        # proxy here
        pass

runner.run(
    # here, we let all agents share the same
    # centralized training interface
    agent_mapping_func=lambda agent: "share",
    # set
    training={
        "interface": {
            "type": CentralizedAgent,
            # ...
        }
    },
    # specify the trainable RL algorithm and trainer
    algorithms={
        "custom_name": {
            "PPO": {
                "trainer": CentralizedTrainer
                # ...
            }
        }
    },
    # ...
)
```

C.2 Single Instance Examples

Though *MALib* is designed to support distributed PB-MARL training tasks, it also supports single instance training via decoupled rollout and training interfaces. We present a single instance of PSRO training here.

Configuration Initialization. At the first, we need to initialize the training and rollout configurations, also the interface instances.

```
exp_cfg = {
    "expr_group": "single_instance_psro",
    "expr_name": f"simple_tag_{time.time()}",
}
env_config = {
    "num_good": 1,
    "num_adversaries": 1,
    "num_obstacles": 2,
    "max_cycles": 75,
}

env = simple_tag_v2.env(**env_config)
possible_agents = env.possible_agents
observation_spaces = env.observation_spaces
action_spaces = env.action_spaces

env_desc = {
    "creator": simple_tag_v2.env,
    "config": env_config,
    "id": "simple_tag_v2",
    "possible_agents": possible_agents,
}

# agent buffer, for sampling
agent_episodes = {
    agent: Episode(
        env_desc["id"],
        policy_id=None,
        capacity=args.buffer_size
    )
    for agent in env.possible_agents
}

rollout_config = {
    "stopper": "simple_rollout",
    "metric_type": args.rollout_metric,
    "fragment_length": 100,
    "num_episodes": 1,
    "terminate": "any",
    "mode": "on_policy",
    "callback": rollout_wrapper(agent_episodes), # online mode
}

algorithm = {
    "name": args.algorithm,
    "model_config": {},
    "custom_config": {}
}
```

```

}

learners = {}
for agent in env.possible_agents:
    learners[agent] = IndependentAgent(
        assign_id=agent,
        env_desc=env_desc,
        training_agent_mapping=None,
        algorithm_candidates={args.algorithm: algorithm},
        observation_spaces=observation_spaces,
        action_spaces=action_spaces,
        exp_cfg=exp_cfg,
    )
    learners[agent].register_env_agent(agent)

rollout_handlers = {
    agent: RolloutWorker(
        worker_index=None,
        env_desc=env_desc,
        metric_type=args.rollout_metric,
        remote=False,
        exp_cfg=exp_cfg,
    )
    for agent in env.possible_agents
}

# global evaluator to control the learning
psro_evaluator = PSROEvaluator()

# payoff manager, maintain agent payoffs and simulation status
payoff_manager = PayoffManager(
    env.possible_agents,
    exp_cfg=exp_cfg
)

```

Training Workflow. The training workflow is comprised of rollout and policy optimization. The implementation is listed as follows.

```

def training_workflow(
    trainable_policy_mapping: Dict[AgentID, PolicyID]
):
    # ...
    for agent in env.possible_agents:
        policy_distribution = payoff_manager.get_equilibrium(
            population)
        policy_distribution[agent] = {
            trainable_policy_mapping[agent]: 1.0}
        rollout_handlers[agent].ready_for_sample(
            policy_distribution)

    for epoch in range(args.num_epoch):
        # ==== collect training data ====
        rollout_feedback[agent], _ = \
            rollout_handlers[agent].sample(
                callback=rollout_config["callback"],
                num_episodes=[rollout_config["num_episodes"]],

```

```

        threaded=False,
        role="rollout",
        trainable_pairs=trainable_policy_mapping,
    )

    # ==== policy optimization ====
    batch = agent_episodes[agent].sample(args.batch_size)
    res = learners[agent].optimize(
        policy_ids={
            agent: trainable_policy_mapping[agent]
        },
        batch={agent: batch},
        training_config={
            "optimizer": "Adam",
            "critic_lr": 1e-4,
            "actor_lr": 1e-4,
            "lr": 1e-4,
            "update_interval": 5,
            "cliprange": 0.2,
            "entropy_coef": 0.001,
            "value_coef": 0.5,
        },
    )
    training_feedback.update(res)

return {
    "rollout": rollout_feedback,
    "training": training_feedback
}

```

Main Loop.

```

# init agent with fixed policy
policy_mapping = extend_policy_pool(trainable=True)
equilibrium: Dict[
    AgentID, Dict[PolicyID, float]
] = run_simulation_and_update_payoff(policy_mapping)

# === Main Loop ===#
iteration = 0
while True:
    # 1. add new trainable policy
    trainable_policy_mapping = extend_policy_pool(trainable=True)

    # 2. do rollout and training workflow
    feedback: Dict = training_workflow(trainable_policy_mapping)

    # 3. simulation and payoff table update
    equilibrium: Dict[
        AgentID, Dict[PolicyID, float]
    ] = run_simulation_and_update_payoff(trainable_policy_mapping)

    # 4. convergence judgement
    nash_payoffs: Dict[AgentID, float] = payoff_manager.aggregate(
        equilibrium=equilibrium
    )

```



```
weighted_payoffs: Dict[AgentID, float] = \  
    payoff_manager.aggregate(  
        equilibrium=equilibrium,  
        brs=trainable_policy_mapping,  
    )  
evaluation_results = psro_evaluator.evaluate(  
    None,  
    weighted_payoffs=weighted_payoffs,  
    oracle_payoffs=nash_payoffs,  
    trainable_mapping=trainable_policy_mapping,  
)  
  
if evaluation_results[EvaluateResult.CONVERGED]:  
    print("converged!")  
    break  
  
iteration += 1
```