
Gossip-based Actor-Learner Architectures for Deep Reinforcement Learning

Mahmoud Assran

Facebook AI Research &
Department of Electrical and Computer Engineering
McGill University
mahmoud.assran@mail.mcgill.ca

Joshua Romoff

Facebook AI Research &
Department of Computer Science
McGill University
joshua.romoff@mail.mcgill.ca

Nicolas Ballas

Facebook AI Research
ballasn@fb.com

Joelle Pineau

Facebook AI Research
jpineau@fb.com

Michael Rabbat

Facebook AI Research
mikerabbat@fb.com

Abstract

Multi-simulator training has contributed to the recent success of Deep Reinforcement Learning by stabilizing learning and allowing for higher training throughputs. We propose Gossip-based Actor-Learner Architectures (GALA) where several actor-learners (such as A2C agents) are organized in a peer-to-peer communication topology, and exchange information through asynchronous gossip in order to take advantage of a large number of distributed simulators. We prove that GALA agents remain within an ϵ -ball of one-another during training when using loosely coupled asynchronous communication. By reducing the amount of synchronization between agents, GALA is more computationally efficient and scalable compared to A2C, its fully-synchronous counterpart. GALA also outperforms A3C, being more robust and sample efficient. We show that we can run several loosely coupled GALA agents in parallel on a single GPU and achieve significantly higher hardware utilization and frame-rates than vanilla A2C at comparable power draws.

1 Introduction

Deep Reinforcement Learning (Deep RL) agents have reached superhuman performance in a few domains [Silver et al., 2016, 2018, Mnih et al., 2015, Vinyals et al., 2019], but this is typically at significant computational expense [Tian et al., 2019]. To both reduce running time and stabilize training, current approaches rely on distributed computation wherein data is sampled from many parallel simulators distributed over parallel devices [Espeholt et al., 2018, Mnih et al., 2016]. Despite the growing ubiquity of multi-simulator training, scaling Deep RL algorithms to a large number of simulators remains a challenging task.

On-policy approaches train a policy by using samples generated from that same policy, in which case data sampling (acting) is entangled with the training procedure (learning). To perform distributed training, these approaches usually introduce multiple learners with a *shared policy*, and multiple actors (each with its own simulator) associated to each learner. The shared policy can either be updated in a synchronous fashion (e.g., learners synchronize gradients before each optimization step [Stooke and Abbeel, 2018]), or in an asynchronous fashion [Mnih et al., 2016]. Both approaches have drawbacks: synchronous approaches suffer from straggler effects (bottlenecked by the slowest individual simulator), and therefore may not exhibit strong scaling efficiency; asynchronous methods are robust to stragglers, but prone to gradient staleness, and may become unstable with a large number of actors [Clemente et al., 2017].

Alternatively, off-policy approaches typically train a policy by sampling from a replay buffer of past transitions [Mnih et al., 2015]. Training off-policy allows for disentangling data-generation from learning, which can greatly increase computational efficiency when training with many parallel actors [Espeholt et al., 2018, Horgan et al., 2018, Kapturovski et al., 2019, Gruslys et al., 2017]. Generally, off-policy updates need to be handled with care as the sampled transitions may not conform to the current policy and consequently result in unstable training [Fujimoto et al., 2018].

We propose Gossip-based Actor-Learner Architectures (GALA), which aim to retain the robustness of synchronous on-policy approaches, while improving both their computational efficiency and scalability. GALA leverages multiple agents, where each agent is composed of one learner and possibly multiple actors/simulators. Unlike classical on-policy approaches, GALA does not require that each agent share the same policy, but rather it inherently enforces (through gossip) that each *agent’s policy remain ϵ -close to all others throughout training*. Relaxing this constraint allows us to reduce the synchronization needed between learners, thereby improving the algorithm’s computational efficiency.

Instead of computing an exact average between all the learners after a local optimization step, gossip-based approaches compute an approximate average using loosely coupled and possibly asynchronous communication (see Nedić et al. [2018] and references therein). While this approximation implicitly injects some noise in the aggregate parameters, we prove that this is in fact a principled approach as the learners’ policies stay within an ϵ -ball of one-another (even with non-linear function approximation), the size of which is directly proportional to the spectral-radius of the agent communication topology and their learning rates.

As a practical algorithm, we propose GALA-A2C, an algorithm that combines gossip with A2C agents. We compare our approach on six Atari games [Machado et al., 2017] following Stooke and Abbeel [2018] with vanilla A2C, A3C and the IMPALA off-policy method [Dhariwal et al., 2017, Mnih et al., 2016, Espeholt et al., 2018]. Our main empirical findings are:

1. Following the theory, GALA-A2C is empirically stable. Moreover, we observe that GALA can be more stable than A2C when using a large number of simulators, suggesting that the noise introduced by gossiping can have a beneficial effect.
2. GALA-A2C has similar sample efficiency to A2C and greatly improves its computational efficiency and scalability.
3. GALA-A2C achieves significantly higher hardware utilization and frame-rates than vanilla A2C at comparable power draws, when using a GPU.
4. GALA-A2C is competitive in term of performance relative to A3C and IMPALA.

Perhaps most remarkably, our empirical findings for GALA-A2C are obtained by simply using the default hyper-parameters from A2C. Our implementation of GALA-A2C is publicly available at <https://github.com/facebookresearch/gala>.

2 Technical Background

Reinforcement Learning. We consider the standard Reinforcement Learning setting [Sutton and Barto, 1998], where the agent’s objective is to maximize the expected value from each state $V(s) = \mathbb{E} [\sum_{i=0}^{\infty} \gamma^i r_{t+i} | s_t = s]$, γ is the discount factor which controls the bias towards nearby rewards. To maximize this quantity, the agent chooses at each discrete time step t an action a_t in the current state s_t based on its policy $\pi(a_t | s_t)$ and transitions to the next state s_{t+1} receiving reward r_t based on the environment dynamics.

Temporal difference (TD) learning [Sutton, 1984] aims at learning an approximation of the expected return parameterized by θ , i.e., the value function $V(s; \theta)$, by iteratively updating its parameters via gradient descent:

$$\nabla_{\theta} (G_t^N - V(s_t; \theta))^2 \tag{1}$$

where $G_t^N = \sum_{i=0}^{N-1} \gamma^i r_{t+i} + \gamma^N V(s_{t+n}; \theta_t)$ is the N -step return. Actor-critic methods [Sutton et al., 2000, Mnih et al., 2016] simultaneously learn both a parameterized policy $\pi(a_t | s_t; \omega)$ with parameters ω and a critic $V(s_t; \theta)$. They do so by training a value function via the TD error defined

in (1) and then proceed to optimize the policy using the policy gradient (PG) with the value function as a baseline:

$$\nabla_{\omega} (-\log \pi(a_t|s_t; \omega) A_t), \quad (2)$$

where $A_t = G_t^N - V(s_t; \theta_t)$ is the advantage function, which represents the relative value the current action has over the average. In order to both speed up training time and decorrelate observations, Mnih et al. [2016] collect samples and perform updates with several asynchronous actor-learners. Specifically, each worker $i \in \{1, 2, \dots, W\}$, where W is the number of parallel workers, collects samples according to its current version of the policy weights ω_i , and computes updates via the standard actor-critic gradient defined in (2), with an additional entropy penalty term that prevents premature convergence to deterministic policies:

$$\nabla_{\omega_i} \left(-\log \pi(a_t|s_t; \omega_i) A_t - \eta \sum_a \pi(a|s_t; \omega_i) \log \pi(a|s_t; \omega_i) \right). \quad (3)$$

The workers then perform HOGWILD! [Recht et al., 2011] style updates (asynchronous writes) to a shared set of master weights before synchronizing their weights with the master’s. More recently, Dhariwal et al. [2017] removed the asynchrony from A3C, referred to as A2C, by instead synchronously collecting transitions in parallel environments $i \in \{1, 2, \dots, W\}$ and then performing a large batched update:

$$\nabla_{\omega} \left[\frac{1}{W} \sum_{i=1}^W \left(-\log \pi(a_t^i|s_t^i; \omega) A_t^i - \eta \sum_a \pi(a|s_t^i; \omega) \log \pi(a|s_t^i; \omega) \right) \right]. \quad (4)$$

Gossip algorithms. Gossip algorithms are used to solve the distributed averaging problem. Suppose there are n agents connected in a peer-to-peer graph topology, each with parameter vector $x_i^{(0)} \in \mathbb{R}^d$. Let $\mathbf{X}^{(0)} \in \mathbb{R}^{n \times d}$ denote the row-wise concatenation of these vectors. The objective is to iteratively compute the average vector $\frac{1}{n} \sum_{i=1}^n x_i^{(0)}$ across all agents. Typical gossip iterations have the form $\mathbf{X}^{(k+1)} = \mathbf{P}^{(k)} \mathbf{X}^{(k)}$, where $\mathbf{P}^{(k)} \in \mathbb{R}^{n \times n}$ is referred to as the mixing matrix and defines the communication topology. This corresponds to the update $x_i^{(k+1)} = \sum_{j=1}^n p_{i,j}^{(k)} x_j^{(k)}$ for an agent v_i . At an iteration k , an agent v_i only needs to receive messages from other agents v_j for which $p_{i,j}^{(k)} \neq 0$, so sparser matrices $\mathbf{P}^{(k)}$ correspond to less communication and less synchronization between agents.

The mixing matrices $\mathbf{P}^{(k)}$ are designed to be row stochastic (each entry is greater than or equal to zero, and each row sums to 1) so that $\lim_{K \rightarrow \infty} \prod_{k=0}^K \mathbf{P}^{(k)} = \mathbf{1} \boldsymbol{\pi}^\top$, where $\boldsymbol{\pi}$ is the ergodic limit of the Markov chain defined by $\mathbf{P}^{(k)}$ and $\mathbf{1}$ is a vector with all entries equal to 1 [Seneta, 1981].¹ Consequently, the gossip iterations converge to a limit $\mathbf{X}^{(\infty)} = \mathbf{1} (\boldsymbol{\pi}^\top \mathbf{X}^{(0)})$; meaning the value at an agent i converges to $x_i^{(\infty)} = \sum_{j=1}^n \pi_j x_j^{(0)}$. In particular, if the matrices $\mathbf{P}^{(k)}$ are symmetric and doubly-stochastic (each row and each column must sum to 1), we obtain an algorithm such that $\pi_j = 1/n$ for all j , and therefore $x_i^{(\infty)} = 1/n \sum_{j=1}^n x_j^{(0)}$ converges to the average of the agents’ initial vectors.

For the particular case of GALA, we only require the matrices $\mathbf{P}^{(k)}$ to be row stochastic in order to show the ϵ -ball guarantees.

3 Gossip-based Actor-Learner Architectures

We consider the distributed RL setting where n agents (each composed of a single learner and several actors) collaborate to maximize the expected return $V(s)$. Each agent v_i has a parameterized policy network $\pi(a_t|s_t; \omega_i)$ and value function $V(s_t; \theta_i)$. Let $x_i = (\omega_i, \theta_i)$ denote agent v_i ’s complete set of trainable parameters. We consider the specific case where each v_i corresponds to a single A2C agent, and the agents are configured in a directed and peer-to-peer communication topology defined by the mixing matrix $\mathbf{P} \in \mathbb{R}^{n \times n}$.

In order to maximize the expected reward, each GALA-A2C agent alternates between one local policy-gradient and TD update, and one iteration of asynchronous gossip with its peers. Pseudocode is

¹Assuming that information from every agent eventually reaches all other agents

where $\beta < 1$ is related to the joint spectral radius of the graph sequence, α is the reference learning rate, $\tilde{\beta} := \beta^{-\frac{\tau+B}{\tau+B+1}}$, and $L := \sup_{s=1,2,\dots} \|\mathbf{G}^{(s)}\|$ denotes an upper bound on the magnitude of the local optimization updates during training.

Proposition 2 states that the agents’ parameters are guaranteed to reside within an ϵ -ball of their average at all iterations $k \geq \tau + B$. The size of this ball is proportional to the reference learning-rate, the spectral radius of the graph topology, and the upper bound on the magnitude of the local gradient updates. One may also be able to control the constant L in practice since Deep RL agents are typically trained with some form of gradient clipping.

4 Related work

Several recent works have approached scaling up RL by using parallel environments. Mnih et al. [2016] used parallel asynchronous agents to perform HOGWILD! [Recht et al., 2011] style updates to a shared set of parameters. Dhariwal et al. [2017] proposed A2C, which maintains the parallel data collection, but performs updates synchronously, and found this to be more stable empirically. While A3C was originally designed as a purely CPU-based method, Babaeizadeh et al. [2016] proposed GA3C, a GPU implementation of the algorithm. Stooke and Abbeel [2018] also scaled up various RL algorithms by using significantly larger batch sizes and distributing computation onto several GPUs. Differently from those works, we propose the use of *Gossip Algorithms* to aggregate information between different agents and thus simulators. Nair et al. [2015], Horgan et al. [2018], Espeholt et al. [2018], Kapturowski et al. [2019], Gruslys et al. [2017] use parallel environments as well, but disentangle the data collection (actors) from the network updates (learners). This provides several computational benefits, including better hardware utilization and reduced straggler effects. By disentangling acting from learning these algorithms must use off-policy methods to handle learning from data that is not directly generated from the current policy (e.g., slightly older policies).

Gossip-based approaches have been extensively studied in the control-systems literature as a way to aggregate information for distributed optimization algorithms [Nedić et al., 2018]. In particular, recent works have proposed to combine gossip algorithms with stochastic gradient descent in order to train Deep Neural Networks [Lian et al., 2017b,a, Assran et al., 2018], but unlike our work, focus only on the supervised classification paradigm.

5 Experiments

We evaluate GALA for training Deep RL agents on Atari-2600 games [Machado et al., 2017]. We focus on the same six games studied in Stooke and Abbeel [2018]. Unless otherwise-stated, all learning curves show averages over 10 random seeds with 95% confidence intervals shaded in. We follow the reproducibility checklist [Pineau, 2018], see Appendix A for details.

We compare A2C [Dhariwal et al., 2017], A3C [Mnih et al., 2016], IMPALA [Espeholt et al., 2018], and GALA-A2C. All methods are implemented in PyTorch [Paszke et al., 2017]. While A3C was originally proposed with CPU-based agents with 1-simulator per agent, Stooke and Abbeel [2018] propose a large-batch variant in which each agent manages 16-simulators and performs batched inference on a GPU. We found this large-batch variant to be more stable and computationally efficient (cf. Appendix C.1). We use the Stooke and Abbeel [2018] variant of A3C to provide a more competitive baseline. We parallelize A2C training via the canonical approach outlined in Stooke and Abbeel [2018], whereby individual A2C agents (running on potentially different devices), all average their gradients together before each update using the ALLREDUCE primitive.² For A2C and A3C we use the hyper-parameters suggested in Stooke and Abbeel [2018]. For IMPALA we use the hyper-parameters suggested in Espeholt et al. [2018]. For GALA-A2C we use the same hyper-parameters as the original (non-gossip-based) method. All GALA agents are configured in a directed ring graph. All implementation details are described in Appendix C. For the IMPALA baseline, we use a prerelease of TorchBeast [Küttler et al., 2019] available at <https://github.com/facebookresearch/torchbeast>.

²This is mathematically equivalent to a single A2C agent with multiple simulators (e.g., n agents, with b simulators each, are equivalent to a single agent with nb simulators).

Table 1: Across all training seeds we select the best final policy produced by each method at the end of training and evaluate it over 10 evaluation episodes (up to 30 no-ops at the start of the episode). Evaluation actions generated from $\arg \max_a \pi(a|s)$. The table depicts the mean and standard error across these 10 evaluation episodes.

	Steps	BeamRider	Breakout	Pong	Qbert	Seaquest	SpaceInvaders
IMPALA ¹	50M	8220	641	21	18902	1717	1727
IMPALA	40M	7118 ± 2536	127 ± 65	21 ± 0	7878 ± 2573	462 ± 2	4071 ± 393
A3C	40M	5674 ± 752	414 ± 56	21 ± 0	14923 ± 460	1840 ± 0	2232 ± 302
A2C	25M	8755 ± 811	419 ± 3	21 ± 0	16805 ± 172	1850 ± 5	2846 ± 22
A2C	40M	9829 ± 1355	495 ± 57	21 ± 0	19928 ± 99	1894 ± 6	3021 ± 36
GALA-A2C	25M	9500 ± 1020	690 ± 72	21 ± 0	18810 ± 37	1874 ± 4	2726 ± 189
GALA-A2C	40M	10188 ± 1316	690 ± 72	21 ± 0	20150 ± 28	1892 ± 6	3074 ± 69

¹ Espeholt et al. [2018] results using shallow network (identical to the network used in our experiments).

Convergence and stability: We begin by empirically studying the convergence and stability properties of A2C and GALA-A2C. Figure 1a depicts the percentage of successful runs (out of 10 trials) of standard policy-gradient A2C when we sweep the number of simulators across six different games. We define a run as successful if it achieves better than 50% of nominal 16-simulator A2C scores. When using A2C, we observe an identical trend across all games in which the number of successful runs decreases significantly as we increase the number of simulators. Note that the A2C batch size is proportional to the number of simulators, and when increasing the number of simulators we adjust the learning rate following the recommendation in Stooke and Abbeel [2018].


Figure 1a also depicts the percentage of successful runs when A2C agents communicate their parameters using gossip algorithms (GALA-A2C). In *every* simulator sweep across the six games (600 runs), the gossip-based architecture increases or maintains the percentage of successful runs relative to vanilla A2C, when using identical hyper-parameters. We hypothesize that exercising slightly different policies at each learner using gossip-algorithms can provide enough decorrelation in gradients to improve learning stability. We revisit this point later on (cf. Figure 3b). We note that Stooke and Abbeel [2018] find that stepping through a random number of uniform random actions at the start of training can partially mitigate this stability issue. We did not use this random start action mitigation in the reported experiments.

While Figure 1a shows that GALA can be used to stabilize multi-simulator A2C and increase the number of successful runs, it does not directly say anything about the final performance of the learned models. Figures 1b and 1c show the rewards plotted against the number of environment steps when training with 64 simulators. Using gossip-based architectures stabilizes and maintains the peak performance and sample efficiency of A2C across all six games (Figure 1b), and also increases the number of convergent runs (Figure 1c).


Figures 1d and 1e compare the wall-clock time convergence of GALA-A2C to vanilla A2C. Not only is GALA-A2C more stable than A2C, but it also runs at a higher frame-rate by mitigating straggler effects. In particular, since GALA-A2C learners do not need to synchronize their gradients, each learner is free to run at its own rate without being hampered by variance in peer stepping times.

Comparison with distributed Deep RL approaches: Figure 1 also compares GALA-A2C to state-of-the-art methods like IMPALA and A3C.³ In each game, the GALA-A2C learners exhibited good sample efficiency and computational efficiency, and achieved highly competitive final game scores. Next we evaluate the final policies produced by each method at the end of training. After training across 10 different seeds, we are left with 10 distinct policies per method. We select the best final policy and evaluate it over 10 evaluation episodes, with actions generated from $\arg \max_a \pi(a|s)$. In almost every single game, the GALA-A2C learners achieved the highest evaluation scores of any method. Notably, the GALA-A2C learners that were trained for 25M steps achieved (and in most cases surpassed) the scores for IMPALA learners trained for 50M steps [Espeholt et al., 2018].


³We report results for both the TorchBeast implementation of IMPALA, and from Table C.1 from Espeholt et al. [2018]




(a) Simulator sweep: Percentage of convergent runs out of 10 trials.




(b) Sample complexity: Best 3 runs for each method.




(c) Sample complexity: Average across 10 runs.




(d) Computational complexity: Best 3 runs for each method.



(e) Computational complexity: Average across 10 runs.



(f) Energy efficiency: Best 3 runs for each method.



(g) Energy efficiency: Average across 10 runs.

Figure 1: (a) GALA increases or maintains the percentage of convergent runs relative to A2C. (b)-(c) GALA maintains the best performance of A2C while being more robust. (d)-(e) GALA achieves competitive scores in each game and in the shortest amount of time. (f)-(g) GALA achieves competitive game scores while being energy efficient.




Figure 2: (a) The radius of the ϵ -ball within which the agents’ parameters reside during training. The theoretical upper bound in Proposition 1 is explicitly calculated and compared to the true empirical quantity. The bound in Proposition 1 is remarkably tight. (b) Average correlation between agents’ gradients during training (darker colors depict low correlation and lighter colors depict higher correlations). Neighbours in the GALA-A2C topology are annotated with the label “peer.” The GALA-A2C heatmap is generally much darker than the A2C heatmap, indicating that GALA-A2C agents produce more diverse gradients with significantly less correlation.




Figure 3: Comparing GALA-A2C hardware utilization to that of A2C when using one NVIDIA V100 GPU and 48 Intel CPUs. (a) Samples of instantaneous GPU utilization and power draw plotted against each other. Bubble sizes indicate frame-rates obtained by the corresponding algorithms; larger bubbles depict higher frame-rates. GALA-A2C achieves higher hardware utilization than A2C at comparable power draws. This translates to much higher frame-rates and increased energy efficiency. (b) Hardware utilization/energy efficiency vs. number of simulators. GALA-A2C benefits from increased parallelism and achieves a 10-fold improvement in GPU utilization over A2C.

Effects of gossip: To better understand the stabilizing effects of GALA, we evaluate the diversity in learner policies during training. Figure 2a shows the distance of the agents’ parameters from consensus throughout training. The theoretical upper bound in Proposition 1 is also explicitly calculated and plotted in Figure 2a. As expected, the learner policies remain within an ϵ -ball of one-another in weight-space, and this size of this ball is remarkably well predicted by Proposition 1.

Next, we measure the diversity in the agents’ gradients. We hypothesize that the ϵ -diversity in the policies predicted by Proposition 1, and empirically observed in Figure 2a, may lead to less correlation in the agents’ gradients. The categorical heatmap in Figure 2b shows the pair-wise cosine-similarity between agents’ gradients throughout training, computed after every 500 local environment steps, and averaged over the first 10M training steps. Dark colors depict low correlations and light colors depict high correlations. We observe that GALA-A2C agents exhibited less gradient correlations than A2C agents. Interestingly, we also observe that GALA-A2C agents’ gradients are more correlated with those of peers that they explicitly communicate with (graph neighbours), and less correlated with those of agents that they do not explicitly communicate with.

Computational performance: Figure 3 showcases the hardware utilization and energy efficiency of GALA-A2C compared to A2C as we increase the number of simulators. Specifically, Figure 3a shows that GALA-A2C achieves significantly higher hardware utilization than vanilla A2C at comparable power draws. This translates to much higher frame-rates and increased energy efficiency. Figure 3b shows that GALA-A2C is also better able to leverage increased parallelism and achieves a 10-fold improvement in GPU utilization over vanilla A2C. Once again, the improved hardware utilization and frame-rates translate to increased energy efficiency. In particular, GALA-A2C steps

through roughly 20 thousand more frames per Kilojoule than vanilla A2C. Figures 1f and 1g compare game scores as a function of energy utilization in Kilojoules. GALA-A2C is distinctly more energy efficient than the other methods, achieving higher game scores with less energy utilization.

6 Conclusion

We propose Gossip-based Actor-Learner Architectures (GALA) for accelerating Deep Reinforcement Learning by leveraging parallel actor-learners that exchange information through asynchronous gossip. We prove that the GALA agents’ policies are guaranteed to remain within an ϵ -ball during training, and verify this empirically as well. We evaluated our approach on six Atari games, and find that GALA-A2C improves the computational efficiency of A2C, while also providing extra stability and robustness by decorrelating gradients. GALA-A2C also achieves significantly higher hardware utilization than vanilla A2C at comparable power draws, and is competitive with state-of-the-art methods like A3C and IMPALA.

Acknowledgments

We would like to thank the authors of TorchBeast for providing their pytorch implementation of IMPALA.

References

- Mahmoud Assran. Asynchronous subgradient push: Fast, robust, and scalable multi-agent optimization. Master’s thesis, McGill University Libraries, 2018.
- Mahmoud Assran and Michael Rabbat. Asynchronous subgradient-push. *arXiv preprint arXiv:1803.08950*, 2018.
- Mahmoud Assran, Nicolas Loizou, Nicolas Ballas, and Michael Rabbat. Stochastic gradient push for distributed deep learning. *arXiv preprint arXiv:1811.10792*, 2018.
- Mohammad Babaeizadeh, Iuri Frosio, Stephen Tyree, Jason Clemons, and Jan Kautz. Ga3c: Gpu-based a3c for deep reinforcement learning. *CoRR abs/1611.06256*, 2016.
- Vincent D Blondel, Julien M Hendrickx, Alex Olshevsky, and John N Tsitsiklis. Convergence in multiagent coordination, consensus, and flocking. In *Proceedings of the 44th IEEE Conference on Decision and Control*, pages 2996–3000. IEEE, 2005.
- Alfredo V Clemente, Humberto N Castejón, and Arjun Chandra. Efficient parallel methods for deep reinforcement learning. *arXiv preprint arXiv:1705.04862*, 2017.
- Prafulla Dhariwal, Christopher Hesse, Oleg Klimov, Alex Nichol, Matthias Plappert, Alec Radford, John Schulman, Szymon Sidor, Yuhuai Wu, and Peter Zhokhov. Openai baselines, 2017.
- Lasse Espeholt, Hubert Soyer, Remi Munos, Karen Simonyan, Volodymir Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, et al. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures. *arXiv preprint arXiv:1802.01561*, 2018.
- Scott Fujimoto, Herke van Hoof, and David Meger. Addressing function approximation error in actor-critic methods. *arXiv preprint arXiv:1802.09477*, 2018.
- Audrunas Gruslys, Mohammad Gheshlaghi Azar, Marc G. Bellemare, and Rémi Munos. The reactor: A sample-efficient actor-critic architecture. *CoRR*, abs/1704.04651, 2017.
- Christoforos N Hadjicostis and Themistoklis Charalambous. Average consensus in the presence of delays in directed graph topologies. *IEEE Transactions on Automatic Control*, 59(3):763–768, 2013.
- Dan Horgan, John Quan, David Budden, Gabriel Barth-Maron, Matteo Hessel, Hado van Hasselt, and David Silver. Distributed prioritized experience replay. *CoRR*, abs/1803.00933, 2018.
- Steven Kapturowski, Georg Ostrovski, Will Dabney, John Quan, and Remi Munos. Recurrent experience replay in distributed reinforcement learning. In *International Conference on Learning Representations*, 2019.
- Heinrich Küttler, Nantas Nardelli, Thibaut Lavril, Marco Selvatici, Viswanath Sivakumar, Tim Rocktäschel, and Edward Grefenstette. Torchbeast: A pytorch platform for distributed rl. *arXiv preprint arXiv:1910.03552*, 2019.

- Xiangru Lian, Ce Zhang, Huan Zhang, Cho-Jui Hsieh, Wei Zhang, and Ji Liu. Can decentralized algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic gradient descent. In *Advances in Neural Information Processing Systems*, pages 5330–5340, 2017a.
- Xiangru Lian, Wei Zhang, Ce Zhang, and Ji Liu. Asynchronous decentralized parallel stochastic gradient descent. *arXiv preprint arXiv:1710.06952*, 2017b.
- Marlos C. Machado, Marc G. Bellemare, Erik Talvitie, Joel Veness, Matthew J. Hausknecht, and Michael Bowling. Revisiting the arcade learning environment: Evaluation protocols and open problems for general agents. *CoRR*, abs/1709.06009, 2017.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidfjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.
- 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, 2016.
- Arun Nair, Praveen Srinivasan, Sam Blackwell, Cagdas Alcicek, Rory Fearon, Alessandro De Maria, Vedavyas Panneershelvam, Mustafa Suleyman, Charles Beattie, Stig Petersen, Shane Legg, Volodymyr Mnih, Koray Kavukcuoglu, and David Silver. Massively parallel methods for deep reinforcement learning. *CoRR*, abs/1507.04296, 2015.
- Angelia Nedić, Alex Olshevsky, and Michael G Rabbat. Network topology and communication-computation tradeoffs in decentralized optimization. *Proceedings of the IEEE*, 106(5):953–976, 2018.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- Joelle Pineau. The machine learning reproducibility checklist (version 1.0), 2018.
- Benjamin Recht, Christopher Re, Stephen Wright, and Feng Niu. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *Advances in Neural Information Processing Systems 24*, pages 693–701. Curran Associates, Inc., 2011.
- Eugene Seneta. *Non-negative Matrices and Markov Chains*. Springer, 1981.
- David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484, 2016.
- David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dhharshan Kumaran, Thore Graepel, et al. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419): 1140–1144, 2018.
- Adam Stooke and Pieter Abbeel. Accelerated methods for deep reinforcement learning. *arXiv preprint arXiv:1803.02811*, 2018.
- Richard S Sutton and Andrew G Barto. *Introduction to reinforcement learning*, volume 135. MIT press Cambridge, 1998.
- Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pages 1057–1063, 2000.
- Richard Stuart Sutton. Temporal credit assignment in reinforcement learning. *Doctoral Dissertations Available from Proquest*, 1984.
- Yuangdong Tian, Jerry Ma, Qucheng Gong, Shubho Sengupta, Zhuoyuan Chen, James Pinkerton, and C Lawrence Zitnick. Elf opengo: An analysis and open reimplement of alphazero. *arXiv preprint arXiv:1902.04522*, 2019.
- John Tsitsiklis, Dimitri Bertsekas, and Michael Athans. Distributed asynchronous deterministic and stochastic gradient optimization algorithms. *IEEE transactions on automatic control*, 31(9): 803–812, 1986.
- O Vinyals, I Babuschkin, J Chung, M Mathieu, M Jaderberg, W Czarnecki, A Dudzik, A Huang, P Georgiev, R Powell, et al. Alphastar: Mastering the real-time strategy game starcraft ii, 2019.

Jacob Wolfowitz. Products of indecomposable, aperiodic, stochastic matrices. *Proceedings of the American Mathematical Society*, 14(5):733–737, 1963.

Appendices

A Reproducibility Checklist

We follow the reproducibility checklist Pineau [2018]: For all algorithms presented, check if you include:

- **A clear description of the algorithm.** See Algorithm 1
- **An analysis of the complexity (time, space, sample size) of the algorithm.** See Figures 1 and 3 for an analysis of sample efficiency, wall-clock time, and energy efficiency.
- **A link to a downloadable source code, including all dependencies.** See the attached zip file.

For any theoretical claim, check if you include:

- **A statement of the result.** See Propositions 1 and 2 in the main text.
- **A clear explanation of any assumptions.** See Appendix B for full details.
- **A complete proof of the claim.** See Appendix B for full details.

For all figures and tables that present empirical results, check if you include:

- **A complete description of the data collection process, including sample size.** We used the Arcade Learning Environment [Machado et al., 2017], specifically we used the gym package - see: github.com/openai/gym
- **A link to downloadable version of the dataset or simulation environment.** See: github.com/openai/gym
- **An explanation of how samples were allocated for training / validation / testing.** We didn't use explicit training / validation/ testing splits - but ran each algorithm with 10 different random seeds.
- **An explanation of any data that were excluded.** We only used 6 atari games due to time constraints - the same 6 games that were used in Stooke and Abbeel [2018].
- **The range of hyper-parameters considered, method to select the best hyper-parameter configuration, and specification of all hyper-parameters used to generate results.** We used standard hyper-parameters from Dhariwal et al. [2017], Stooke and Abbeel [2018], Espeholt et al. [2018].
- **The exact number of evaluation runs.** We used 10 seeds for the Atari experiments.
- **A description of how experiments were run.** See Appendix C for full details.
- **A clear definition of the specific measure or statistics used to report results.** 95% confidence intervals are used in all plots / tables unless otherwise stated.
- **Clearly defined error bars.** 95% confidence intervals are used in all plots / tables unless otherwise stated.
- **A description of results with central tendency (e.g. mean) and variation (e.g. stddev).** 95% confidence intervals are used in all plots / tables unless otherwise stated.
- **A description of the computing infrastructure used.** See Appendix C for full details.

B Proofs

Setting and Notation Before presenting the theoretical guarantees, we define some notation. Suppose we have n learners (e.g., actor-critic agents) configured in a peer-to-peer communication topology represented by a directed and potentially time-varying graph (the non-zero entries in the mixing matrix $P^{(k)}$ define the communication topology at each iteration k).

Learners constitute vertices in the graph, denoted by v_i for all $i \in [n]$, and edges constitute directed communication links. Let $\mathcal{N}_i^{\text{out}}$ denote agent v_i 's *out-peers*, the set of agents that v_i can send messages

to, and let $\mathcal{N}_i^{\text{in}}$ denote agent v_i 's *in-peers*, the set of agents that can send messages to v_i . If the graph is time-varying, these sets are annotated with time indices. Let $x_i \in \mathbb{R}^d$ denote the agent v_i 's complete set of trainable parameters, and let the training function $\mathcal{T}_i : \mathbb{R}^d \mapsto \mathbb{R}^d$ define agent v_i 's training dynamics (i.e., agent v_i optimizes its parameters by iteratively computing $x_i \leftarrow \mathcal{T}_i(x_i)$).

For each agent v_i we define send- and receive-buffers, \mathcal{B}_i and \mathcal{R}_i respectively, which are used by the underlying communication system (standard in the gossip literature [Tsitsiklis et al., 1986]). When an agent wishes to broadcast a message to its out-peers, it simply copies the message into its broadcast buffer. Similarly, when agent receives a message, it is automatically copied into the receive buffer. For convenience, we assume that each learner v_i can hold at most one message from in-peer v_j in its receive buffer, \mathcal{R}_i at any time k ; i.e., a newly received message from agent v_j overwrites the older one in the receive buffer.

Let $k \in \mathbb{N}$ denote the global iteration counter. That is, k increments whenever any agent (or subset of agents) completes one loop in Algorithm 1. Consequently, at each global iteration k , there is a set of agents \mathcal{I} that are activated, and within this set there is a (possibly non-empty) subset of agents $\mathcal{C} \subseteq \mathcal{I}$ that gossip in the same iteration. If a message from agent v_j is received by agent v_i at time k , let $\tau_{j,i}^{(k)}$ denote the time at which this message was sent. Let $\tau \geq \tau_{j,i}^{(k)}$ for all $i, j \in [n]$ and $k > 0$ denote an upper bound on the message delays. For analysis purposes, messages are sent with an effective delay such that they arrive right when the agent is ready to process the messages. That is, a message that is sent by agent v_j at iteration k' and processed by agent v_i at iteration k , where $k \geq k'$, is treated as having experienced a delay $\tau_{j,i}^{(k)} = k - k'$, even if the message actually arrives before iteration k and waits in the receive-buffer.

Let $\alpha g_i^{(k)} := \mathcal{T}_i(x_i^{(k)}) - x_i^{(k)}$ denote agent v_i 's local computation update at iteration k after scaling by some reference learning rate $\alpha > 0$, and define $g_i^{(k)} := 0$ if agent v_i is not active at iteration k . Algorithm 1 can thus be written as follows. If agent v_i does not gossip at iteration k , then its update is simply

$$x_i^{(k+1)} = x_i^{(k)} + \alpha g_i^{(k)}. \quad (5)$$

If agent v_i does gossip at iteration k , then its update is

$$x_i^{(k+1)} = \frac{1}{1 + |\mathcal{N}_i^{\text{in}}|} \left(x_i^{(k)} + \sum_{j \in \mathcal{N}_i^{\text{in}}} x_j^{\tau_{j,i}^{(k)}} + \alpha g_i^{(k)} \right), \quad (6)$$

where $x_j^{\tau_{j,i}^{(k)}}$ is the parameter value of the agent v_j , at the time where the message was sent, i.e., $\tau_{j,i}^{(k)}$.

We can analyze Algorithm 1 in matrix form by stacking all n agents' parameters, $x_i^{(k)} \in \mathbb{R}^d$, into a matrix $\mathbf{X}^{(k)}$, and equivalently stacking all of the update vectors, $g_i^{(k)} \in \mathbb{R}^d$, into a matrix $\mathbf{G}^{(k)}$. In order to represent the state of messages that are in transit (sent but not yet received), for analysis purposes, we augment the graph topology with virtual nodes using a standard graph augmentation [Hadjicostis and Charalambous, 2013] (we add τ virtual nodes for each non-virtual agent, where each virtual node stores a learner's parameters at a specific point within the last τ iterations). Let $\tilde{n} := n(\tau + 1)$ denote the cardinality of the augmented graph's vertex set. Equation (6) can be re-written as

$$\mathbf{X}^{(k+1)} = \tilde{\mathbf{P}}^{(k)} \left(\mathbf{X}^{(k)} + \alpha \mathbf{G}^{(k)} \right), \quad (7)$$

where $\mathbf{X}^{(k)}, \mathbf{G}^{(k)} \in \mathbb{R}^{\tilde{n} \times d}$, and the mixing matrix $\tilde{\mathbf{P}}^{(k)} \in \mathbb{R}^{\tilde{n} \times \tilde{n}}$ corresponding to the augmented graph is row-stochastic for all iterations k , i.e., all entries are non-negative, and all rows sum to 1. Mapping (6) to (7) may not be obvious, but is quite standard in the recent literature. We refer the interested reader to Assran and Rabbat [2018], Hadjicostis and Charalambous [2013].

Proof of Proposition 1. The proof is very similar to the proofs in Assran et al. [2018] and Assran and Rabbat [2018], and makes use of the graph augmentations in Hadjicostis and Charalambous [2013], the lower dimensional stochastic matrix dynamics in Blondel et al. [2005], and the ergodic matrix results in Wolfowitz [1963]. Since the matrices $\tilde{\mathbf{P}}^{(k)}$ are row-stochastic, their largest singular value is 1, which corresponds to singular vectors in $\text{span}\{1_{\tilde{n}}\}$. Let the matrix $\mathbf{Q} \in \mathbb{R}^{(\tilde{n}-1) \times \tilde{n}}$ define an orthogonal projection onto the space orthogonal to $\text{span}\{1_{\tilde{n}}\}$. Associated to each matrix

$\tilde{\mathbf{P}}^{(k)} \in \mathbb{R}^{\tilde{n} \times \tilde{n}}$ there is a unique matrix $\tilde{\mathbf{P}}'^{(k)} \in \mathbb{R}^{(\tilde{n}-1) \times (\tilde{n}-1)}$ such that $\mathbf{Q}\tilde{\mathbf{P}}^{(k)} = \tilde{\mathbf{P}}'^{(k)}\mathbf{Q}$. Let $\tilde{\mathbf{P}}'$ denote the collection of matrices $\tilde{\mathbf{P}}'^{(k)}$ for all k . The spectrum of the matrices $\tilde{\mathbf{P}}'^{(k)}$ is the spectrum of $\tilde{\mathbf{P}}^{(k)}$ after removing one multiplicity of the singular value 1. From (7), we have

$$\begin{aligned} \mathbf{Q}\mathbf{X}^{(k+1)} &= \mathbf{Q}\tilde{\mathbf{P}}^{(k)} \dots \tilde{\mathbf{P}}^{(1)}\tilde{\mathbf{P}}^{(0)}\mathbf{X}^{(0)} + \alpha \sum_{s=0}^k \mathbf{Q}\tilde{\mathbf{P}}^{(k)} \dots \tilde{\mathbf{P}}^{(s+1)}\tilde{\mathbf{P}}^{(s)}\mathbf{G}^{(s)} \\ &= \tilde{\mathbf{P}}'^{(k)} \dots \tilde{\mathbf{P}}'^{(1)}\tilde{\mathbf{P}}'^{(0)}\mathbf{Q}\mathbf{X}^{(0)} + \alpha \sum_{s=0}^k \tilde{\mathbf{P}}'^{(k)} \dots \tilde{\mathbf{P}}'^{(s+1)}\tilde{\mathbf{P}}'^{(s)}\mathbf{Q}\mathbf{G}^{(s)}. \end{aligned} \quad (8)$$

Note that $\mathbf{Q}(\mathbf{X}^{(k+1)} - \bar{\mathbf{X}}^{(k+1)}) = 0$ and $(\mathbf{X}^{(k+1)} - \bar{\mathbf{X}}^{(k+1)})^T \mathbf{1}_{\tilde{n}} = 0$. Thus

$$\begin{aligned} \left\| \mathbf{X}^{(k+1)} - \bar{\mathbf{X}}^{(k+1)} \right\| &= \left\| \mathbf{Q}(\mathbf{X}^{(k+1)} - \bar{\mathbf{X}}^{(k+1)}) \right\| \\ &\leq \left\| \tilde{\mathbf{P}}'^{(k)} \dots \tilde{\mathbf{P}}'^{(s+1)}\tilde{\mathbf{P}}'^{(0)}\mathbf{Q}\mathbf{X}^{(0)} \right\| + \alpha \sum_{s=0}^k \left\| \tilde{\mathbf{P}}'^{(k)} \dots \tilde{\mathbf{P}}'^{(s+1)}\tilde{\mathbf{P}}'^{(s)}\mathbf{Q}\mathbf{G}^{(s)} \right\|, \end{aligned}$$

where we have implicitly also made use of (8).

Defining $\beta := \sup_{s=0,1,\dots,k} \sigma_2(\tilde{\mathbf{P}}'^{(k)} \dots \tilde{\mathbf{P}}'^{(s+1)}\tilde{\mathbf{P}}'^{(s)})$, it follows that

$$\left\| \mathbf{X}^{(k+1)} - \bar{\mathbf{X}}^{(k+1)} \right\| \leq \beta^{k+1} \left\| \mathbf{Q}\mathbf{X}^{(0)} \right\| + \alpha \sum_{s=0}^k \beta^{k+1-s} \left\| \mathbf{Q}\mathbf{G}^{(s)} \right\|. \quad (9)$$

Assuming all learners are initialized with the same parameters, the first exponentially decay term on the right hand side of (9) vanishes and we have

$$\left\| \mathbf{X}^{k+1} - \bar{\mathbf{X}}^{(k+1)} \right\| \leq \alpha \sum_{s=0}^k \beta^{k+1-s} \left\| \mathbf{G}^{(s)} \right\|.$$

□

Proof of Proposition 2. The proof extends readily from Proposition 1. Given the assumptions on the graph sequence, the product of the matrices $\tilde{\mathbf{P}}^{(k)} \dots \tilde{\mathbf{P}}^{(s+1)}\tilde{\mathbf{P}}^{(s)}$ is ergodic for any $k - s \geq \tau + B$ (cf. [Assran, 2018]). Letting $\beta := \sup_{s=0,1,\dots,k} \sigma_2(\tilde{\mathbf{P}}'^{(k)} \dots \tilde{\mathbf{P}}'^{(s)})$, it follows from Wolfowitz [1963] and Blondel et al. [2005] that $\beta < 1$. □

C Implementation Details


C.1 A3C Implementation Comparison

While A3C was originally proposed with CPU-based agents with 1-simulator per agent, Stooke and Abbeel [2018] propose a variant in which each agent manages 16-simulators and performs batched inference on a GPU. Figure 4 compares 64-simulator learning curves using A3C as originally proposed in Mnih et al. [2016] to the large-batch variant in Stooke and Abbeel [2018]. The large-batch variant appears to be more robust and computationally efficient, therefore we use this GPU-based version of A3C in our main experiments to provide a more competitive baseline.


C.2 Experimental Setup

All experiments use the network suggested by Dhariwal et al. [2017]. Specifically, the network contains 3 convolutional layers and one hidden layer, followed by a linear output layer for the policy/linear output layer for the critic. The hyper-parameters for A2C, A3C and GALA-A2C are summarized in Table 2. IMPALA hyperparameters are the same as reported in Espeholt et al. [2018] (cf. table G.1 in their appendix).

In all GALA experiments we used 16 environments per learner, e.g., in the 64 simulator experiments in Section 5 we use 4 learners. GALA agents communicate using a 1-peer ring network. Figure 5 shows an example of such a ring network. The non-zero weight $p_{i,j}$ of the mixing matrix \mathbf{P} corresponding to the 1-peer ring are set to $\frac{1}{1+|\mathcal{N}_i^{\text{in}}|}$, which is equal to $1/2$ as $|\mathcal{N}_i^{\text{in}}| = 1$ for all i in the 1-peer ring graph.



(a) Sample complexity: Average across 10 runs.



(b) Computational complexity: Average across 10 runs.

Figure 4: Comparing large-batch A3C [Stooke and Abbeel, 2018] to original A3C [Mnih et al., 2016]. Running A3C with larger-batches provides more stable and sample efficient convergence (top row), while also maintaining computational efficiency by leveraging GPU acceleration (bottom row).

Hyper-parameter	Value
Image Width	84
Image Height	84
Grayscale	Yes
Action Repetitions	4
Max-pool over last k action repeat frames	2
Frame Stacking	4
End of episode when life lost	Yes
Reward Clipping	$[-1, 1]$
RMSProp momentum	0.0
RMSProp ϵ	0.01
Clip global gradient norm	0.5
Number of simulators per learner(B)	16
Base Learning Rate(α)	7×10^{-4}
Learning Rate Scaling	$\sqrt{\text{number of learners}}$ (not used in A3C)
VF coeff	0.5
Horizon(N)	5
Entropy Coeff.	0.01
Discount (γ)	.99
Max number of no-ops at the start of the episode	30

Table 2: Hyperparameters for both A2C, GALA-A2C, and A3C




Figure 5: Example of an n-agents/1-peer ring communication topology used in our experiment