



Scaling up stochastic gradient descent for non-convex optimisation

Saad Mohamad¹ · Hamad Alamri² · Abdelhamid Bouchachia¹

Received: 5 July 2020 / Revised: 12 April 2022 / Accepted: 7 July 2022
© The Author(s) 2022

Abstract

Stochastic gradient descent (SGD) is a widely adopted iterative method for optimizing differentiable objective functions. In this paper, we propose and discuss a novel approach to scale up SGD in applications involving non-convex functions and large datasets. We address the bottleneck problem arising when using both shared and distributed memory. Typically, the former is bounded by limited computation resources and bandwidth whereas the latter suffers from communication overheads. We propose a unified distributed and parallel implementation of SGD (named DPSGD) that relies on both asynchronous distribution and lock-free parallelism. By combining two strategies into a unified framework, DPSGD is able to strike a better trade-off between local computation and communication. The convergence properties of DPSGD are studied for non-convex problems such as those arising in statistical modelling and machine learning. Our theoretical analysis shows that DPSGD leads to speed-up with respect to the number of cores and number of workers while guaranteeing an asymptotic convergence rate of $O(1/\sqrt{T})$ given that the number of cores is bounded by $T^{1/4}$ and the number of workers is bounded by $T^{1/2}$ where T is the number of iterations. The potential gains that can be achieved by DPSGD are demonstrated empirically on a stochastic variational inference problem (Latent Dirichlet Allocation) and on a deep reinforcement learning (DRL) problem (advantage actor critic - A2C) resulting in two algorithms: DPSVI and HSA2C. Empirical results validate our theoretical findings. Comparative studies are conducted to show the performance of the proposed DPSGD against the state-of-the-art DRL algorithms.

Keywords Stochastic gradient descent · Large scale non-convex optimisation · Distributed and parallel computation · Variational inference · Deep reinforcement learning

Editors: João Gama, Alípio Jorge, Salvador García.

✉ Abdelhamid Bouchachia
abouchachia@bournemouth.ac.uk

¹ Department of Computing, Bournemouth University, Poole, UK

² WMG, Warwick University, Coventry, UK

1 Introduction

Stochastic gradient descent (SGD) is a general iterative algorithm for solving large-scale optimisation problems such as minimising a differentiable objective function $f(\mathbf{v})$ parameterised in $\mathbf{v} \in \mathcal{V}$:

$$\min_{\mathbf{v}} f(\mathbf{v}) \quad (1)$$

In several statistical models and machine learning problems, $f(\mathbf{v}) = \frac{1}{n} \sum_{i=1}^n f_i(\mathbf{v})$ is the empirical average loss where each $f_i(\mathbf{v})$ indicates that the function has been evaluated at a data instance \mathbf{x}_i . SGD updates \mathbf{v} using the gradients computed on some (or single) data points. In this work, we are interested in problems involving non-convex objective functions such as variational inference (Blei et al., 2017; Hoffman et al., 2013; Wainwright et al., 2008; Jordan et al., 1999) and artificial neural networks (Chilimbi et al., 2014; Dean et al., 2012; Li et al., 2014a; Xing et al., 2015; Zhang et al., 2015) amongst many others. Non-convex problems abound in ML and are often characterised by a very large number of parameters (e.g., deep neural nets), which hinders their optimisation. This challenge is often compounded by the sheer size of the training datasets, which can be in order of millions of data points. As the size of the available data increases, it becomes more essential to boost SGD scalability by distributing and parallelising its sequential computation. The need for scalable optimisation algorithms is shared across application domains.

Many studies have proposed to scale-up SGD by distributing the computation over different computing units taking advantage of the advances in hardware technology. The main existing paradigms exploit either *shared memory* or *distributed memory* architectures. While shared memory is usually used to run an algorithm in parallel on a single multi-core machine (Recht et al., 2011; Zhao et al., 2017; Lian et al., 2015; Huo & Huang, 2016), distributed memory, on the other hand, is used to distribute the algorithm on multiple machines (Agarwal & Duchi, 2011; Lian et al., 2015; Zinkevich et al., 2010; Langford et al., 2009). Distributed SGD (DSGD) is appropriate for very large-scale problems where data can be distributed over massive (theoretically unlimited) number of machines each with its own computational resources and I/O bandwidth. However DSGD efficiency is bounded by the communication latency across machines. Parallel SGD (PSGD) takes advantage of the multiple and fast processing units within a single machine with higher bandwidth communication; however, the computational resources and I/O bandwidth are limited.

In this paper, we set out to explore the potential gains that can be achieved by leveraging the advantages of both the distributed and parallel paradigms in a unified approach. The proposed algorithm, DPSGD, curbs the communication cost by updating a local copy of the parameter vector being optimised, \mathbf{v} , multiple times during which each machine performs parallel computation. The distributed computation of DPSGD among multiple machines is carried out in an *asynchronous* fashion whereby workers compute their local updates independently (Lian et al., 2015). A master aggregates these updates to amend the global parameters. The parallel implementation of DPSGD on each local machine is *lock-free* whereby multiple cores are allowed equal access to the shared memory to read and update the variables without locking (Zhao et al., 2017) (i.e., they can read and write the shared memory simultaneously). We provide a theoretical analysis of the convergence rate of DPSGD for non-convex optimisation problems and prove that linear speed-up with respect to the number of cores and workers is achievable while they are bounded by $T^{1/4}$ and $T^{1/2}$, respectively, where T is the total number of iterations. Furthermore, we empirically

validate these results by developing two inferential algorithms relying on DPSGD. The first one is an asynchronous lock-free stochastic variational inference algorithm (DPSVI) that can be deployed on a wide family of Bayesian models (see the Appendix); its potential is demonstrated here on a Latent Dirichlet Allocation problem. The second one is DPSGD-based Deep Reinforcement Learning (DRL) that can be used to scale up the training of DRL networks for multiple tasks (see the Appendix).

The rest of the paper is organised as follows. Section 2 presents the related work. Section 3 presents the proposed algorithm and the theoretical study. In Sect. 4, we carry out experiments and discuss the empirical results. Finally, Sect. 5 draws some conclusions and suggests future work. The appendices present the proofs, the asynchronous distributed lock-free parallel SVI and the highly-scalable actor critic algorithm.

2 Related work

We divide this section into two parts. In the first part, we discuss the relevant literature to distributed and parallel SGD focusing on the theoretical aspects (Lian et al., 2015; Zhao et al., 2017; Fang & Lin, 2017; Huo & Huang, 2017; Bottou, 2010; Niu et al., 2011; Tsitsiklis et al., 1986; Elgabri et al., 2020; Wang et al., 2019; Recht et al., 2011; Leblond et al., 2017; Dean et al., 2012; Zhou & Cong, 2017; Yu et al., 2019; Lin et al., 2018; Stich, 2018). To keep the paper centred and due to space limitation, only SGD-based methods for non-convex problems are covered. The second part covers related work focusing on the implementation/application of the distributed and parallel algorithms (Hoffman et al., 2013; Mohamad et al., 2018; Neiswanger et al., 2015; Dean et al., 2012; Paine et al., 2013; Ruder, 2016; Li et al., 2014a; Abadi et al., 2016; Paszke et al., 2017; Babaeizadeh et al., 2016; Mnih et al., 2016; Clemente et al., 2017; Horgan et al., 2018; Espeholt et al., 2018; Nair et al., 2015; Adamski et al., 2018).

2.1 Theoretical aspects

A handful of SGD-based methods have been proposed recently for large-scale non-convex optimisation problems (De Sa et al., 2015; Lian et al., 2015; Zhao et al., 2017; Fang & Lin, 2017; Huo & Huang, 2017), which embrace either a distributed or parallel paradigm. HOGWILD presented by Niu et al. (2011) proposes several asynchronous parallel SGD variants with locks and lock-free shared memory. Theoretical convergence analysis for convex objectives presented in that study was inspiring and adopted for most of the recent literature on asynchronous parallel optimisation algorithms. Similarly Leblond et al. (2017), De Sa et al. (2015) provided convergence analysis for SGD asynchronous lock-free parallel optimisation with shared memory for convex objectives where they provide convergence analysis with relaxed assumptions on the sparsity of the problem. De Sa et al. (2015) also analysed the HOGWILD convergence for non-convex objectives. Asynchronous distributed and lock-free parallel SGD algorithms for non-convex objectives have also been studied in Lian et al. (2015) showing that linear speed-up with respect to the number of workers is achievable when bounded by $O(\sqrt{T})$. Improved versions using variance reduction techniques have recently been proposed in Huo and Huang (2017), Fang and Lin (2017) to accelerate the convergence rate with a linear rate being achieved instead of the sub-linear one of SGD. Although these and other algorithmic implementations are

lock-free (Lian et al., 2015; Huo & Huang, 2017; Fang & Lin, 2017; De Sa et al., 2015), the theoretical analysis of the convergence was based on the assumption that no overwriting happens. Hence, write-lock or atomic operation for the memory are needed to prove the convergence. In contrast, Zhao et al. (2017) proposed a completely parallel lock-free implementation and analysis.

Different implementations exploiting both parallelism with shared memory and distributed computation across multiple machines have been proposed (Dean et al., 2012; Zhou & Cong, 2017; Yu et al., 2019; Lin et al., 2018; Stich, 2018). Except for Stich (2018), Dean et al. (2012), these methods adopted synchronous SGD implementation where Dean et al. (2012), Lin et al. (2018) focused on the implementation aspects, providing extensive empirical study on deep learning models. While the implementation ideas are very similar to ours, we consider lock-free local parallelism with asynchronous distribution and we provide theoretical analysis. We also evaluate our approach on different ML problems i.e., SVI and DRL (to be discussed in the next part). Instead of using a parameter server, the local learners in Zhou and Cong (2017) compute the average of their copies of parameters at regular intervals through global reduction. Communication overhead is controlled by introducing a communication interval parameter into the algorithm. However, the provided theoretical analysis in Zhou and Cong (2017) does not establish a speedup and synchronisation is required for global reduction. Authors in Yu et al. (2019) provide theoretical study for the model averaging introduced in Zhou and Cong (2017) showing that linear speedup of local SGD on non-convex objectives can be achieved as long the averaging interval is carefully controlled. Similar study for convex problems with asynchronous worker communication by Stich (2018) shows linear speedup in the number of workers and the mini-batch size with reduced communication. The multiple steps local SGD update by Stich (2018) aimed at reducing the communication overhead is similar to our proposed algorithm. Nonetheless, we adopt local lock-free parallelism and asynchronous distribution with parameter server scheme instead of model averaging. Finally, we point out that although we focus on SGD first-order method (Bottou, 2010; Niu et al., 2011; Tsitsiklis et al., 1986; Elgabli et al., 2020; Wang et al., 2019), our study can be extended to second-order methods (Shamir et al., 2014; Jahani et al., 2020a; Ba et al., 2016; Crane & Roosta, 2019; Jahani et al., 2020b) and variance reduction methods (Huo & Huang, 2017; Fang & Lin, 2017), where the high noise of our local multiple-steps update can be reduced contributing to further speed-up. We leave this for future work.

2.2 Implementation aspects

The first effort to scale-up variational inference is described in Hoffman et al. (2013) where gradient descent updates are replaced with SGD updates. Inspired by this work, Mohamad et al. (2018) replaces the SGD updates with asynchronous distributed SGD ones. This was achieved by computing the SVI stochastic gradient on each worker based on few (mini-batched or single) data points acquired from distributed sources. The update steps are then aggregated to form the global update. In Neiswanger et al. (2015), the strategy consists of distributing the entire dataset across workers and letting each one of them perform VI updates in parallel. This requires that, at each iteration, the workers must

be synchronised to combine their parameters. However, this synchronisation requirement limits the scalability so the maximum speed achievable is bounded by the slowest worker. Approaches for scaling up VI that rely on Bayesian filtering techniques have been reviewed in Mohamad et al. (2018).

Asynchronous SGD (ASYSG) (Lian et al., 2015), an implementation of SGD that distributes the training over multiple workers, has been adopted by DistBelief (Dean et al., 2012) (a parameter server-based algorithm for training neural networks) and Project Adam (Chilimbi et al., 2014) (another DL framework for training neural networks). Paine et al. (2013) showed that ASYSG can achieve noticeable speedups on small GPU clusters. Other similar work (Ruder, 2016; Li et al., 2014a) have also employed ASYSG to scale up deep neural networks. The two most popular and recent DL frameworks TensorFlow (Abadi et al., 2016) and Pytorch (Paszke et al., 2017) have embraced the Hogwild (Recht et al., 2011; Zhao et al., 2017) and ASYSG (Lian et al., 2015) implementations to scale-up DL problems.¹

Distributed and parallel SGD have also been employed in deep reinforcement learning (DRL) (Babaeizadeh et al., 2016; Mnih et al., 2016; Clemente et al., 2017; Horgan et al., 2018; Espeholt et al., 2018; Nair et al., 2015; Adamski et al., 2018). In Babaeizadeh et al. (2016), a hybrid CPU/GPU version of the Asynchronous Advantage ActorCritic (A3C) algorithm (Mnih et al., 2016) was introduced. The study focused on mitigating the severe under-utilisation of the GPU computational resources in DRL caused by its sequential nature of data generation. Unlike (Mnih et al., 2016), the agents in RL do not compute the gradients themselves. Instead, they send data to central learners that update the network on the GPU accordingly. However, as the number of core increases, the central GPU learner becomes unable to cope with the data. Furthermore, large amount of data requires large storage capacity. Also the internal communication overhead can affect the speed-up when the bandwidth reaches its ceiling. We note that a similar way for paralleling DRL is proposed by Clemente et al. (2017). Similarly, Horgan et al. (2018) propose to generate data in parallel using multi-cores CPU's where experiences are accumulated in a shared experience replay memory. Along the same trend, Espeholt et al. (2018) proposed to accumulate data by distributed actors and communicate it to the centralised learner where the computation is done. The architecture of these studies (Horgan et al., 2018; Espeholt et al., 2018) allows the distribution of the generation and selection of data instead of distributing locally computed gradients as in Nair et al. (2015). Hence, it requires sending large size information over the network in case of large size batch of data making the communication more demanding. Furthermore, the central learner has to perform most of the computation limiting the scalability.

The work in Adamski et al. (2018) is the most similar to ours, where SGD based hybrid distributed-parallel actor critic is studied. The parallel algorithm of Mnih et al. (2016) is combined with parameter server architecture of Nair et al. (2015) to allow parallel distributed implementation of A3C on a computer cluster with multi-core nodes. Each node applies the algorithm in Babaeizadeh et al. (2016) to queue data in batches, which are used to compute local gradients. These gradients are then gathered from all workers, averaged and applied to update the global parameters. To reduce the communication overhead,

¹ See, for instance, <https://github.com/pytorch/examples> and <https://github.com/tmulc18/Distributed-TensorFlow-Guide>

Table 1 List of symbols

Variable	Description
\boldsymbol{v}	The parameter vector to be optimised, called the <i>global parameter</i> . It is maintained and updated by the master machine
\boldsymbol{u}	Copy of the global parameter maintained and is updated by the workers
$\hat{\boldsymbol{u}}$	Copy of the local parameter \boldsymbol{u} and is stored in the worker's shared memory
$\ \boldsymbol{x}\ $	The Euclidean norm of vector \boldsymbol{x}
$f_i(\cdot)$	The objective function defined on the i^{th} instance
$\nabla f(\boldsymbol{v})$	The gradient vector of $f(\boldsymbol{v})$
t	The global unique iterate used in the synthetic sequence - not required by the algorithm
b	The local unique iterate used in the synthetic sequence - not required by the algorithm
m	Index referring to the update vector computed by the worker n_m and collected by the master
M	The master batch size which is the number of local updates coming from the workers for each global update by the master
p	The total number of threads (cores)
nW	The total number of workers (nodes)
S and P	Binary diagonal matrices used to denote whether local over-writing happens (see Sect. 3.2)

authors carried out careful reexamination of Adam optimiser's hyper-parameters allowing large-size batches to be used. Detailed discussion of these methods and comparison to our implementation is provided in the appendix.

3 The DPSGD algorithm and its properties

Before delving into the details of the proposed algorithm, we introduce the list of symbols in Table 1 that are used in the rest of the text.

3.1 Overview of the algorithm

The proposed DPSGD algorithms assumes a *star-shaped* computer network architecture: a master maintains the global parameter \boldsymbol{v} (Algorithm 1) and the other machines act as workers which independently and simultaneously compute local parameters \boldsymbol{u} (Algorithm 2). The workers communicate only with the master in order to access the state of the global parameter (line 3 in Algorithm 2) and provide the master with their local updates (computed based on local parameters) (line 10 in Algorithm 2). Each worker is assumed to be a multi-core machine, and the local parameter are obtained by running a lock-free parallel SGD (see Algorithm 2). This is achieved by allowing all cores equal access to the shared memory to read and update at any time with no restriction at all (Zhao et al., 2017). The master aggregates M predefined amounts of local updates coming from the workers (line 3 in Algorithm 1), and then computes its global parameter. The update step is performed as an atomic operation such that the workers are locked out and cannot read the global parameter during this step (see Algorithm 1).

Algorithm 1 DPSGD-Master: Updates performed at the master machine

-
- 1: **initialise:** number of iteration T , global variable \mathbf{v} , global learning rate $\{\rho_t\}_{t=0, \dots, T-1}$
 - 2: **for** $t = 0, 1, 2, \dots, T - 1$ **do**
 - 3: Collect M updating vectors $\mathbf{w}_1, \dots, \mathbf{w}_M$ from the workers.
 - 4: Update the current estimate of the global parameter $\mathbf{v} \leftarrow \mathbf{v} + \rho_t \sum_m \mathbf{w}_m$
 - 5: $t \leftarrow t + 1$
 - 6: **end for**
-

Algorithm 2 DPSGD-Worker: Updates performed at each worker machine

-
- 1: **initialise:** number of iterations of per-worker loop B , learning rate η , number of threads p
 - 2: **while** (MasterIsRun) **do**
 - 3: Pull a global parameter \mathbf{v} from the master and put it in the shared memory.
 - 4: Fork p threads
 - 5: **for** $c=0$ to $B-1$ **do**
 - 6: Read current values of \mathbf{u} , denoted as $\hat{\mathbf{u}}$, from the shared memory.
 - 7: Randomly pick i from $\{1, \dots, n\}$ and compute the gradient $\nabla f_i(\hat{\mathbf{u}})$.
 - 8: $\mathbf{u} \leftarrow \mathbf{u} - \eta \nabla f_i(\hat{\mathbf{u}})$
 - 9: **end for**
 - 10: Push the update vector $\mathbf{u} - \mathbf{v}$ from the shared memory to the master
 - 11: **end while**
-

Note that the local distributed computations are done in an asynchronous style, i.e., DPSGD does not lock the workers until the master starts updating the global parameter. That is, the workers might compute some of the stochastic gradients based on early values of the global parameter. Similarly, the lock-free parallel implies that local parameter can be updated by other cores in the time after being read and before being used for the update computation. Given this non-synchronisation among workers and among cores, the results of parameter update seem to be totally disordered, which makes the convergence analysis very difficult.

Following Zhao and Li (2016), we introduce a synthetic process to generate the final value of local and global parameters after all threads, workers have completed their updates as shown in Algorithm 1 and 2. That is, we generate a sequence of synthetic values of \mathbf{v} and \mathbf{u} with some order to get the final value of \mathbf{v} . These synthetic values are used for DPSGD convergence proof. The synthetic generation process is explained in the following section.

3.2 Synthetic process

Let t be the global unique iterate attached to the loop in Algorithm 1; b is the local unique iterate attached to the inner loop in Algorithm 2 and m is an index referring to the update vector computed by a worker $n_m \in \{1, \dots, nW\}$. If we omit the outer loop of

Algorithm 2, the key steps in Algorithm 2 are those related to the writing (updating) or reading the local parameter.

3.2.1 Local synthetic write (update) sequence

As in Zhao and Li (2016), we assume all threads will update the elements in \mathbf{u} in the order from 1 to \tilde{B} , where $\tilde{B} = B * p$ with p is the number of threads. Thus, $\{u_1, \dots, u_{\tilde{B}-1}\}$ is the synthetic sequence which may never occur in the shared memory. However, it is employed to obtain the final value $u_{\tilde{B}}$ after all threads have completed their updates in the inner-loop of Algorithm 2. In other terms, this ordered synthetic update sequence generating the same final value as that of the disordered lock-free update process. At iterate b , the synthetic update done by a thread can be written as follows:

$$\mathbf{u}_b = \mathbf{u}_0 - \sum_{j=0}^{b-1} \eta S_j \nabla f_{i_j}(\hat{\mathbf{u}}_j) \quad (2)$$

where S_j is a diagonal matrix whose entries are 0 or 1, determining which dimensions of the parameter vector u_b have been successfully updated by the j_{th} gradient computed on the shared local parameter $\hat{\mathbf{u}}_j$. That is $S_j(k, k) = 0$ if dimension k is over-written by another thread and $S_j(k, k) = 1$ if dimension k is successfully updated by $\nabla f_{i_j}(\hat{\mathbf{u}}_j)$ without over-writing. Equation 2 can be rearranged in an iterative form as:

$$\mathbf{u}_{b+1} = \mathbf{u}_b - \eta S_b \nabla f_{i_b}(\hat{\mathbf{u}}_b) \quad (3)$$

Including the outer loop and the global update in Algorithm 1, we define the synthetic sequence $\{\mathbf{u}_{t,m,b}\}$ equivalent to the updates for the b^{th} per-worker loop of the m^{th} update vector associated with the t^{th} master loop:

Algorithm 2, line 3 refers to: $\mathbf{u}_{t,m,0} = \mathbf{v}_{t-1}$

Algorithm 2, line 8 refers to: $\mathbf{u}_{t,m,b+1} = \mathbf{u}_{t,m,b} - \eta S_{i+\tau_{t,m},m,b} \nabla f_{i+\tau_{t,m},m,b}(\hat{\mathbf{u}}_{t,m,b})$ (4)

Algorithm 1, line 4 refers to: $\mathbf{v}_t = \mathbf{v}_{t-1} + \rho_{t-1} \left(\sum_{m=1}^M \mathbf{u}_{t-\tau_{t,m},m,\tilde{B}} - \mathbf{v}_{t-1-\tau_{t,m}} \right)$

where $\tau_{t,m}$ is the delay of the m^{th} global update for the t^{th} iteration caused by the asynchronous distribution. To compute $\nabla f_{i+\tau_{t,m},m,b}(\hat{\mathbf{u}}_{t,m,b})$, $\hat{\mathbf{u}}_{t,m,b}$ is read from the shared memory by a thread.

3.2.2 Local memory read

As denoted earlier, $\hat{\mathbf{u}}_b$ is the local parameter read from the shared memory which is used to compute $\nabla f_{i_b}(\hat{\mathbf{u}}_b)$ by a thread. Using the synthetic sequence $\{u_1, \dots, u_{\tilde{B}-1}\}$, $\hat{\mathbf{u}}_b$ can be written as:

$$\hat{\mathbf{u}}_b = \mathbf{u}_{a(b)} + \sum_{j=a(b)}^{b-1} P_{b,j-a(b)} \nabla f_{i_j}(\hat{\mathbf{u}}_j) \quad (5)$$

where $a(b) < b$ is the step in the inner-loop whose updates have been completely written in the shared memory. $P_{b,j-a(b)}$ are diagonal matrices whose diagonal entries are 0 or 1.

$\sum_{j=a(b)}^{b-1} P_{b,j-a(b)} \nabla f_{i_j}(\hat{u}_j)$ determine what dimensions of the new gradient updates, $\nabla f_{i_j}(\hat{u}_j)$, from time $a(b)$ to $b - 1$ have been added to $u_{a(b)}$ to obtain \hat{u}_b . That is, \hat{u}_b may read some dimensions of new gradients between time $a(b)$ to $b - 1$ including those which might have been over-written by some other threads. Including the outer loop and the global update in Algorithm 1, the local read becomes:

$$\hat{u}_{t,m,b} = u_{t,m,a(b)} - \eta \sum_{j=a(b)}^{b-1} P_{b,j-a(b)}^{t+\tau_{t,m},m} \nabla f_{i_{t+\tau_{t,m},m,j}}(\hat{u}_{t,m,j}) \tag{6}$$

The partial updates of the remaining steps between $a(b)$ and $b - 1$ are now defined by $\{P_{b,j-a(b)}^{t,m}\}_{a(b)}^{b-1}$.

3.3 Convergence analysis

Using the synthetic sequence, we develop the theoretical results of DPSGD showing that under some assumptions, we can guarantee linear speed-up with respect to the number of cores (threads) and number of nodes (workers). Before presenting the studies, we introduce and explain the require assumptions:

Assumption 1 The function $f(\cdot)$ is smooth, that is to say, the gradient of $f(\cdot)$ is *Lipschitzian*: there exists a constant $L > 0, \forall \mathbf{x}, \mathbf{y}$,

$$||\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})|| \leq L ||\mathbf{x} - \mathbf{y}||$$

or equivalently,

$$f(\mathbf{y}) \leq f(\mathbf{x}) + \nabla f(\mathbf{x})^T (\mathbf{y} - \mathbf{x}) + \frac{L}{2} ||\mathbf{y} - \mathbf{x}||^2.$$

Assumption 2 The per-dimension over-writing defined by $S_{t,m,b}$ is a random variate, independent of $i_{t,m,j}$.²

This assumption is reasonable since $S_{t,m,b}$ is affected by the hardware, while $i_{t,m,j}$ is independent thereof.

Assumption 3 The conditional expectation of the random matrix $S_{t,m,b}$ on $u_{t,m,b}$ and $\hat{u}_{t,m,b}$ is a strictly positive definite matrix, i.e., $\mathbb{E}[S_{t,m,b} | u_{t,m,b}, \hat{u}_{t,m,b}] = S > 0$ with the minimum eigenvalue $\alpha > 0$.

Assumption 4 The gradients are unbiased and bounded: $\nabla f(\mathbf{x}) = \mathbb{E}_i[\nabla f_i(\mathbf{x})]$ and $||\nabla f_i(\mathbf{x})|| \leq V, \forall i \in \{1, \dots, n\}$.

Then, it follows that the variance of the stochastic gradient is bounded. $\mathbb{E}_i[||\nabla f_i(\mathbf{x}) - \nabla f(\mathbf{x})||^2] \leq \sigma^2, \forall \mathbf{x}$, where $\sigma^2 = V^2 - ||\nabla f(\mathbf{x})||^2$

² Note that i can be a set of indices for a per-worker mini-batch. In this paper, i refers to a single index for simplicity.

Assumption 5 Delays between old local stochastic gradients and the new ones in the shared memory are bounded: $0 \leq b - a(b) \leq D$ and the delays between stale distributed update vectors and the current ones are bounded $0 \leq \max_{t,m} \tau_{t,m} \leq D'$

Assumption 6 All random variables in $\{i_{t,m,j}\}_{\forall t, \forall m, \forall j}$ are independent of each other.

Note that we are aware that this independence assumption is not fully accurate due to the potential dependency between selected data samples for computing gradients at the same shared parameters. For example, samples with fast computation of gradients for the same shared variable leads to more frequent selection of these samples as they likely to finish their gradient computation while the shared memory has not been overwritten. Hence, the selected samples can be correlated. This can also affect the independence assumption between the overwriting matrix and the selected sample (Assumption 2). However, we follow existing studies (Zhao & Li, 2016; Zhao et al., 2017; Lian et al., 2015; Reddi et al., 2015; Duchi et al., 2015; De Sa et al., 2015; Lian et al., 2018; Hsieh et al., 2015), assuming DPSGD maintains the required conditions for independence via Assumptions 2 and 6.

We are now ready to state the following convergence rate for any non-convex objective:

Theorem 1 *If Assumptions 1 to 6 hold and the following inequalities are true:*

$$M^2 \tilde{B}^2 \eta^2 L^2 \rho_{t-1} D' \sum_{n=1}^{D'} \rho_{t+n} \leq 1 \quad (7)$$

$$\frac{1}{1 - \eta - \frac{9\eta(D+1)L^2(\mu^{D+1}-1)}{\mu-1}} \leq \mu \quad (8)$$

then, we can obtain the following results:

$$\begin{aligned} & \frac{1}{\sum_{t=1}^T \rho_{t-1}} \sum_{t=1}^T \rho_{t-1} \mathbb{E}[\|\nabla f(\mathbf{v}_{t-1})\|^2] \leq \frac{2(f(\mathbf{v}_0) - f(\mathbf{v}_*))}{M\tilde{B}\eta\alpha \sum_{t=1}^T \rho_{t-1}} + \\ & \frac{\eta^2 L^2}{\tilde{B} \sum_{t=1}^T \rho_{t-1}} \sum_{t=1}^T \rho_{t-1} \left[V^2 \left(\sum_{b=0}^{\tilde{B}-1} \frac{\mu(\mu^b - 1)}{\mu - 1} + \tilde{B} \frac{\mu(\mu^D - 1)}{\mu - 1} \right) + \right. \\ & \left. M\tilde{B}^2 \sigma^2 \sum_{j=t-1-D'}^{t-2} \rho_{j-1}^2 \right] + \frac{L\eta V^2}{\alpha \sum_{t=1}^T \rho_{t-1}} \sum_{t=1}^T \rho_{t-1}^2 \end{aligned}$$

where $\tilde{B} = pB$ and \mathbf{v}_* is the global optimum of the objective function in Eq. 1.

We denote the expectation of all random variables in Algorithm 2 by $\mathbb{E}[\cdot]$. Theorem 1 shows that the weighted average of the l_2 norm of all gradients $\|\nabla f(\mathbf{v}_{t-1})\|^2$ can be bounded, which indicates an ergodic convergence rate. It can be seen that it is possible to achieve speed-up by increasing the number of cores and workers. Nevertheless to reach such speed-up, the learning rates η and ρ_t have to be set properly (see Corollary 1).

Corollary 1 *By setting the learning rates to be equal and constant:*

$$\rho^2 = \eta^2 = \frac{\sqrt{f(\mathbf{v}_0) - f(\mathbf{v}_*)}}{A\alpha\sqrt{TM\tilde{B}}} \tag{9}$$

such that $A = LV^2 \left(\frac{1}{\alpha} + \frac{1}{\alpha^2} + \frac{2L\mu}{(1-\mu)\alpha} \right)$, $V > 0$ and μ is a constant where $0 < \mu < 1$, then the bound in Eqs. 7 and 8 can lead to the following bound:

$$T \geq \max \left\{ \frac{M\tilde{B}L^2D^2(f(\mathbf{v}_0) - f(\mathbf{v}_*))}{A^2\alpha^2}, \frac{(f(\mathbf{v}_0) - f(\mathbf{v}_*))(\mu(\mu - 1) + 9L^2\mu(D + 1)(\mu^{D+1} - 1))^4}{M\tilde{B}A^2\alpha^2(\mu - 1)^8} \right\} \tag{10}$$

and Theorem 1 gives the following convergence rate:

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla f(\mathbf{v}_{t-1})\|^2] \leq 3A\sqrt{\frac{f(\mathbf{v}_0) - f(\mathbf{v}_*)}{TM\tilde{B}}} \tag{11}$$

This corollary shows that by setting the learning rates to certain values and setting the number of iterations T to be greater than a bound depending on the maximum delay allowed, a convergence rate of $O(1/\sqrt{TMp\tilde{B}})$ can be achieved and this is delay-independent. The negative effect of using old parameters (asynchronous distribution) and over-writing the shared memory (lock-free parallel) vanish asymptotically. Hence, to achieve speed-up, the number of iterations has to exceed a bound controlled by the maximum delay parameters, the number of iterations B (line 5 in Algorithm 2), the number of global updates M (line 3 in Algorithm 1) and the number of parallel threads (cores) p .

3.4 Discussion

Using Corollary 1, we can derive the result of lock-free parallel optimisation algorithm (Zhao et al., 2017) and the asynchronous distributed optimisation algorithm (Lian et al., 2015) as particular cases. By setting the number of threads $p = 1$ and the number of local update $B = 1$, we end up with the distributed asynchronous algorithm presented in Lian et al. (2015). The convergence bound of Corollary 1 then becomes $O(1/\sqrt{TM})$ which is equivalent to that of Corollary 2 in Lian et al. (2015). By synchronising the global learning $D' = 0$, setting the master batch size $M = 1$ and the number of global iteration $T = 1$, we end up with the parallel lock free algorithm presented in Zhao et al. (2017). The convergence bound of Corollary 1 then becomes $O(1/\sqrt{p\tilde{B}})$ which is equivalent to that of Theorem 1 in Zhao et al. (2017). The experiments below will empirically demonstrate these two parallel and distributed particular cases of DPSGD.

Since D' and D are related to the number of workers and cores (threads) respectively, bounding the latter allows speed-up with respect to the number of workers and cores with no loss of accuracy. The satisfaction of Eq. 10 is guaranteed if:

$$T \geq \frac{M\tilde{B}L^2D^2(f(\mathbf{v}_0) - f(\mathbf{v}_*))}{A^2\alpha^2}$$

and

$$T \geq \frac{(f(\mathbf{v}_0) - f(\mathbf{v}_*))(\mu(\mu - 1) + 9L^2\mu(D + 1)(\mu^{D+1} - 1))^4}{M\bar{B}A^2\alpha^2(\mu - 1)^8}$$

The first inequality leads to $O(T^{1/2}) > D'$. Thus, the upper bound on the number of workers is $O(T^{1/2})$. Since $0 < \mu < 1$, the second inequality can be written as follows: $O(T^{1/4}) \geq (\mu(1 - \mu) + 9L^2\mu(D + 1)(1 - \mu^{D+1}))$. Hence, $O(T^{1/4}) \geq D$. Thus, the upper bound on the number of number of cores (threads) is $O(T^{1/4})$. The convergence rate for serial and synchronous parallel stochastic gradient (SG) is consistent with $O(1/\sqrt{T})$ (Ghadimi & Lan, 2013; Dekel et al., 2012; Nemirovski et al., 2009). While the workload for each worker running DPSGD is almost the same as the workload of the serial or synchronous parallel SG, the progress done by DPSGD is Mp times faster than that of serial SG.

In addition to the speed-up, DPSGD allows one to steer the trade-off between multi-core local computation and multi-node communication within the cluster. This can be done by controlling the parameter B . Traditional methods reduce the communication cost by increasing the batch size which decreases the convergence rate, increase local memory load and decrease local input bandwidth. On the contrary, increasing B for DPSGD can increase the speed-up if some assumptions are met (see Theorem 1 and Corollary 1). This ability makes DPSGD easily adaptable to diverse spectrum of large-scale computing systems with no loss of speed-up.

Denote T_c the communication time need for each master-worker exchange. For simplification, we assume that T_c is fixed and is the same for all nodes. If the time needed for computing one update $T_u \leq T_c$, then the total time needed by the distributed algorithm $DTT = T * (Tu + Tc)$ could be higher than that of the sequential SGD $STT = M * T * Tu$. In such cases, existing distributed algorithms increases the local batch size so that T_u increases, resulting in lower stochastic gradient variance and allowing higher learning rate to be used, hence better convergence rate. This introduces a trade-off between computational efficiency and sample efficiency. Increasing the batch size by a factor of k increases the time need for local computation by $O(k)$ and reduces the variance proportionally to $1/k$ (Bottou et al., 2018). Thus, higher learning rate can be used. However, there is a limit on the size of the learning rate. In another words, maximising the learning speed with respect to the learning rate and the batch size has a global solution. This maximum learning speed can be improved using DPSGD, performing B times less communication steps. For the mini-batch SGD with minibatch size G , the convergence rate can be written as $O(1/\sqrt{GT})$. Since the total number of examples examined is GT and there is only \sqrt{G} times improvement, the convergence speed degrades as mini-batch size increases. The convergence rate of DPSGD with mini-batch G can be easily deduced from Theorem 1 as $O(1/\sqrt{BMGT})$. Hence, \sqrt{BM} better convergence rate than mini-batch SGD and \sqrt{BM} better convergence rate than standard asynchronous SGD with B times less communication. These improvements are studied in the following.

4 Experiments

In this section, we empirically verify the potential speed-up gains expected from the theoretical analysis. First, we apply distributed parallel stochastic variational inference (DPSVI) algorithm on a Latent Dirichlet Allocation (LDA) analysis problem. DPSVI

is derived from DPSGD by replacing the SG of SVI by DPSG to scale up the inference computation over a multi-core cluster (see appendix for more details). For the Latent Dirichlet Allocation analysis problem, we use the SVI algorithm (Hoffman et al., 2013) as benchmark. The evaluation is done on 300, 000 news articles from the *New York Times* corpus.

Furthermore, we use DPSGD to scale up the training of DRL algorithm, namely *Advantage Actor Critic (A2C)* algorithm, implementing *highly scalable A2C (HSA2C)* (details in the appendix). We compare HSA2C against other distributed A2C implementations using a testbed of six Atari games and demonstrate an average training time of 21.95 min compared to over 13.75 h by the baseline A3C. In particular, HSA2C shows a significant speed-up on Space invaders with learning time below 10 min compared to the 30 min achieved by the best competitor.

4.1 Variational inference

The development of the proposed DPSVI algorithm follow from DPSGD, but in the context of VI. In Appendix 7, we characterise the entire family of models where DPSVI is applicable, which is shown to be equivalent to the models for which SVI applies. Next, DPSVI is derived from DPSGD. Finally, we derive an asynchronous distributed lock-free parallel inference algorithm for LDA as a case study for DPSVI.

Datasets We use the *NYTimes* corpus (Lichman, 2013) containing 300, 000 news articles from the *New York Times* corpus. The data is pre-processed by removing all the words not found in a dictionary containing 102, 660 most frequent words - see (Lichman, 2013) for more information. We reserve 5, 000 documents from *NYTimes* data as a validation set and another 5, 000 documents as a testing set.

Performance The performance of the LDA model is assessed using a model fit measure, *perplexity*, which is defined as the geometric mean of the inverse marginal probability of each word in the held-out set of documents (Blei et al., 2003). We also compute the running time speed-up (TSP) (Lian et al., 2015) defined as

$$TSP = \frac{T(\text{SVI})}{T(\text{DPSVI})} \quad (12)$$

where $T(\cdot)$ denotes the running time and is taken when both models achieve the same final held-out perplexity of 5000 documents.

Parameters In all experiments, the LDA number of topics is $K = 50$. SVI LDA is run on the training set for $\kappa \in \{0.5, 0.7, 0.9\}$, $\tau_0 \in \{1, 24, 256, 1024\}$, and $batch \in \{16, 64, 256, 1024\}$. The best performing parameters $batch = 1024$, $\kappa = 0.5$ and $\tau_0 = 1$ providing perplexity of 5501 are used (Table 1 in Mohamad et al. (2018) summarises the best settings with the resulting perplexity on the test set). As for the DPSGD LDA version, the local learning rate G (see Eq. 52) is set to 64 and M equal to 16. We evaluate a range of learning rates $\eta = \rho \in \{0.2, 0.1, 0.05, 0.01\}$ where M , p and B are set to 1. The best learning rate 0.1 providing held-out perplexity of 5501 was used. For different B , M and p , the learning rate is changed according to Corollary 1:

$$\rho' = \rho \left(\frac{pBM}{p'B'M'} \right)^{0.25} = \frac{0.1}{(p'B'M')^{0.25}} \quad (13)$$

Fig. 1 LDA analysis: Running time speed-up (TSP) with respect to the number of workers

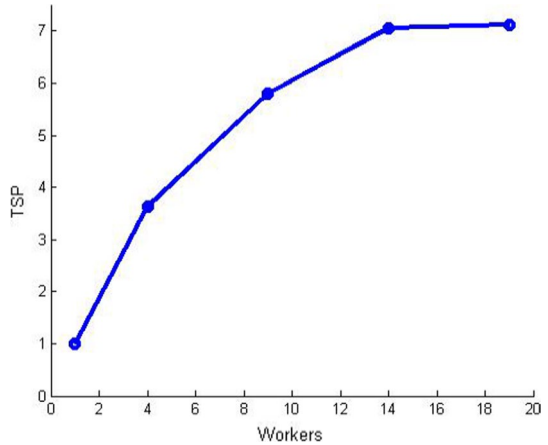
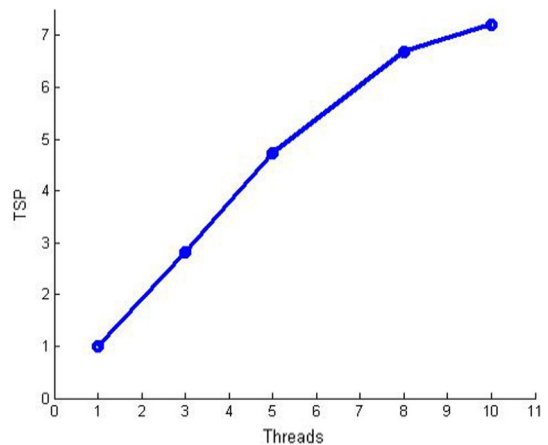


Fig. 2 LDA analysis: Running time speed-up (TSP) with respect to the number of threads



All DPSGD LDA experiments were performed on a high-performance computing (HPC) environment using message passing interface (MPI) for python (MPI4py). The *cluster* consists of 10 nodes, including the head node, with each node being a 1-sockets-6-cores-2-thread processor.

4.1.1 Node speed-up

Here, we study the speed-up of DPSVI with respect to the number of workers where $p = 1$ and $B = 1$. DSPSVI LDA is then compared against serial SVI ($B = 1$, $p = 1$ and $nW = 1$). We run DPSVI for various numbers of workers $nW \in \{4, 9, 14, 19\}$. The number of nodes is nW as long as nW is less than 9. As nW becomes higher than the available nodes, the processors' cores of nodes are employed as workers until all cores (threads) of each node are used i.e., $9 \times 12 = 108$. The batch size M is fixed to 36.

Fig. 3 LDA analysis: Running time speed-up (TSP) with respect to the number of workers and threads

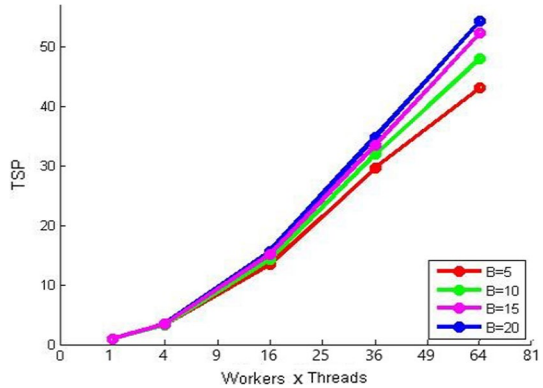


Fig. 4 LDA analysis using DPSVI: perplexity (model fit) with respect to running logarithmic time in seconds

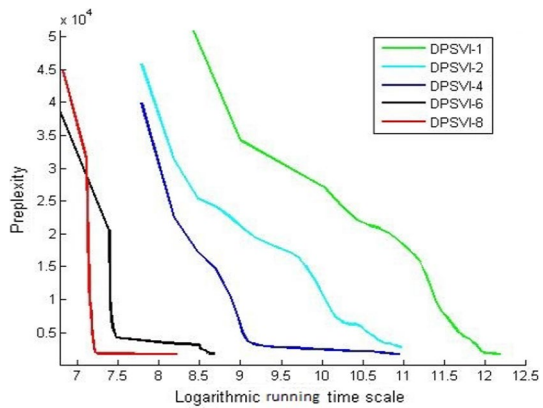


Figure 1 summarises the total speed-up (i.e., TSP measured at the end of the algorithm) with respect to the number of workers where the achieved pre-perplexity is almost the same. The result shows linear speed-up as long as the number of workers is less than 14. Then, linear speed-up slowly converts to sub-linear and is expected to drop for higher number of workers due to reaching the maximum communication bandwidth.

4.1.2 Thread speed-up

In this section, we study the speed-up of DPSVI with respect to the number of threads where $nW = 1$. We empirically set B to 15. Similar to the node-related speed-up analysis, experiments are run for different $p \in \{3, 5, 8, 10\}$. Then, DPSVI is compared against serial SVI. The results are shown in Fig. 2. The outcome shows linear speed-up as long as the number of threads is less than 8. Then, the speed-up slowly converts to sub-linear and is expected to become worse for higher number of threads. This drop in the speed-up is due to hardware communication and other factors affecting the CPU power.

4.1.3 Node-thread speed-up

Finally, we study the speed-up of DPSVI with respect to the number of nodes and threads. To simplify the experiments, we take the number of cores to be equal to the number of nodes. Experiments are run for different $p = nW \in \{2, 4, 6, 8\}$. We also present results with different $B \in \{5, 10, 15, 20\}$ in order to show the effect of steering the trade-off between local computation and communication. DPSVI is compared against serial SVI and the results are shown in Fig. 3. The result shows speed-up whose speed slows down as the number of threads and nodes exceed 6. This is due to communication and other hardware factors. However, the rate of this slowing down for higher B is less significant which illustrates the advantage of reducing the communication overhead when reaching its ceiling point. Note that for very high number of workers, increasing B might not be very helpful as our theoretical results show that high B tightens the bound on the number of workers allowed for the speed-up to hold. Figure 4 reports the perplexity on the training set with respect to running time in seconds (logarithmic scale) with $B = 15$. Five curves are drawn for different nodes-threads number, where DPSVI- n denotes our DPSVI with n nodes and threads. The convergence and speed-up of DPSVI are clearly illustrated.

4.2 Deep reinforcement learning

We use six different Atari games to study the performance gains that can be achieved by the proposed HSA2C algorithm using the Atari 21600 emulator (Bellemare et al., 2013) provided by the OpenAI Gym framework (Brockman et al., 2016). This emulator is one of the most commonly used benchmark environments for RL algorithms. Here, we use Pong, Boxing, Seaquest Space invaders, Amidar and Qbert which have been included in related work (Mnih et al., 2016; Adamski et al., 2017, 2018; Babaeizadeh et al., 2016). These games are used to evaluate the effects of reducing the communication bottleneck when using an increasingly higher number of steps, B , with different numbers of nodes. We also study the speed-up achieved by HSA2C with respect to the number of nodes. Finally, we compare the performance reported by various state-of-the-art algorithms (Mnih et al., 2016; Adamski et al., 2017, 2018; Babaeizadeh et al., 2016).

4.2.1 Implementation details

HSA2C has been implemented and tested on a high-performance computing (HPC) environment using message passing interface (MPI) for Python (MPI4py 3.0.0) and Pytorch 0.4.0. Our cluster consists of 60 nodes consisting of 28 2.4 GHz CPUs per node. In our experiments, we used the same input pre-processing as Mnih et al. (2015). Each experiment was repeated 5 times (each with an action repeat of 4) and the average results are reported. The agents used the neural network architectures described in Mnih et al. (2013): a convolutional layer with 16 filters of size 8×8 with stride 4, followed by a convolutional layer with 32 filters of size 4×4 with stride 2, followed by a fully connected layer with 256 hidden units. All three hidden layers were followed by a rectifier nonlinearity. The network has two sets of outputs – a softmax output with one entry per action representing the probability of selecting the action and a single linear output representing the value function. Local learning rate, mini-batch size and the optimiser setting are contrasted with those reported in Adamski et al. (2018) in order to provide a fair comparison with their asynchronous mini-batch implementation. The global learning rate was set to 0.01 for the

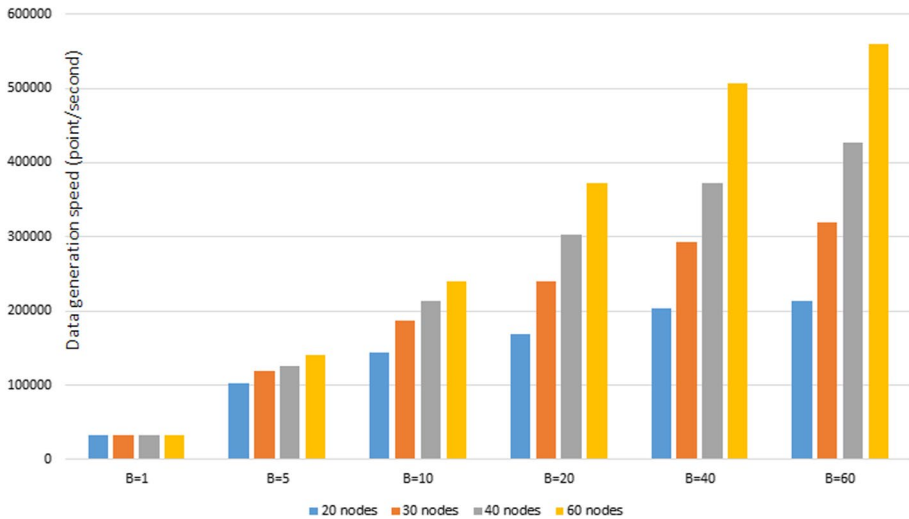


Fig. 5 The average data generation speed of HSA2C measured in points per second within 30 min run on Space invader game

SGD optimiser with 0.5 momentum. The global batch was set to the number of utilised nodes.

4.2.2 Speed-up analysis

In this section, we study the effect of B on various aspects of the scalability of HSA2C with respect to the number of nodes. Figure 5 shows the average speed of data generated by the distributed actors measured in data-points per seconds on the Space invader game. Comparable figures can be obtained for other games, but they are not reported here. It is noticeable that at $B = 1$, the speed of data generation is about the same as the number of nodes increases. This is due to the communication cost which increases with the number of nodes used. That is, the expected waiting time of each node's exchange with the master increases. Increasing B will reduce the number of exchanges while performing the same number of updates locally. This is illustrated in Fig. 5 where the data generation speed increases with the number of nodes as B increases.

Figure 6 shows the time (in seconds) required to reach the highest score of Pong, Boxing, Seaquest, Space invader, Amidar and Qbert reached with $B = 1$ in a 30-minute run over different number of nodes and B . The aim of these figures is to demonstrate the potential performance gains achieved by HSA2C as B increases in comparison to distributed deep reinforcement learning (DDRL) A3C, which is algorithmically equivalent to HSA2C when $B = 1$ (our baseline). In order to produce these figures, we initially carried out a search to empirically determine the highest score that can be achieved by HSA2C within a time period of 30 minutes when $B = 1$. The search for the highest score is performed on four different cluster sizes: 20, 30, 40, and 60. The figures present the time required for various HSA2C parameters to reach that benchmark score. These experimental results clearly show the impact of B on the communication costs and confirm the findings of Fig. 5. By using a larger number of nodes, more communication exchanges are required for each update and performing more local update (i.e., increasing B) reduces the communication exchanges

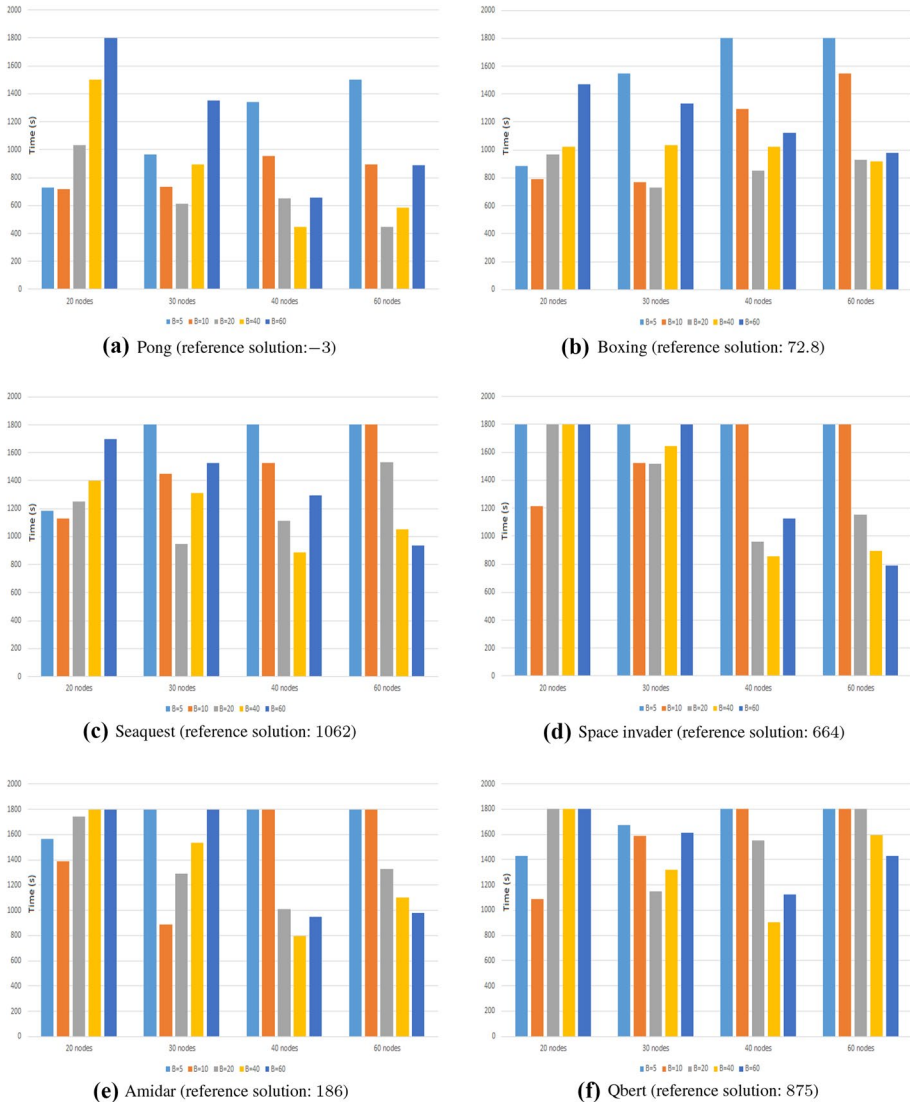


Fig. 6 The time (in seconds) required to reach reference solution (the highest score with $B = 1$ in a 30 min run) over a range of node numbers and B

needed to reach certain score without much sacrificing the learning performance. Thus, a better speed-up can be achieved. On the other hand, with a smaller number of nodes, increasing B does not make a significant difference in reducing the communication whilst the negative effect of an increased variance becomes significant as the size of the learning batch becomes smaller (depending on the number of nodes). Overall, there is evidence for a variance-communication trade-off controlled by B .

Table 2 Best scores and the corresponding time in minutes achieved by HSA2C using the best B from Table 2 and considering 20, 30, 40 and 60 nodes compared to the best reported results by competitors

Algorithms	Pong	Boxing	Seaquest	Space invaders	Amidar	Qbert
HSA2C (20 nodes)	20 (63 m)	98 (19.8 m)	1858 (21.7 m)	684 (17.9 m)	251 (26 m)	1210 (20 m)
HSA2C (30 nodes)	19 (40 m)	98 (19.8 m)	1874 (19.1 m)	671 (14.7 m)	234 (15 m)	1420 (27 m)
HSA2C (40 nodes)	20 (32.6 m)	95 (20.7 m)	1894 (20.5 m)	681 (9.9 m)	282 (20 m)	1620 (28 m)
HSA2C (60 nodes)	19 (38 m)	96 (23.2 m)	1874 (22.3 m)	667 (9.4 m)	243 (24 m)	800 (30 m)
DDRL A3C	20 (240 m)	98 (30 m)	1832 (30 m)	650 (30 m)	–	–
GA3C	18 (60 m)	92 (120 m)	1706 (1440 m)	600 (1440 m)	218 (1440 m)	395 (30 m)
BA3C	17 (1440 m)	–	1840 (1440 m)	700 (1440 m)	–	–
A3C	20 (480 m)	95 (660 m)	2300 (1440 m)	1400 (900 m)	280 (1440 m)	400 (30 m)

The multiple local updates B mitigates the speed-up limit caused by higher communication cost with higher number of nodes. Choosing the right B for different numbers of nodes allows HSA2C to scale better than HSA2C when $B = 1$. For all the six games, we have found that increasing the number of nodes over 40 does not lead to better performance. This is due to the higher variance entailed by a higher B . That is, the performance improvement coming from the communication reduction is overtaken by the entailed variance when using more than 40 nodes. This limit could be overcome in two different ways, either by decreasing the communication without further increasing the variance or by directly mitigating the variance problems.

4.2.3 Comparison

We show the effectiveness of the proposed approach by comparing it against similar scalable Actor-critic optimisation approaches. The most similar work in speeding up Atari games training is presented in Adamski et al. (2018), Mnih et al. (2016), Adamski et al. (2017), Babaeizadeh et al. (2016). The algorithm in Adamski et al. (2018) (DDRL A3C) is a particular case of HSA2C, where the communication is synchronised and the number of iterations of per-worker loop is set to one ($B = 1$). GA3C is a hybrid GPU/CPU algorithm which is a flavour of A3C focusing on batching the data points in order to better utilise the massively parallel nature of GPU computations. This is similar to the single-node algorithm called BA3C (Adamski et al., 2017).

Table 2 presents the best score and time (in minutes) HSA2C archives using the best B values found in Fig. 6 in comparison to the competitors. The reported scores are taken from the original papers. As GA3C, BA3C and A3C are parallel single-node algorithms, their experimental settings are not comparable to ours. This comparison shows that our approach achieves a better score than all competitors. In particular, we achieve an average score of 665 in 30.63 minutes average time using 560 total CPU cores compared to the DDRL A3C score of 650 in 82.5 average time with 778 total CPU cores. Most importantly, this comparison validates the effectiveness of our proposed approach to reduce communication while preserving performance. This is clearly shown in the

comparison between our approach and our implementation of DDRL A3C (the top competitor in Table 2) using the same setting (see Sect. 4.2.2).

For this study, we have decided not to include GPU-based implementations such as Stooke and Abbeel (2018) as our focus here is on CPU-enabled methods. However, HSA2C is generic and lends itself to GPU-based implementations whereby each node consists of multiple CPUs and a GPU. In such a case, local computation and simulation can be done using CPUs/GPU units, where our multiple local update approach can further speed up the standard DA3C communication-based (Stooke & Abbeel, 2018). The empirical work reported here provides an initial validation of the underlying idea.

5 Conclusion

We have proposed a novel asynchronous distributed and lock-free parallel optimisation algorithm. The algorithm has been implemented on a computer cluster with multi-core nodes. Both theoretical and empirical results have shown that DPSGD leads to speed-up on non-convex problems. The paper shows how DPSVI and HSA2C have been derived from DPSGD. Both are an asynchronous distributed and lock-free parallel implementation for respectively stochastic variational inference (SVI) and advantage actor critic (A2C). Empirical results have allowed to validate the theoretical findings and to compare against similar state-of-the-art methods.

Going forward, further improvements and validations could be achieved by pursuing research along five directions: (1) employing variance reduction techniques to improve the convergence rate (from sub-linear to linear) while guaranteeing multi-node and multi-core speed-up; (2) proposing a framework enabling dynamic trade-offs between local computation and communication; (3) proposing techniques to improve the local optimum of the distributed parallel algorithms; (4) applying DPSVI to other members of the family of models stated in the appendix; (5) applying DPSGD to other large-scale deep learning problems.

6 Proofs

Let $q(\mathbf{x}) = \frac{1}{n} \sum_{i=1}^n \|\nabla f_i(\mathbf{x})\|^2$. We have $\mathbb{E}_i[\|\nabla f_i(\mathbf{x})\|^2] = q(\mathbf{x})$. Hence, $\mathbb{E}_i[q(\mathbf{x})] = \mathbb{E}_i[\|\nabla f_i(\mathbf{x})\|^2]$. Taking the full expectation on both sides, we get $\mathbb{E}[q(\mathbf{x})] = \mathbb{E}[\|\nabla f_i(\mathbf{x})\|^2]$. It can be proven that:

$$\mathbb{E}_t[q(\hat{\mathbf{u}}_{t,m,j})] < \mu \mathbb{E}_t[q(\hat{\mathbf{u}}_{t,m,j+1})] \quad (14)$$

Lemma 1

given μ and η satisfying

$$\frac{1}{1 - \eta - \frac{9\eta(D+1)L^2(\mu^{D+1}-1)}{\mu-1}} \leq \mu \quad (15)$$

where $\mathbb{E}_t[\cdot]$ denotes $\mathbb{E}_{i_{t,*}, s_{t,*}}[\cdot]$. The proof can be derived from that in Zhao et al. (2017) using Assumptions 1 and 5. The stars means for all, \forall .

Proof to Theorem 1 From the smoothness Assumption 1, we have:

$$\begin{aligned}
 f(\mathbf{v}_t) - f(\mathbf{v}_{t-1}) &\leq \langle \nabla f(\mathbf{v}_{t-1}), \mathbf{v}_t - \mathbf{v}_{t-1} \rangle + \frac{L}{2} \|\mathbf{v}_t - \mathbf{v}_{t-1}\|^2 \\
 &= -\langle \nabla f(\mathbf{v}_{t-1}), \eta \rho_{t-1} \sum_{m=1}^M \sum_{b=0}^{\tilde{B}-1} S_{t,m,b} \nabla f_{i_{t,m,b}}(\hat{\mathbf{u}}_{t-\tau_{t,m},m,b}) \rangle \\
 &\quad + \frac{L}{2} \|\eta \rho_{t-1} \sum_{m=1}^M \sum_{b=0}^{\tilde{B}-1} S_{t,m,b} \nabla f_{i_{t,m,b}}(\hat{\mathbf{u}}_{t-\tau_{t,m},m,b})\|^2
 \end{aligned}
 \tag{16}$$

where the last equality uses the update in Eq. (4). Taking expectation of the above inequality with respect to $i_{t,*,*}$, and $S_{t,*,*}$, we obtain:

$$\begin{aligned}
 E_t[f(\mathbf{v}_t)] - f(\mathbf{v}_{t-1}) &\leq -M\tilde{B}\eta\rho_{t-1} \langle \nabla f(\mathbf{v}_{t-1}), \frac{1}{M\tilde{B}} \sum_{m=1}^M \sum_{b=0}^{\tilde{B}-1} S \nabla f(\hat{\mathbf{u}}_{t-\tau_{t,m},m,b}) \rangle \\
 &\quad + \frac{L\eta^2\rho_{t-1}^2}{2} E_t \|\sum_{m=1}^M \sum_{b=0}^{\tilde{B}-1} S_{t,m,b} \nabla f_{i_{t,m,b}}(\hat{\mathbf{u}}_{t-\tau_{t,m},m,b})\|^2
 \end{aligned}
 \tag{17}$$

where we used Assumptions 2, 3, 4 and 6. Since S is a strictly definite matrix with the largest *eigenvalue* less or equal than 1 and the minimum *eigenvalue* is $\alpha > 0$ and from the fact $\langle a, b \rangle = \frac{1}{2}(\|a\|^2 + \|b\|^2 - \|a - b\|^2)$, we have:

$$\begin{aligned}
 E_t[f(\mathbf{v}_t)] - f(\mathbf{v}_{t-1}) &\leq -\frac{M\tilde{B}\eta\rho_{t-1}\alpha}{2} \left(\|\nabla f(\mathbf{v}_{t-1})\|^2 + \left\| \frac{1}{M\tilde{B}} \sum_{m=1}^M \sum_{b=0}^{\tilde{B}-1} \nabla f(\hat{\mathbf{u}}_{t-\tau_{t,m},m,b}) \right\|^2 - \right. \\
 &\quad \left. \underbrace{\|\nabla f(\mathbf{v}_{t-1}) - \frac{1}{M\tilde{B}} \sum_{m=1}^M \sum_{b=0}^{\tilde{B}-1} \nabla f(\hat{\mathbf{u}}_{t-\tau_{t,m},m,b})\|^2}_{\text{H1}} \right) + \\
 &\quad \frac{L\eta^2\rho_{t-1}^2}{2} E_t \underbrace{\|\sum_{m=1}^M \sum_{b=0}^{\tilde{B}-1} S_{t,m,b} \nabla f_{i_{t,m,b}}(\hat{\mathbf{u}}_{t-\tau_{t,m},m,b})\|^2}_{\text{H2}}
 \end{aligned}
 \tag{18}$$

Next, we obtain an upper bound for H1. Using the triangular inequality and Assumption 1, we can write the following:

$$\begin{aligned}
 \text{H1} &= \|\nabla f(\mathbf{v}_{t-1}) - \frac{1}{M\tilde{B}} \sum_{m=1}^M \sum_{b=0}^{\tilde{B}-1} \nabla f(\hat{\mathbf{u}}_{t-\tau_{t,m},m,b})\|^2 \\
 &\leq \frac{1}{M\tilde{B}} \sum_{m=1}^M \sum_{b=0}^{\tilde{B}-1} \|\nabla f(\mathbf{v}_{t-1}) - \nabla f(\hat{\mathbf{u}}_{t-\tau_{t,m},m,b})\|^2 \\
 &\leq \frac{L^2}{\tilde{B}} \sum_{b=0}^{\tilde{B}-1} \|\mathbf{v}_{t-1} - \hat{\mathbf{u}}_{t-\tau_{t,y},y,b}\|^2
 \end{aligned}
 \tag{19}$$

where $y = \operatorname{argmax}_{m \in \{1, \dots, M\}} \|\mathbf{v}_{t-1} - \hat{\mathbf{u}}_{t-\tau_{t,m},m,b}\|^2$. Using triangular inequality and the updates in Eq. (4), we obtain:

$$\begin{aligned}
 H1 &\leq \frac{L^2}{\bar{B}} \sum_{b=0}^{\bar{B}-1} \left(\|\mathbf{v}_{t-1} - \mathbf{u}_{t-\tau_{t,y},y,a(b)}\|^2 + \eta^2 \left\| \sum_{j=a(b)}^{b-1} \nabla f_{i_{t,y,j}}(\hat{\mathbf{u}}_{t-\tau_{t,y},y,j}) \right\|^2 \right) \\
 &\leq \frac{L^2}{\bar{B}} \sum_{b=0}^{\bar{B}-1} \left(\|\mathbf{v}_{t-1} - \mathbf{v}_{t-1-\tau_{t,y}}\|^2 + \eta^2 \left\| \sum_{j=0}^{b-1} \nabla f_{i_{t,y,j}}(\hat{\mathbf{u}}_{t-\tau_{t,y},y,j}) \right\|^2 + \eta^2 \left\| \sum_{j=a(b)}^{b-1} \nabla f_{i_{t,y,j}}(\hat{\mathbf{u}}_{t-\tau_{t,y},y,j}) \right\|^2 \right) \\
 &= \frac{L^2}{\bar{B}} \sum_{b=0}^{\bar{B}-1} \left(\underbrace{\eta^2 \left\| \sum_{j=t-1-\tau_{t,y}}^{t-2} \rho_{j-1}^2 \sum_{m=1}^M \sum_{b=0}^{\bar{B}-1} \nabla f_{i_{j,m,b}}(\hat{\mathbf{u}}_{j-\tau_{j,m},m,b}) \right\|^2}_{H11} + \eta^2 \left\| \sum_{j=0}^{b-1} \nabla f_{i_{t,y,j}}(\hat{\mathbf{u}}_{t-\tau_{t,y},y,j}) \right\|^2 \right) \\
 &\quad + \eta^2 \left\| \sum_{j=a(b)}^{b-1} \nabla f_{i_{t,y,j}}(\hat{\mathbf{u}}_{t-\tau_{t,y},y,j}) \right\|^2
 \end{aligned} \tag{20}$$

Using the triangular inequality again, we have the following:

$$\begin{aligned}
 H11 &= \left\| \sum_{j=t-1-\tau_{t,y}}^{t-2} \rho_{j-1}^2 \sum_{m=1}^M \sum_{b=0}^{\bar{B}-1} \nabla f_{i_{j,m,b}}(\hat{\mathbf{u}}_{j-\tau_{j,m},m,b}) \right\|^2 \\
 &= \left\| \sum_{j=t-1-\tau_{t,y}}^{t-2} \rho_{j-1}^2 \sum_{m=1}^M \sum_{b=0}^{\bar{B}-1} \left(\nabla f_{i_{j,m,b}}(\hat{\mathbf{u}}_{j-\tau_{j,m},m,b}) - \nabla f(\hat{\mathbf{u}}_{j-\tau_{j,m},m,b}) + \nabla f(\hat{\mathbf{u}}_{j-\tau_{j,m},m,b}) \right) \right\|^2 \\
 &\leq \left\| \sum_{j=t-1-\tau_{t,y}}^{t-2} \rho_{j-1}^2 \sum_{m=1}^M \sum_{b=0}^{\bar{B}-1} \left(\nabla f_{i_{j,m,b}}(\hat{\mathbf{u}}_{j-\tau_{j,m},m,b}) - \nabla f(\hat{\mathbf{u}}_{j-\tau_{j,m},m,b}) \right) \right\|^2 \\
 &\quad + \left\| \sum_{j=t-1-\tau_{t,y}}^{t-2} \rho_{j-1}^2 \sum_{m=1}^M \sum_{b=0}^{\bar{B}-1} \nabla f(\hat{\mathbf{u}}_{j-\tau_{j,m},m,b}) \right\|^2
 \end{aligned} \tag{21}$$

By taking the expectation on both sides with respect to all random variables associated with $k \in \{t-1-\tau_{t,y}, \dots, t-2\}$ and using Assumptions 4, 5 and 6 and , we obtain:

$$\begin{aligned}
 E_k[H11] &\leq \sum_{j=t-1-\tau_{t,y}}^{t-2} \rho_{j-1}^2 E_k \left\| \sum_{m=1}^M \sum_{b=0}^{\bar{B}-1} \left(\nabla f_{i_{j,m,b}}(\hat{\mathbf{u}}_{j-\tau_{j,m},m,b}) - \nabla f(\hat{\mathbf{u}}_{j-\tau_{j,m},m,b}) \right) \right\|^2 \\
 &\quad + D' \sum_{j=t-1-\tau_{t,y}}^{t-2} \rho_{j-1}^2 E_k \left\| \sum_{m=1}^M \sum_{b=0}^{\bar{B}-1} \nabla f(\hat{\mathbf{u}}_{j-\tau_{j,m},m,b}) \right\|^2 \\
 &\leq M\bar{B}\sigma^2 \sum_{j=t-1-\tau_{t,y}}^{t-2} \rho_{j-1}^2 + D' \sum_{j=t-1-\tau_{t,y}}^{t-2} \rho_{j-1}^2 E_k \left\| \sum_{m=1}^M \sum_{b=0}^{\bar{B}-1} \nabla f(\hat{\mathbf{u}}_{j-\tau_{j,m},m,b}) \right\|^2
 \end{aligned} \tag{22}$$

Taking the full expectation, we have

$$\begin{aligned}
 E[H1] \leq & \frac{L^2\eta^2}{\tilde{B}} \sum_{b=0}^{\tilde{B}-1} \left(M\tilde{B}\sigma^2 \sum_{j=t-1-\tau_{t,y}}^{t-2} \rho_{j-1}^2 + D' \sum_{j=t-1-\tau_{t,y}}^{t-2} \rho_{j-1}^2 E_k \left\| \sum_{m=1}^M \sum_{b=0}^{\tilde{B}-1} \nabla f(\hat{\mathbf{u}}_{j-\tau_{j,m},m,b}) \right\|^2 \right. \\
 & \left. + \left\| \sum_{j=0}^{b-1} \nabla f_{i_{t,y,j}}(\hat{\mathbf{u}}_{t-\tau_{t,y},y,j}) \right\|^2 + \left\| \sum_{j=a(b)}^{b-1} \nabla f_{i_{t,y,j}}(\hat{\mathbf{u}}_{t-\tau_{t,y},y,j}) \right\|^2 \right)
 \end{aligned} \tag{23}$$

Using Assumption 4 and Lemma 1 we have the following:

$$\begin{aligned}
 E[H1] \leq & \frac{L^2\eta^2}{\tilde{B}} \sum_{b=0}^{\tilde{B}-1} \left(M\tilde{B}\sigma^2 \sum_{j=t-1-\tau_{t,y}}^{t-2} \rho_{j-1}^2 + D' \sum_{j=t-1-\tau_{t,y}}^{t-2} \rho_{j-1}^2 E \left\| \sum_{m=1}^M \sum_{b=0}^{\tilde{B}-1} \nabla f(\hat{\mathbf{u}}_{j-\tau_{j,m},m,b}) \right\|^2 \right. \\
 & \left. + \frac{\mu(\mu^b - 1)}{\mu - 1} V^2 + \frac{\mu(\mu^D - 1)}{\mu - 1} V^2 \right)
 \end{aligned} \tag{24}$$

Using Assumption 4, the full expectation of $H2$ can be bounded as follows:

$$E[H2] \leq \frac{L\eta^2\rho_{t-1}^2}{2} M\tilde{B}V^2 \tag{25}$$

By taking the full expectation of Eq. 18 and applying the upper bound of $E[H1]$ and $E[H2]$, we obtain:

$$\begin{aligned}
 E[f(\mathbf{v}_t)] - E[f(\mathbf{v}_{t-1})] \leq & -\frac{M\tilde{B}\eta\rho_{t-1}\alpha}{2} \left[E[\|\nabla f(\mathbf{v}_{t-1})\|^2] + E \left[\left\| \frac{1}{M\tilde{B}} \sum_{m=1}^M \sum_{b=0}^{\tilde{B}-1} \nabla f(\hat{\mathbf{u}}_{t-\tau_{t,m},m,b}) \right\|^2 \right] \right] \\
 & - \frac{L^2\eta^2}{\tilde{B}} \sum_{b=0}^{\tilde{B}-1} \left(M\tilde{B}\sigma^2 \sum_{j=t-1-\tau_{t,y}}^{t-2} \rho_{j-1}^2 + D' \sum_{j=t-1-\tau_{t,y}}^{t-2} \rho_{j-1}^2 E \left\| \sum_{m=1}^M \sum_{b=0}^{\tilde{B}-1} \nabla f(\hat{\mathbf{u}}_{j-\tau_{j,m},m,b}) \right\|^2 \right. \\
 & \left. + \frac{\mu(\mu^b - 1)}{\mu - 1} V^2 + \frac{\mu(\mu^D - 1)}{\mu - 1} V^2 \right) + \frac{L\eta^2\rho_{t-1}^2}{2} M\tilde{B}V^2 \\
 \leq & -\frac{M\tilde{B}\eta\rho_{t-1}\alpha}{2} E[\|\nabla f(\mathbf{v}_{t-1})\|^2] + \frac{M\eta^3\rho_{t-1}\alpha L^2V^2}{2} \sum_{b=0}^{\tilde{B}-1} \frac{\mu(\mu^b - 1)}{\mu - 1} \\
 & + \frac{M\tilde{B}\eta^3L^2\rho_{t-1}\alpha V^2}{2} \frac{\mu(\mu^D - 1)}{\mu - 1} + \frac{L\eta^2\rho_{t-1}^2}{2} M\tilde{B}V^2 + \frac{M^2\tilde{B}^2\eta^3L^2\rho_{t-1}\alpha}{2} \sigma^2 \sum_{j=t-1-D'}^{t-2} \rho_{j-1}^2 - \\
 & \frac{\eta\rho_{t-1}\alpha}{2M\tilde{B}} E \left[\left\| \sum_{m=1}^M \sum_{b=0}^{\tilde{B}-1} \nabla f(\hat{\mathbf{u}}_{t-\tau_{t,m},m,b}) \right\|^2 \right] + \frac{M\tilde{B}\eta^3L^2\rho_{t-1}\alpha D'}{2} \sum_{j=t-1-D'}^{t-2} \\
 & \rho_{j-1}^2 E \left[\left\| \sum_{m=1}^M \sum_{b=0}^{\tilde{B}-1} \nabla f(\hat{\mathbf{u}}_{j-\tau_{j,m},m,b}) \right\|^2 \right]
 \end{aligned} \tag{26}$$

Summarising the inequality Eq. 26, from $t = 1$ to T , we end up with:

$$\begin{aligned}
E[f(\mathbf{v}_T)] - f(\mathbf{v}_0) &\leq -\frac{M\tilde{B}\eta\alpha}{2} \sum_{t=1}^T \rho_{t-1} E[\|\nabla f(\mathbf{v}_{t-1})\|^2] \\
&+ \left(\sum_{b=0}^{\tilde{B}-1} \frac{\mu(\mu^b - 1)}{\mu - 1} + \tilde{B} \frac{\mu(\mu^D - 1)}{\mu - 1} \right) \frac{M\eta^3 L^2 V^2 \alpha}{2} \sum_{t=1}^T \rho_{t-1} \\
&+ \frac{L\eta^2}{2} M\tilde{B}V^2 \sum_{t=1}^T \rho_{t-1}^2 + \frac{M^2 \tilde{B}^2 \eta^3 L^2 \alpha}{2} \sigma^2 \sum_{t=1}^T \rho_{t-1} \sum_{j=t-1-D'}^{t-2} \rho_{j-1}^2 \\
&- \frac{\eta\alpha}{2M\tilde{B}} \sum_{t=1}^T \rho_{t-1} E[\|\sum_{m=1}^M \sum_{b=0}^{\tilde{B}-1} \nabla f(\hat{\mathbf{u}}_{t-\tau_{t,m},m,b})\|^2] \\
&+ \frac{M\tilde{B}\eta^3 L^2 \alpha D'}{2} \sum_{t=1}^T \rho_{t-1}^2 E[\|\sum_{m=1}^M \sum_{b=0}^{\tilde{B}-1} \nabla f(\hat{\mathbf{u}}_{t-\tau_{t,m},m,b})\|^2] \sum_{n=1}^{D'} \rho_{t+n} \\
&= -\frac{M\tilde{B}\eta\alpha}{2} \sum_{t=1}^T \rho_{t-1} E[\|\nabla f(\mathbf{v}_{t-1})\|^2] \\
&+ \frac{M\eta^3 L^2 \alpha}{2} \sum_{t=1}^T \rho_{t-1} \left[V^2 \left(\sum_{b=0}^{\tilde{B}-1} \frac{\mu(\mu^b - 1)}{\mu - 1} + \tilde{B} \frac{\mu(\mu^D - 1)}{\mu - 1} \right) \right. \\
&+ \left. M\tilde{B}^2 \sigma^2 \sum_{j=t-1-D'}^{t-2} \rho_{j-1}^2 \right] + \frac{L\eta^2}{2} M\tilde{B}V^2 \sum_{t=1}^T \rho_{t-1}^2 \\
&+ \frac{\eta\alpha}{2} \sum_{t=1}^T \rho_{t-1} E[\|\sum_{m=1}^M \sum_{b=0}^{\tilde{B}-1} \nabla f(\hat{\mathbf{u}}_{t-\tau_{t,m},m,b})\|^2] \\
&\left(M\tilde{B}\eta^2 L^2 \rho_{t-1} D' \sum_{n=1}^{D'} \rho_{t+n} - \frac{1}{M\tilde{B}} \right)
\end{aligned} \tag{27}$$

Assuming that $\forall t \in \{1, 2, \dots\}$:

$$M^2 \tilde{B}^2 \eta^2 L^2 \rho_{t-1} D' \sum_{n=1}^{D'} \rho_{t+n} \leq 1 \tag{28}$$

We have Theorem 1;

$$\begin{aligned}
\sum_{t=1}^T \rho_{t-1} E[\|\nabla f(\mathbf{v}_{t-1})\|^2] &\leq \frac{2(f(\mathbf{v}_0) - f(\mathbf{v}_*))}{M\tilde{B}\eta\alpha} + \frac{\eta^2 L^2}{\tilde{B}} \sum_{t=1}^T \rho_{t-1} \left[V^2 \left(\sum_{b=0}^{\tilde{B}-1} \frac{\mu(\mu^b - 1)}{\mu - 1} + \tilde{B} \frac{\mu(\mu^D - 1)}{\mu - 1} \right) \right. \\
&+ \left. M\tilde{B}^2 \sigma^2 \sum_{j=t-1-D'}^{t-2} \rho_{j-1}^2 \right] + \frac{L\eta V^2}{\alpha} \sum_{t=1}^T \rho_{t-1}^2
\end{aligned} \tag{29}$$

□

Proof to Corollary 1 According to Theorem 1, we have:

$$\frac{1}{\sum_{t=1}^T \rho_{t-1}} \sum_{t=1}^T \rho_{t-1} E[|\nabla f(\mathbf{v}_{t-1})|^2] \leq \frac{2(f(\mathbf{v}_0) - f(\mathbf{v}_*))}{M\tilde{B}\eta\alpha \sum_{t=1}^T \rho_{t-1}} + \frac{\eta^2 L^2}{\tilde{B} \sum_{t=1}^T \rho_{t-1}} \sum_{t=1}^T \rho_{t-1} \left[V^2 \left(\sum_{b=0}^{\tilde{B}-1} \frac{\mu(\mu^b - 1)}{\mu - 1} + \tilde{B} \frac{\mu(\mu^D - 1)}{\mu - 1} \right) + M\tilde{B}^2 \sigma^2 \sum_{j=t-1-D'}^{t-2} \rho_{j-1}^2 \right] + \frac{L\eta V^2}{\alpha \sum_{t=1}^T \rho_{t-1}} \sum_{t=1}^T \rho_{t-1}^2 \tag{30}$$

By setting the learning rate ρ_t to be a constant and equal to η , we obtain the following:

$$\frac{1}{T} \sum_{t=1}^T E[|\nabla f(\mathbf{v}_{t-1})|^2] \leq \frac{2(f(\mathbf{v}_0) - E[f(\mathbf{v}_*)])}{TM\tilde{B}\eta^2\alpha} + \eta^2 LV^2 \left(\frac{1}{\alpha} - \frac{L\mu(1 - \mu^{\tilde{B}})}{(1 - \mu)^2 \tilde{B}} + \frac{L\mu}{(1 - \mu)} + L \frac{\mu(1 - \mu^D)}{1 - \mu} \right) + \eta^4 L^2 M\tilde{B} \sigma^2 D' \tag{31}$$

The inequalities shown in Theorem 1 can be re-arranged as follows:

$$\eta^4 \leq \frac{1}{M^2 \tilde{B}^2 L^2 D'^2} \tag{32}$$

Assuming $\eta \leq 1$ and $0 < \mu < 1$,

$$\eta^4 \leq \frac{(\mu - 1)^8}{(\mu(\mu - 1) + 9L^2\mu(D + 1)(\mu^{D+1} - 1))^4} \tag{33}$$

By setting the learning rate as follows:

$$\eta^2 = \frac{\sqrt{(f(\mathbf{v}_0) - f(\mathbf{v}_*))}}{A\alpha\sqrt{TM\tilde{B}}} \tag{34}$$

where A is a constant:

$$A = LV^2 \left(\frac{1}{\alpha} + \frac{1}{\alpha^2} + \frac{2L\mu}{(1 - \mu)\alpha} \right)$$

We end up with the following bound on the delayed parameters that if holds, Theorem 1 is satisfied:

$$T \geq \frac{M\tilde{B}L^2D'^2(f(\mathbf{v}_0) - f(\mathbf{v}_*))}{A^2\alpha^2} \tag{35}$$

$$T \geq \frac{(f(\mathbf{v}_0) - f(\mathbf{v}_*))(\mu(\mu - 1) + 9L^2\mu(D + 1)(\mu^{D+1} - 1))^4}{M\tilde{B}A^2\alpha^2(\mu - 1)^8} \tag{36}$$

Hence, from Theorem 1, we obtain:

$$\begin{aligned}
\frac{1}{T} \sum_{t=1}^T E[|\|\nabla f(\mathbf{v}_{t-1})\|^2] &\leq \frac{2(f(\mathbf{v}_0) - f(\mathbf{v}_*))}{TM\tilde{B}\eta^2\alpha} + \eta^2 LV^2 \left(\frac{1}{\alpha} + \frac{2L\mu}{(1-\mu)} \right) + \eta^2 L\sigma^2 \\
&\leq \frac{2(f(\mathbf{v}_0) - f(\mathbf{v}_*))}{TM\tilde{B}\eta^2\alpha} + \eta^2 \alpha A \\
&= 3A \sqrt{\frac{f(\mathbf{v}_0) - f(\mathbf{v}_*)}{TM\tilde{B}}}
\end{aligned} \tag{37}$$

Therefore, Corollary 1 has been proven. \square

7 Asynchronous Distributed Lock-free Parallel Stochastic Variational Inference

In this section, we describe our proposed distributed parallel implementation of the asynchronous lock-free stochastic variational inference algorithm (DPSVI) on a *computer cluster* with multi-core nodes. The steps of the algorithm follow from the proposed DPSGD but in the context of VI. First, we derive the model family applicable with DPSVI and review SVI following the same steps in Hoffman et al. (2013). Then, we derive DPSVI from DPSGD.

Model family The family of models considered here consists of three random variables: observations $\mathbf{x} = \mathbf{x}_{1:n}$, local hidden variables $\mathbf{z} = \mathbf{z}_{1:n}$, global hidden variables $\boldsymbol{\beta}$ and fixed parameters $\boldsymbol{\zeta}$. The model assumes that the distribution of the n pairs of $(\mathbf{x}_i, \mathbf{z}_i)$ is conditionally independent given $\boldsymbol{\beta}$. Furthermore, their distribution and the prior distribution of $\boldsymbol{\beta}$ are in the exponential family.

$$p(\boldsymbol{\beta}, \mathbf{x}, \mathbf{z} | \boldsymbol{\zeta}) = p(\boldsymbol{\beta} | \boldsymbol{\zeta}) \prod_{i=1}^n p(\mathbf{z}_i, \mathbf{x}_i | \boldsymbol{\beta}) \tag{38}$$

$$p(\mathbf{z}_i, \mathbf{x}_i | \boldsymbol{\beta}) = h(\mathbf{x}_i, \mathbf{z}_i) \exp(\boldsymbol{\beta}^T t(\mathbf{x}_i, \mathbf{z}_i) - a(\boldsymbol{\beta})) \tag{39}$$

$$p(\boldsymbol{\beta} | \boldsymbol{\zeta}) = h(\boldsymbol{\beta}) \exp(\boldsymbol{\zeta}^T t(\boldsymbol{\beta}) - a(\boldsymbol{\zeta})) \tag{40}$$

Here, we overload the notation for the base measures $h(\cdot)$, sufficient statistics $t(\cdot)$ and log normaliser $a(\cdot)$. While the proposed approach is generic, for simplicity we assume a conjugacy relationship between $(\mathbf{x}_i, \mathbf{z}_i)$ and $\boldsymbol{\beta}$. That is, the distribution $p(\boldsymbol{\beta} | \mathbf{x}, \mathbf{z})$ is in the same family as the prior $p(\boldsymbol{\beta} | \boldsymbol{\zeta})$. Note that this family of models includes, but is not limited to, latent Dirichlet allocation, Bayesian Gaussian mixture, probabilistic matrix factorisation, hidden Markov models, hierarchical linear and probit regression and many Bayesian non-parametric models.

Mean-field variational inference Variational inference (VI) approximates intractable posterior $p(\boldsymbol{\beta}, \mathbf{z} | \mathbf{x})$ by positing a family of simple distributions $q(\boldsymbol{\beta}, \mathbf{z})$ and find the

member of the family that is closest to the posterior (closeness is measured with KL divergence). The resulting optimisation problem is equivalent to maximising the evidence lower bound (ELBO).

$$\mathcal{L}(q) = E_q[\log p(\mathbf{x}, \mathbf{z}, \boldsymbol{\beta})] - E_q[\log p(\mathbf{z}|\boldsymbol{\beta})] \leq \log p(\mathbf{x}) \tag{41}$$

Mean-field is the simplest family of distribution, where the distribution over the hidden variables factorises as follows:

$$q(\boldsymbol{\beta}, \mathbf{z}) = q(\boldsymbol{\beta}|\lambda) \prod_{i=1}^n p(z_i|\boldsymbol{\phi}_i) \tag{42}$$

Further, each variational distribution is assumed to come from the same family of the true one. Mean-field variational inference optimises the new ELBO with respect to the local and global variational parameters $\boldsymbol{\phi}$ and λ .

$$\mathcal{L}(\lambda, \boldsymbol{\phi}) = E_q \left[\log \frac{p(\boldsymbol{\beta})}{q(\boldsymbol{\beta})} \right] + \sum_{i=1}^n E_q \left[\log \frac{p(\mathbf{x}_i, \mathbf{z}_i|\boldsymbol{\beta})}{q(\mathbf{z}_i)} \right] \tag{43}$$

It iteratively updates each variational parameter holding the other parameters fixed. With the assumptions taken so far, each update has a closed form solution. The local parameters are a function of the global parameters.

$$\boldsymbol{\phi}(\lambda_t) = \arg \max_{\boldsymbol{\phi}} \mathcal{L}(\lambda_t, \boldsymbol{\phi}) \tag{44}$$

We are interested in the global parameters which summarise the whole dataset (clusters in the Bayesian Gaussian mixture, topics in LDA).

$$\mathcal{L}(\lambda) = \max_{\boldsymbol{\phi}} \mathcal{L}(\lambda, \boldsymbol{\phi}) \tag{45}$$

To find the optimal value of λ given that $\boldsymbol{\phi}$ is fixed, we compute the natural gradient of $\mathcal{L}(\lambda)$ and set it to zero by setting:

$$\lambda^* = \zeta + \sum_{i=1}^n E_{\boldsymbol{\phi}(\lambda_t)}[t(\mathbf{x}_i, \mathbf{z}_i)] \tag{46}$$

Thus, the new optimal global parameters are $\lambda_{t+1} = \lambda^*$. The algorithm works by iterating between computing the optimal local parameters given the global ones (Eq. 44) and computing the optimal global parameters given the local ones (Eq. 46).

Stochastic variational inference. Instead of analysing all of the data to compute λ^* at each iteration, stochastic optimisation can be used. Assuming that the data samples are uniformly randomly selected from the dataset, an unbiased noisy estimator of $\mathcal{L}(\lambda, \boldsymbol{\phi})$ can be developed based on a single data point:

$$\mathcal{L}_i(\lambda, \boldsymbol{\phi}_i) = E_q \left[\log \frac{p(\boldsymbol{\beta})}{q(\boldsymbol{\beta})} \right] + nE_q \left[\log \frac{p(\mathbf{x}_i, \mathbf{z}_i|\boldsymbol{\beta})}{q(\mathbf{z}_i)} \right] \tag{47}$$

The unbiased stochastic approximation of the ELBO as a function of λ can be written as follows:

$$\mathcal{L}_i(\lambda) = \max_{\phi_i} \mathcal{L}_i(\lambda, \phi_i) \quad (48)$$

Following the same steps in the previous section, we end up with a noisy unbiased estimate of Eq. 45:

$$\hat{\lambda} = \zeta + nE_{\phi_i(\lambda_t)}[t(\mathbf{x}_i, \mathbf{b}z_i)] \quad (49)$$

At each iteration, we move the global parameters a step-size ρ_t (learning rate) in the direction of the noisy natural gradient:

$$\lambda_{t+1} = (1 - \rho_t)\lambda_t + \rho_t\hat{\lambda} \quad (50)$$

With certain conditions on ρ_t , the algorithm converges ($\sum_{t=1}^{\infty} \rho_t = \infty$, $\sum_{t=1}^{\infty} \rho_t^2 < \infty$) (Robbins & Monro, 1951).

Based on a batch of data points, the unbiased noisy estimator of $\mathcal{L}(\lambda, \phi)$ can be written as follows:

$$\mathcal{L}_g(\lambda, \phi_g) = E_q \left[\log \frac{p(\beta)}{q(\beta)} \right] + \frac{n}{G} \sum_{i \in G_g} E_q \left[\log \frac{p(\mathbf{x}_i, z_i | \beta)}{q(z_i)} \right] \quad (51)$$

where $G_g = \{(g-1)G+1, \dots, gG\}$. Equation 49 can be written as follows:

$$\hat{\lambda}_g = \zeta + \frac{n}{G} \sum_{i \in G_g} E_{\phi_i(\lambda_t)}[t(\mathbf{x}_i, z_i)] \quad (52)$$

Following this derivation, DPSVI, Algorithm 3 and 4, can be simply obtained from DPSGD by amending line (7) and (8) of Algorithm 2. Specifically, in line (7) we compute the local variational parameters $\phi_i(\lambda^*)$ corresponding to the data point \mathbf{x}_i and the global variational parameter λ^* , $\phi_i(\lambda^*) = \arg \max_{\phi} \mathcal{L}_i(\lambda^*, \phi)$. The index i is randomly picked from $\{1, \dots, n\}$ and λ^* replaces $\hat{\mu}$. In line (8) we perform the update after replacing the stochastic gradient $\nabla f_i(\hat{\mu})$ by negative stochastic natural gradient with respect to the global parameter λ^* , $\mathbf{g}_i(\lambda^*) = -\left(\zeta + \frac{n}{G} \sum_{j \in G_i} E_{\phi_j(\lambda^*)}[t(\mathbf{x}_j, z_j)] - \lambda^*\right)$.

Finally, the derivation of LDA from the presented model family can be found in Mohamad et al. (2018).

Algorithm 3 DPSVI-Master: Updates performed at the master machine

- 1: **initialise:** number of iteration T , global variable \mathbf{v} , global learning rate $\{\rho_t\}_{t=0, \dots, T-1}$
 - 2: **for** $t = 0, 1, 2, \dots, T-1$ **do**
 - 3: Collect M updating vectors $\mathbf{w}_1, \dots, \mathbf{w}_M$ from the workers.
 - 4: Update the current estimate of the global parameter $\mathbf{v} \leftarrow \mathbf{v} + \rho_t \sum_m \mathbf{w}_m$
 - 5: $t \leftarrow t + 1$
 - 6: **end for**
-

Algorithm 4 DPVI-Worker: Updates performed at each worker machine

```

1: initialise: number of iterations of per-worker loop  $B$ , learning rate  $\eta$ , number of threads  $p$ 
2: while (MasterIsRun) do
3:   Pull a global parameter  $\boldsymbol{v}$  from the master and put it in the shared memory.
4:   Fork  $p$  threads
5:   for  $c=0$  to  $B-1$  do
6:     Read current values of  $\boldsymbol{\lambda}^*$ , denoted as  $\hat{\boldsymbol{\lambda}}^*$ , from the shared memory.
7:     Randomly pick  $i$  from  $\{1, \dots, n\}$ 
8:     Compute the local variational parameters  $\phi_i^*(\hat{\boldsymbol{\lambda}}^*)$  corresponding to the data point  $\boldsymbol{x}_i$  and the
       global variational parameter  $\hat{\boldsymbol{\lambda}}^*$ ,  $\phi_i^*(\hat{\boldsymbol{\lambda}}^*) = \arg \max_{\phi_i} \mathcal{L}_i(\hat{\boldsymbol{\lambda}}^*, \phi_i)$ 
9:     Compute the stochastic natural gradient with respect to the global parameter  $\boldsymbol{\lambda}$ ,  $\boldsymbol{g}_i(\boldsymbol{\lambda}, \phi) =$ 
       
$$-\left( \zeta + \frac{\eta}{G} \sum_{j \in G_i} E_{\phi_j(\boldsymbol{\lambda})} [t(\boldsymbol{x}_j, \boldsymbol{z}_j)] - \boldsymbol{\lambda} \right)$$

10:    Update the local version of the global variational parameter  $\boldsymbol{\lambda}^* \leftarrow \boldsymbol{\lambda}^* + \eta \boldsymbol{g}_i(\hat{\boldsymbol{\lambda}}^*, \phi^*)$ 
11:    end for
12:    Push the update vector  $\boldsymbol{\lambda}^* - \boldsymbol{v}$  from the shared memory to the master
13: end while

```

8 Highly scalable advantage actor critics

The impressive results that have been achieved by deep artificial neural networks in several application domains are often driven by the availability of very large training data sets (Krizhevsky et al., 2012). In reinforcement learning (RL) (Sutton & Barto, 2018), an agent learns how to behave by interacting with its environment and has to experiment by trial and error, over and over again, often accumulating millions of repeated experiences. In order to enable learning in complex, real-world environments, recent advances in RL have successfully incorporated function approximation through deep networks resulting in deep reinforcement learning (DRL) (Sutton & Barto, 2018). The already data hungry DL function is then aggravated by the data inefficiency of RL motivating the development of more scalable learning algorithms.

Policy gradient methods directly maximise the expected rewards of a parameterised policy using gradient-based iterative methods such as the Stochastic Gradient Descent method (SGD). Advantage actor-critic incorporates control variate techniques to reduce the variance of the approximated gradient (Sutton et al., 2000). Thus, different versions of SGD can be used directly for learning. Various SGD implementations have been exploited to scale up DRL including distributed (Agarwal & Duchi, 2011; Lian et al., 2015) and parallel algorithms (Recht et al., 2011; Zhao et al., 2017) resulting in different scalable DRL algorithms (Ong et al., 2015; Nair et al., 2015; Adamski et al., 2018; Mnih et al., 2016; Babaeizadeh et al., 2016; Clemente et al., 2017; Horgan et al., 2018). These developments are particularly relevant to DRL due to its inherently sequential nature and the massive amount of data required for learning complex tasks such as playing games (Silver et al., 2016), controlling robots (Abbeel et al., 2007), optimising memory control (Ipek et al., 2008), and personalising web services (Theodorou et al., 2015), amongst others.

More recently, a few hybrid DRL algorithms have been proposed that combine both aspects of parallel (i.e. shared) and distributed computation (i.e., distributed memory) (Adamski et al., 2018; Babaeizadeh et al., 2016). A typical problem of distributed learning is the communication overhead arising from the necessity to share the weight

updates between nodes. Traditionally, large batches and step sizes have been used to curb the communication while preserving scalability. These strategies introduce a trade-off between computational and sample efficiency: a large batch increases the time needed to calculate the gradient locally, but decreases its variance allowing higher learning rate to be used. However, there is a limit on the speed-up that can be achieved by tuning the learning rate and batch sizes (Li et al., 2014b; Bottou et al., 2018).

Here, we apply DPSGD to steer off communication computation in a hybrid distributed-parallel implementation of DRL, with initial focus on advantage actor critic (A2C), which has been extensively studied (Ong et al., 2015; Adamski et al., 2018; Mnih et al., 2016; Babaeizadeh et al., 2016; Clemente et al., 2017; Horgan et al., 2018). We use a cluster network with a given number of computational nodes equipped with multiple CPUs. Each node maintains a copy of the actor and critic's models; for instance, the weights of the corresponding neural network implementing those models. A few of these nodes act as a *master* and the remaining nodes are the *workers*. The models maintained by the master are called global models. Within each worker node, the A2C's models are shared locally among its multiple CPUs, and are denoted as *local*.

Each worker performs multiple lock-free parallel updates (Zhao & Li, 2016) for the models shared across the CPUs. The global model is then updated by the master using the asynchronously aggregated local multiple steps updates. This simple strategy yields a highly-scalable advantage actor-critics (HSA2C), and harnesses distributed and local computation and storage resources. By updating the local variables multiple times, HSA2C mitigates the communication cost converting time speed-up.

8.1 Background

In reinforcement learning (RL), an agent interacts sequentially with an environment, with the goal of maximising cumulative rewards. At each step t the agent observes a state s_t , selects an action a_t according to its policy $\pi(a_t|s_t)$, and receives the next state s_{t+1} along with a reward r_t . This continues until the agent reaches a terminal state at $t = T$. The cumulative rewards, called return, at time t can be then written as $R_t = \sum_{i=0}^{\infty} \gamma^i r_{t+i}$, where the goal is to learn a policy that maximises the expected return from each state s_t : $E[R_t|s_t = s]$. The action value function of policy π , $Q^\pi(s, a) = E[R_t|s_t = s, a_t = a, \pi]$ is the expected return for taking action a in state s and following policy π . The value function $V^\pi(s) = E[R_t|s_t = s]$ is the expected return of policy π from state s .

Two main approaches of RL have been studied: value-based and policy based RL. In value-based RL, the policy is inferred from the value function which is represented by a function approximator such as a neural network. Hence, the value function can be written as $Q(s, a; w)$, where w represents the approximator parameters. Then, the goal of the RL algorithms is to iteratively update w to find the optimal action value function representing the optimal policy π . Alternatively, policy based DRL directly parameterises the policy $\pi(a|s, \theta)$ and updates its parameter by performing, typically approximate, gradient ascent on

$$L(\theta) = E[Q^{\pi_\theta}(a, s)|\theta]$$

e.g. see (Williams, 1992). Hence, the gradient of the objective function can be expressed as follows:

$$\nabla L(\theta) = E[\nabla_{\theta} \log \pi(\mathbf{a}|\mathbf{s}, \theta) Q^{\pi_{\theta}}(\mathbf{a}, \mathbf{s}) | \theta] \quad (53)$$

An unbiased estimate of the gradient in Eq. 53 can be obtained by computing the update from randomly sampled tuples of form $(\mathbf{s}, \mathbf{a}, r)$.

8.1.1 Advantage actor critic (A2C) methods

To reduce the variance of the estimate in Eq. (53) while keeping it unbiased, the Q function can be replaced with an advantage function, $A^{\pi_{\theta}}(\mathbf{a}, \mathbf{s}) = (Q^{\pi_{\theta}}(\mathbf{a}, \mathbf{s}) - V^{\pi_{\theta}}(\mathbf{s}))$ (Sutton et al., 2000) where V is the state value function. This approach can be viewed as an actor-critic architecture where the policy is the actor and the advantage function is the critic (Sutton & Barto, 1998).

Deep neural networks are used to approximate the actor and critics functions. Typically, two neural networks are deployed, one parameterised by θ approximating the actor and the other parameterised by θ_v approximating the critic. Hence, the gradient of the objective function can be expressed as follows:

$$\nabla L(\theta, \theta_v) = \begin{pmatrix} \nabla_{\theta} L(\theta, \theta_v) \\ \nabla_{\theta_v} L(\theta, \theta_v) \end{pmatrix} \quad (54)$$

where

$$\begin{aligned} \nabla_{\theta} L(\theta, \theta_v) &= E[\nabla_{\theta} \log \pi(\mathbf{a}|\mathbf{s}, \theta) A^{\theta_v}(\mathbf{a}, \mathbf{s}) | \theta] \\ \nabla_{\theta_v} L(\theta, \theta_v) &= E[\nabla_{\theta_v} A^{\theta_v}(\mathbf{a}, \mathbf{s}) | \theta] \end{aligned} \quad (55)$$

In the next section we provide a brief overview of existing scalable SGD approaches and how they have been adopted to scale-up A2C.

8.1.2 Scalable SGD algorithms for DRL

Stochastic gradient descent (SGD) and its variants are used to optimise the A2C objective function (Eq. (54)) (Sutton & Barto, 2018). SGD updates the actor and the critics network weights (θ, θ_v) based on the approximate gradient of the objective function $L(\theta, \theta_v)$ computed using an experience trajectory (or batch of trajectories) sampled using policy π_{θ} . Various scalable SGD-based approaches have recently been proposed to scale up DRL algorithms (Ong et al., 2015; Nair et al., 2015; Mnih et al., 2016; Babaeizadeh et al., 2016; Clemente et al., 2017; Horgan et al., 2018; Adamski et al., 2018). In general, these approaches can be viewed as derivations from either distributed SGD (DSGD) or parallel SGD (PSGD).

A DSGD-like architecture called Gorila, proposed by Nair et al. (2015), relies on asynchronous training of RL agents in a distributed setting. Gorila distributes Deep Q-Network DQN (Mnih et al., 2015) across multiple machines. Each machine runs an actor that interacts with the environment, samples data from the replay memory and computes the gradients of the DQN loss with respect to the policy parameters. The gradients are asynchronously sent to a central parameter server which updates a central copy of the model. The updated policy parameters are sent to the actor-learners at fixed intervals. In Mnih et al. (2016), PSGD-like parameter updates and data generation have been used within a

single-machine, in a multi-threaded rather than a distributed context. The shared parameter is then updated in an asynchronous DSGD-like fashion.

Algorithm 5 PA3C

```

1: initialise: number of threads  $p$ , thread step counter  $t \leftarrow 0$ , iteration counter  $T \leftarrow 0$ , thread gradient
   accumulator  $\mathbf{g} \leftarrow 0$  and  $\mathbf{g}_v \leftarrow 0$ , shared update vector accumulator  $\mathbf{G} \leftarrow 0$  and  $\mathbf{G}_v \leftarrow 0$ , mini-batch
   size  $m$ , mini-batch counter  $n$ , shared parameter vector  $\theta$  and  $\theta_v$ , learning rate.
2: Fork  $p$  threads
3: repeat
4:    $t_{start} = t$ 
5:   Get state  $\mathbf{s}_t$ 
6:   repeat
7:     Read current values of  $\theta$ , denoted as  $\hat{\theta}$ , from the shared memory.
8:     Take action  $\mathbf{a}_t$  according to policy  $\pi(\mathbf{a}|\mathbf{s}_t; \hat{\theta})$ 
9:     Receive new state  $\mathbf{s}_{t+1}$  and reward  $r_t$ 
10:     $t \leftarrow t + 1$ 
11:   until terminal  $\mathbf{s}_t$  or  $t - t_{start} == t_{max}$ 
12:   Read current values of  $\theta_v$ , denoted as  $\hat{\theta}_v$ , from the shared memory.
13:    $R = \begin{cases} 0 & \text{for terminal } \mathbf{s}_t \\ V(\mathbf{s}_t, \hat{\theta}_v) & \text{for non-terminal } \mathbf{s}_t \end{cases}$ 
14:   for  $i \in \{t - 1, \dots, t_{start}\}$  do
15:      $R \leftarrow r_i + \gamma R$ 
16:     Accumulate gradients wrt  $\hat{\theta}$ :  $\mathbf{g} \leftarrow \mathbf{g} + \nabla_{\hat{\theta}} \log \pi(\mathbf{a}_i | \mathbf{s}_i; \hat{\theta})(R - V(\mathbf{s}_i; \hat{\theta}_v))$ 
17:     Accumulate gradients wrt  $\hat{\theta}_v$ :  $\mathbf{g}_v \leftarrow \mathbf{g}_v + \frac{\partial (R - V(\mathbf{s}_i; \hat{\theta}_v))^2}{\partial \hat{\theta}_v}$ 
18:   end for
19:   Accumulate mini-batch gradients:  $\mathbf{G} \leftarrow \mathbf{G} + \mathbf{g}$  and  $\mathbf{G}_v \leftarrow \mathbf{G}_v + \mathbf{g}_v$ 
20:    $\mathbf{g} \leftarrow 0$ 
21:    $\mathbf{g}_v \leftarrow 0$ 
22:    $n \leftarrow n + 1$ 
23:   if  $n == m$  then
24:      $n = 0$ 
25:     Update the shared network weights  $\theta$  using  $\mathbf{G}$  and the learning rate
26:     Update the shared network weights  $\theta_v$  using  $\mathbf{G}_v$  and the learning rate
27:      $\mathbf{G} = 0$  and  $\mathbf{G}_v = 0$ 
28:     Update local steps:  $T \leftarrow T + 1$ 
29:   end if
30: until  $T > T_{max}$ 

```

Scaling-up A2C using this two architectures lead to the asynchronous distributed A2C (DA3C) in Algorithm 6 and Algorithm 7 and the lock-free parallel A2C (PA3C) in Algorithm 5. The lock-free PSGD presented in Algorithm 5 employs (Recht et al., 2011) style of updates for training A2C. The actors and critics networks are stored in a shared memory where different threads can update their parameters without any memory locking. Therefore, in theory, linear speed-ups with respect to the number of threads can be achieved. On the other hand, the asynchronous DSGD presented in Algorithm 6 and Algorithm 7 employs (Agarwal & Duchi, 2011) style of updates for training A2C. Here, the algorithms are deployed on a cluster where a master machine maintains a copy of A2C's networks (the global networks (θ, θ_v) , see Algorithm 6). Other machines serve as workers which independently and simultaneously compute the local stochastic gradients of a copy of A2C's networks (the local networks, see Algorithm 7). The workers only communicate with the master to exchange information in which they access the state of the global networks and

provide the master with the stochastic gradients. The master aggregates predefined amounts of gradients from the workers. Then, it updates its global networks. Note that the local distributed computations are done in an asynchronous style where the workers are not locked until the master starts updating the global networks. That is, the workers might compute some stochastic gradients based on early value of the global networks.

Algorithm 6 DA3C-Master

- 1: **initialise:** global network parameters θ and θ_v , batch size M , global learning rate
 - 2: **for** $t = 0, 1, 2, \dots, T - 1$ **do**
 - 3: Collect M updating vectors $(\mathbf{w}_1, \mathbf{w}_{v1}), \dots, (\mathbf{w}_M, \mathbf{w}_{vM})$ from the workers.
 - 4: Update the current estimate of the global parameters (θ, θ_v) using $\sum_{m=1}^M (\mathbf{w}_m, \mathbf{w}_{vm})$
 - 5: **end for**
-

Algorithm 7 DA3C-Worker

- 1: **initialise:** batch counter $n \leftarrow 0$, counter $t \leftarrow 0$, gradient accumulator $\mathbf{g} \leftarrow 0$ and $\mathbf{g}_v \leftarrow 0$, gradient vector $\mathbf{G} \leftarrow 0$ and $\mathbf{G}_v \leftarrow 0$, mini-batch size m , local networks parameters θ and θ_v
 - 2: **repeat**
 - 3: Pull global networks from the master and assign them to (θ, θ_v)
 - 4: $t_{start} = t$
 - 5: Get state \mathbf{s}_t
 - 6: **repeat**
 - 7: Take action \mathbf{a}_t according to policy $\pi(\mathbf{a}|\mathbf{s}_t; \theta)$
 - 8: Receive new state \mathbf{s}_{t+1} and reward r_t
 - 9: $t \leftarrow t + 1$
 - 10: **until** terminal \mathbf{s}_t or $t - t_{start} == t_{max}$
 - 11:
$$R = \begin{cases} 0 & \text{for terminal } \mathbf{s}_t \\ V(\mathbf{s}_t, \hat{\theta}_v) & \text{for non-terminal } \mathbf{s}_t \end{cases}$$
 - 12: **for** $i \in \{t - 1, \dots, t_{start}\}$ **do**
 - 13: $R \leftarrow r_i + \gamma R$
 - 14: Accumulate gradients wrt θ : $\mathbf{g} \leftarrow \mathbf{g} + \nabla_{\theta} \log \pi(\mathbf{a}_i | \mathbf{s}_i; \theta) (R - V(\mathbf{s}_i; \theta_v))$
 - 15: Accumulate gradients wrt θ_v : $\mathbf{g}_v \leftarrow \mathbf{g}_v + \frac{\partial (R - V(\mathbf{s}_i; \theta_v))^2}{\partial \theta_v}$
 - 16: **end for**
 - 17: Accumulate mini-batch gradients: $\mathbf{G} \leftarrow \mathbf{G} + \mathbf{g}$ and $\mathbf{G}_v \leftarrow \mathbf{G}_v + \mathbf{g}_v$
 - 18: $\mathbf{g} \leftarrow 0$
 - 19: $\mathbf{g}_v \leftarrow 0$
 - 20: $n \leftarrow n + 1$
 - 21: **if** $n == m$ **then**
 - 22: $n = 0$
 - 23: Push the gradient vectors \mathbf{G} and \mathbf{G}_v to the master
 - 24: $\mathbf{G} = 0$ and $\mathbf{G}_v = 0$
 - 25: **end if**
 - 26: **until** Master Stop
-

Inspired by these two approaches, different improvements have been proposed. In Babaeizadeh et al. (2016), a hybrid CPU/GPU version of the Asynchronous Advantage

Actor Critic (A3C) algorithm (Mnih et al., 2016) was introduced. This study focused on mitigating the severe under-utilization of the GPU computational resources in DRL caused by its sequential nature of data generation. In this work, each agent queues policy requests in a *Prediction Queue* before each action, and periodically submits a batch of reward experiences to a *Training Queue*. Thus, unlike (Mnih et al., 2016), the agents do not compute the gradients themselves. Instead, they send experiences to central learners that update the network on the GPU accordingly. Such architecture reduces the GPU idle during the training. However, as the number of core increases, the learner becomes unable to cope with the data. Furthermore, such amount of data requires large storage capacity. Besides, the internal communications can affect the speed-up when bandwidth reach its ceiling. These issues limit the scalability of Babaeizadeh et al. (2016). The approach is also limited to Off-policy DRL methods. We also note that a similar way for parallelisation of DRL is proposed by Clemente et al. (2017).

Similarly, Horgan et al. (2018) proposes to generate experience data in parallel using multi-cores CPUs where experiences are accumulated by actors in a shared experience replay memory. Each actor interacts with the environment takes actions, gets rewards and states. The learner, then, replays samples of experience and updates the shared neural network. The architecture relies on prioritised experience replay (Schaul et al., 2015) to focus only on the most significant data generated by the actors. Along the same trend, in Espeholt et al. (2018), all the data (states, actions and rewards) is also accumulated by distributed actors and communicated to a centralised learner where the computation takes place. To correct the policy-lag between the learner and actors generating the data, these authors introduce the V-trace off-policy actor-critic algorithm. It is an importance weighting technique used to learn the target policy (learner policy) from behaviour ones (actors policy). The architecture of these studies (Horgan et al., 2018; Espeholt et al., 2018) allows distributing the generation and selection of experience data instead of distributing locally computed gradients as in Nair et al. (2015). Hence, it requires sending large size information over the network in case of large size batch of data making the communications more problematic. Furthermore, the central learner has to perform most of the computation which limits the scalability. Moreover, these approach is limited to Off-policy DRL methods.

The work in Adamski et al. (2018) presents an SGD-based hybrid distributed-parallel actor critic approach, and is the most closely related to our approach. These authors combine PA3C with DA3C to allow parallel distributed implementation of A3C on a computer cluster with multi-core node. Each node applies (Babaeizadeh et al., 2016) to queue data in batches, which are used to compute local gradients. These gradients are gathered from all workers, averaged and applied to update the global network parameters. To reduce the communication overhead, a careful reexamination of Adam optimizer's hyper-parameters is carried out allowing large batch sizes to be used.

8.2 Scalable actor-critic through multiple local updates

We propose a hybrid distributed-parallel actor critic algorithm whereby, unlike (Adamski et al., 2018), workers communicate update vectors rather than gradients to the master. These update vectors are the discrepancy between the initial copy of the advantage actor critic (A2C)'s local network parameters pulled from the master and the results of multiple

local iterations performed over the networks using PSGD-like updating style. As in DA3C-like algorithms, the master aggregates predefined amounts of these local updates, updates its global networks and broadcast the updated networks to the workers. Such multi-steps updating vectors allow large global update steps with less communication exchanges. The local A2C's actor model keeps interacting with the environment, taking actions and receiving rewards and next states. The data being collected is constantly used to compute the gradient of the objective function of A2C's actor and critic networks and updating these networks. The data collection and updates are done in PA3C-like style. Once a pre-defined number of updates is reached the local update vector is computed and pushed to the master in DA3C style preventing any locking or synchronisation (unlike (Adamski et al., 2018)) from halting the actor interaction with the environment.

One issue of this approach is that the variance of policy gradient DRL combinatorially increases with each step because of their Markovian dependency. Hence, there is a trade-off between variance and communication cost. Higher number of multi-step update results in higher variance, but reduces the communication cost. In our approach, the variance issue is slightly mitigated by the control variate technique used by the A2C algorithm. We also adopt a local mini-batch along with multi-step TD error to further reduce the variance. The variance can also be reduced when increasing the global batch size (i.e., the number of aggregated update vectors). Thus, adding more nodes reduces the variance allowing a higher number of multi-step updates. As the number of nodes increases, communication exchanges also increase. Therefore, our approach is beneficial when the number of nodes increases, as this involves increased communication, up to a certain breaking point after which the higher variance associated with higher number of multi-step updates counterbalances the benefits.

8.2.1 Overview of the algorithm

HSA2C is presented in Algorithm 8 and 9. On a single node agents work in parallel without any locking for the shared memory. Each agent interacts with an instance of the game environment producing experience data (state, action and reward) (line 8 to 10 in Algorithm 9). HSA2C uses k steps TD error along with the function approximation to allow better bias-variance trade-off (line 14 to 19 in Algorithm 9). Since the steps are taken in parallel by multiple independent actors, data correlation is reduced resulting in lower bias without the need for experience replay making the proposed approach applicable for both off-policy and on-policy DRL methods (Mnih et al., 2016). HSA2C uses local mini-batch (line 20 to line 24 in Algorithm 9) to reduce the stochastic gradient variance which is crucial to achieve the speed-up. That is HSA2C uses mini-batch of stochastic gradients, computed in parallel from different experiences, to estimate lower variance gradient. Such gradients allow more stable multiple updates (line 24 to 29 in Algorithm 9). Finally, HSA2C's workers send the computed multiple steps update vectors (in lock-free parallel) to the master (line 30 to 32 in Alg 9) which asynchronously collect these vectors and update its global parameters (Alg 8). The workers then pull the update global parameters and update their local parameters (line 33 to 34 in Alg 9).

The convergence rate for serial and synchronous parallel stochastic gradient (SG) is consistent with $O(1/\sqrt{T})$ (Ghadimi & Lan, 2013; Dekel et al., 2012; Nemirovski

et al., 2009). As per corollary 1, HSA2C achieves $O(1/\sqrt{T})$ convergence rate with almost $p * nw$ times less computation than A2C, where p is the number of threads and nw is the number of nodes. Hence, HSA2C achieves $p * nw$ iteration-speed-up. This iteration-speed-up converts to time-speed-up (real speed-up) provided that the communication between the nodes takes no time. The total communication time for HSA2C algorithms can be bounded as follows:

$$T * Tc \leq DTT \leq M * T * Tc \quad (56)$$

where T is the number of iterations, Tc is the communication time need for each master - worker exchange and m is the global batch (see Algorithm 8). For simplicity, we assume that the Tc is fixed and the same for all nodes. If the time needed for computing one update $Tu \leq Tc$, the total time needed by the distributed algorithm DTT could be higher than that of the sequential A2C SST :

$$T * Tc + SST/M \leq DTT \leq M * T * Tc + SST/nw \quad (57)$$

where $SST \leq T * Tc$, hence

$$SST < SST(1 + 1/M) \leq DTT \quad (58)$$

In such cases, exiting distributed algorithms like (Adamski et al., 2018) increase the local batch size so that Tu increases resulting in lower stochastic gradient variance and allowing for higher learning rate to be used, hence better convergence rate. This introduces a trade-off between computational efficiency and sample efficiency. Increasing the batch size by a factor of k increases the time needed for local computation by $O(k)$ and reduces the variance proportionally to $1/k$ (Bottou et al., 2018). Thus, a higher learning rate can be used. However, there is a limit on the size of the learning rate. This maximum learning speed can be improved on using HSA2C (Algorithm 8 and Algorithm 9) which performs B times less communication steps. These improvements are empirically studied in the main paper.

Algorithm 8 HSA2C-Master

- 1: **initialise:** global network weights θ and θ_v , batch size M , global learning rate
 - 2: **for** $t = 0, 1, 2, \dots, T - 1$ **do**
 - 3: Collect M updating vectors $(w_1, w_{v1}), \dots, (w_M, w_{vM})$ from the workers.
 - 4: Update the current estimate of the global parameters (θ, θ_v) using $\sum_{m=1}^M (w_m, w_{vm})$
 - 5: **end for**
-

Algorithm 9 HSA2C-Worker

```

1: initialise: number of iterations of per-worker loop  $B$ , number of threads  $p$ , thread step counter  $t \leftarrow 0$ ,
   shared counter  $T \leftarrow 0$ , thread gradient accumulator  $\mathbf{g} \leftarrow 0$  and  $\mathbf{g}_v \leftarrow 0$ , shared update vector
   accumulator  $\mathbf{G} \leftarrow 0$  and  $\mathbf{G}_v \leftarrow 0$ , mini-batch size  $m$ , mini-batch counter  $n$ , local shared parameter
   vector  $\boldsymbol{\theta}$  and  $\boldsymbol{\theta}_v$ , local learning rate.
2: Synchronise initial parameters  $\mathbf{u}_v \leftarrow \boldsymbol{\theta}_v$  and  $\mathbf{u} \leftarrow \boldsymbol{\theta}$ 
3: Fork  $p$  threads
4: while (MasterIsRun) do
5:    $t_{start} = t$ 
6:   Get state  $\mathbf{s}_t$ 
7:   repeat
8:     Read current values of  $\mathbf{u}$ , denoted as  $\hat{\mathbf{u}}$ , from the shared memory.
9:     Take action  $\mathbf{a}_t$  according to policy  $\pi(\mathbf{a}|\mathbf{s}_t; \hat{\mathbf{u}})$ 
10:    Receive new state  $\mathbf{s}_{t+1}$  and reward  $r_t$ 
11:     $t \leftarrow t + 1$ 
12:    until terminal  $\mathbf{s}_t$  or  $t - t_{start} == t_{max}$ 
13:    Read current values of  $\mathbf{u}_v$ , denoted as  $\hat{\mathbf{u}}_v$ , from the shared memory.
14:     $R = \begin{cases} 0 & \text{for terminal } \mathbf{s}_t \\ V(\mathbf{s}_t, \hat{\mathbf{u}}_v) & \text{for non-terminal } \mathbf{s}_t \end{cases}$ 
15:    for  $i \in \{t - 1, \dots, t_{start}\}$  do
16:       $R \leftarrow r_i + \gamma R$ 
17:      Accumulate gradients wrt  $\hat{\mathbf{u}}$ :  $\mathbf{g} \leftarrow \mathbf{g} + \nabla_{\hat{\mathbf{u}}} \log \pi(\mathbf{a}_i|\mathbf{s}_i; \hat{\mathbf{u}})(R - V(\mathbf{s}_i; \hat{\mathbf{u}}_v))$ 
18:      Accumulate gradients wrt  $\hat{\mathbf{u}}_v$ :  $\mathbf{g}_v \leftarrow \mathbf{g}_v + \frac{\partial(R - V(\mathbf{s}_i; \hat{\mathbf{u}}_v))^2}{\partial \hat{\mathbf{u}}_v}$ 
19:    end for
20:    Accumulate mini-batch gradients:  $\mathbf{G} \leftarrow \mathbf{G} + \mathbf{g}$  and  $\mathbf{G}_v \leftarrow \mathbf{G}_v + \mathbf{g}_v$ 
21:     $\mathbf{g} \leftarrow 0$ 
22:     $\mathbf{g}_v \leftarrow 0$ 
23:     $n \leftarrow n + 1$ 
24:    if  $n == m$  then
25:       $n = 0$ 
26:      Update the shared network weights  $\mathbf{u}$  using  $\mathbf{G}$  and the local learning rate
27:      Update the shared network weights  $\mathbf{u}_v$  using  $\mathbf{G}_v$  and the local learning rate
28:       $\mathbf{G} = 0$  and  $\mathbf{G}_v = 0$ 
29:      Update local steps:  $T \leftarrow T + 1$ 
30:      if  $T == B$  then
31:         $T = 0$ 
32:        Push the update vectors  $\mathbf{u} - \boldsymbol{\theta}$  and  $\mathbf{u}_v - \boldsymbol{\theta}_v$  to the master
33:        Pull global parameters  $\boldsymbol{\theta}$  and  $\boldsymbol{\theta}_v$  from the master
34:        Synchronise initial parameters  $\mathbf{u}_v \leftarrow \boldsymbol{\theta}_v$  and  $\mathbf{u} \leftarrow \boldsymbol{\theta}$ 
35:      end if
36:    end if
37: end while

```

Acknowledgements The authors thank the HPC team at Warwick University for providing the computational resources to run the experiments.

Author contributions SM developed and implemented the concepts. He wrote the draft and supported the subsequent revision. HB initiated the idea of the research and followed up its development. He worked on the manuscript throughout the publication process. HA contributed to the implementation of the reinforcement learning part of the manuscript.

Funding This work was partly funded through the European Horizon 2020 Framework Programme under Grant 687691 related to the Project: PROTEUS: Scalable Online Machine Learning for Predictive Analytics and Real-Time Interactive Visualization.

Data availability Not applicable.

Declarations

Conflict of interest None.

Ethics approval Not applicable.

Consent to participate Not applicable.

Consent for publication Not applicable.

Code availability Code could be made available at a later stage.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., et al. (2016). Tensorflow: A system for large-scale machine learning. *OSDI*, 16, 265–283.
- Abbeel, P., Coates, A., Quigley, M., & Ng, A. Y. (2007). An application of reinforcement learning to aerobatic helicopter flight. In *Advances in neural information processing systems* (pp. 1–8).
- Adamski, I., Adamski, R., Grel, T., Jędrzych, A., Kaczmarek, K., & Michalewski, H. (2018). Distributed deep reinforcement learning: Learn how to play atari games in 21 minutes. arXiv preprint [arXiv:1801.02852](https://arxiv.org/abs/1801.02852).
- Adamski, R., Grel, T., Klimek, M., & Michalewski, H. (2017). Atari games and intel processors. *Workshop on Computer Games* (pp. 1–18). Springer.
- Agarwal, A., & Duchi, J. C. (2011). Distributed delayed stochastic optimization. In *Neural Information Processing Systems*.
- Ba, J., Grosse, R., & Martens, J. (2016). Distributed second-order optimization using kronecker-factored approximations.
- Babaeizadeh, M., Frosio, I., Tyree, S., Clemons, J., & Kautz, J. (2016). Reinforcement learning through asynchronous advantage actor-critic on a gpu. arXiv preprint [arXiv:1611.06256](https://arxiv.org/abs/1611.06256).
- Bellemare, M. G., Naddaf, Y., Veness, J., & Bowling, M. (2013). The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47, 253–279.
- Blei, D. M., Kucukelbir, A., & McAuliffe, J. D. (2017). Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112, 859–877.
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of Machine Learning research*, 3, 993–1022.
- Bottou, L. (2010). Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010* (pp. 177–186). Springer.
- Bottou, L., Curtis, F. E., & Nocedal, J. (2018). Optimization methods for large-scale machine learning. *SIAM Review*, 60(2), 223–311.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., & Zaremba, W. (2016). Openai gym. arXiv preprint [arXiv:1606.01540](https://arxiv.org/abs/1606.01540).
- Chilimbi, T. M., Suzue, Y., Apacible, J., & Kalyanaraman, K. (2014). Project adam: Building an efficient and scalable deep learning training system. In *OSDI*.

- Clemente, A.V., Castejón, H.N., & Chandra, A. (2017). Efficient parallel methods for deep reinforcement learning. arXiv preprint [arXiv:1705.04862](https://arxiv.org/abs/1705.04862).
- Crane, R., & Roosta, F. (2019). Dingo: Distributed newton-type method for gradient-norm optimization. arXiv preprint [arXiv:1901.05134](https://arxiv.org/abs/1901.05134).
- De Sa, C., Zhang, C., Olukotun, K., & Ré, C. (2015). Taming the wild: A unified analysis of hogwild!-style algorithms. *Advances in Neural Information Processing Systems*, 28, 2656.
- Dean, J., Corrado, G., Monga, R., Chen, K., Devin, M., Mao, M., Senior, A., Tucker, P., Yang, K., & Le, Q.V., et al. (2012). Large scale distributed deep networks. In *Advances in neural information processing systems* (pp. 1223–1231).
- Dekel, O., Gilad-Bachrach, R., Shamir, O., & Xiao, L. (2012). Optimal distributed online prediction using mini-batches. *Journal of Machine Learning Research*, 13, 165–202.
- Duchi, J.C., Chaturapruek, S., & Ré, C. (2015). Asynchronous stochastic convex optimization. arXiv preprint [arXiv:1508.00882](https://arxiv.org/abs/1508.00882).
- Elgabli, A., Park, J., Bedi, A. S., Issaid, C. B., Bennis, M., & Aggarwal, V. (2020). Q-GADMM: Quantized group ADMM for communication efficient decentralized machine learning. *IEEE Transactions on Communications*, 69, 164–181.
- Espeholt, L., Soyer, H., Munos, R., Simonyan, K., Mnih, V., Ward, T., Doron, Y., Firoiu, V., Harley, T., & Dunning, I., et al. (2018). Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures. arXiv preprint [arXiv:1802.01561](https://arxiv.org/abs/1802.01561).
- Fang, C., & Lin, Z. (2017). Parallel asynchronous stochastic variance reduction for nonconvex optimization. In *AAAI*.
- Ghadimi, S., & Lan, G. (2013). Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23, 2341–2368.
- Hoffman, M. D., Blei, D. M., Wang, C., & Paisley, J. (2013). Stochastic variational inference. *Journal of Machine Learning Research*, 14, 1303–1347.
- Horgan, D., Quan, J., Budden, D., Barth-Maron, G., Hessel, M., van Hasselt, H., & Silver, D. (2018). Distributed prioritized experience replay. arXiv preprint [arXiv:1803.00933](https://arxiv.org/abs/1803.00933).
- Hsieh, C.-J., Yu, H.-F., & Dhillon, I. (2015). Passcode: Parallel asynchronous stochastic dual co-ordinate descent. In *International Conference on Machine Learning* (pp. 2370–2379). PMLR.
- Huo, Z., & Huang, H. (2016). Asynchronous stochastic gradient descent with variance reduction for non-convex optimization. arXiv preprint [arXiv:1604.03584](https://arxiv.org/abs/1604.03584).
- Huo, Z., & Huang, H. (2017). Asynchronous mini-batch gradient descent with variance reduction for non-convex optimization. In *AAAI*.
- Ipek, E., Mutlu, O., Martínez, J. F., & Caruana, R. (2008). Self-optimizing memory controllers: A reinforcement learning approach. *ACM SIGARCH Computer Architecture News*, 36, 39–50.
- Jahani, M., He, X., Ma, C., Mokhtari, A., Mudigere, D., Ribeiro, A., & Takáč, M. (2020a). Efficient distributed hessian free algorithm for large-scale empirical risk minimization via accumulating sample strategy. In *International Conference on Artificial Intelligence and Statistics* (pp. 2634–2644). PMLR.
- Jahani, M., Nazari, M., Rusakov, S., Berahas, A. S., & Takáč, M. (2020b). Scaling up quasi-newton algorithms: Communication efficient distributed sr1. In *International Conference on Machine Learning, Optimization, and Data Science* (pp. 41–54). Springer.
- Jordan, M. I., Ghahramani, Z., Jaakkola, T. S., & Saul, L. K. (1999). An introduction to variational methods for graphical models. *Machine Learning*, 37, 183–233.
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems* (pp. 1097–1105).
- Langford, J., Smola, A.J., & Zinkevich, M. (2009). Slow learners are fast. *Neural Information Processing Systems*.
- Leblond, R., Pedregosa, F., & Lacoste-Julien, S. (2017). Asaga: asynchronous parallel saga. In *Artificial Intelligence and Statistics* (pp. 46–54). PMLR.
- Li, M., Andersen, D. G., Park, J. W., Smola, A. J., Ahmed, A., Josifovski, V., Long, J., Shekita, E. J., & Su, B.-Y. (2014a). Scaling distributed machine learning with the parameter server. In *OSDI*.
- Li, M., Zhang, T., Chen, Y., & Smola, A. J. (2014b). Efficient mini-batch training for stochastic optimization. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 661–670). ACM.
- Lian, X., Huang, Y., Li, Y., & Liu, J. (2015). Asynchronous parallel stochastic gradient for nonconvex optimization. In *Neural Information Processing Systems*.
- Lian, X., Zhang, W., Zhang, C., & Liu, J. (2018). Asynchronous decentralized parallel stochastic gradient descent. In *International Conference on Machine Learning* (pp. 3043–3052). PMLR.
- Lichman, M. (2013). UCI machine learning repository.

- Lin, T., Stich, S. U., Patel, K. K., & Jaggi, M. (2018). Don't use large mini-batches, use local sgd. arXiv preprint [arXiv:1808.07217](https://arxiv.org/abs/1808.07217).
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., & Kavukcuoglu, K. (2016). Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning* (pp. 1928–1937).
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., & Riedmiller, M. (2013). Playing atari with deep reinforcement learning. arXiv preprint [arXiv:1312.5602](https://arxiv.org/abs/1312.5602).
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, *518*(7540), 529.
- Mohamad, S., Bouchachia, A., & Sayed-Mouchaweh, M. (2018). Asynchronous stochastic variational inference. arXiv preprint [arXiv:1801.04289](https://arxiv.org/abs/1801.04289).
- Nair, A., Srinivasan, P., Blackwell, S., Alcicek, C., Fearon, R., De Maria, A., Panneershelvam, V., Suleyman, M., Beattie, C., & Petersen, S., et al. (2015). Massively parallel methods for deep reinforcement learning. arXiv preprint [arXiv:1507.04296](https://arxiv.org/abs/1507.04296).
- Neiswanger, W., Wang, C., & Xing, E. (2015). Embarrassingly parallel variational inference in nonconjugate models. arXiv preprint [arXiv:1510.04163](https://arxiv.org/abs/1510.04163).
- Nemirovski, A., Juditsky, A., Lan, G., & Shapiro, A. (2009). Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, *19*(4), 1574–1609.
- Niu, F., Recht, B., Ré, C., & Wright, S. J. (2011). Hogwild!: A lock-free approach to parallelizing stochastic gradient descent. arXiv preprint [arXiv:1106.5730](https://arxiv.org/abs/1106.5730).
- Ong, H. Y., Chavez, K., & Hong, A. (2015). Distributed deep q-learning. arXiv preprint [arXiv:1508.04186](https://arxiv.org/abs/1508.04186).
- Paine, T., Jin, H., Yang, J., Lin, Z., & Huang, T. (2013). Gpu asynchronous stochastic gradient descent to speed up neural network training. arXiv preprint [arXiv:1312.6186](https://arxiv.org/abs/1312.6186).
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., & Lerer, A. (2017). Automatic differentiation in pytorch.
- Recht, B., Re, C., Wright, S., & Niu, F. (2011). Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *Neural Information Processing Systems*.
- Reddi, S. J., Hefny, A., Sra, S., Póczos, B., & Smola, A. (2015). On variance reduction in stochastic gradient descent and its asynchronous variants. arXiv preprint [arXiv:1506.06840](https://arxiv.org/abs/1506.06840).
- Robbins, H., & Monro, S. (1951). A stochastic approximation method. *The Annals of Mathematical Statistics*, *22*, 400–407.
- Ruder, S. (2016). An overview of gradient descent optimization algorithms. arXiv preprint [arXiv:1609.04747](https://arxiv.org/abs/1609.04747).
- Schaul, T., Quan, J., Antonoglou, I., & Silver, D. (2015). Prioritized experience replay. arXiv preprint [arXiv:1511.05952](https://arxiv.org/abs/1511.05952).
- Shamir, O., Srebro, N., & Zhang, T. (2014). Communication-efficient distributed optimization using an approximate newton-type method. In *International Conference on Machine Learning* (pp. 1000–1008). PMLR.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., et al. (2016). Mastering the game of go with deep neural networks and tree search. *Nature*, *529*(7587), 484–489.
- Stich, S. U. (2018). Local sgd converges fast and communicates little. arXiv preprint [arXiv:1805.09767](https://arxiv.org/abs/1805.09767).
- Stooke, A., & Abbeel, P. (2018). Accelerated methods for deep reinforcement learning. arXiv preprint [arXiv:1803.02811](https://arxiv.org/abs/1803.02811).
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction* (Vol. 1). MIT Press.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
- Sutton, R. S., McAllester, D. A., Singh, S. P., & Mansour, Y. (2000). Policy gradient methods for reinforcement learning with function approximation. *Advances in Neural Information Processing Systems*, *12*, 1057–1063.
- Theocharous, G., Thomas, P.S., & Ghavamzadeh, M. (2015). Personalized ad recommendation systems for life-time value optimization with guarantees. In *IJCAI* (pp. 1806–1812).
- Tsitsiklis, J., Bertsekas, D., & Athans, M. (1986). Distributed asynchronous deterministic and stochastic gradient optimization algorithms. *IEEE Transactions on Automatic Control*, *31*(9), 803–812.
- Wainwright, M.J., & Jordan, M.I., et al. (2008). Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*.
- Wang, J., Sahu, A. K., Yang, Z., Joshi, G., & Kar, S. (2019). Matcha: Speeding up decentralized sgd via matching decomposition sampling. In *2019 Sixth Indian Control Conference (ICC)* (pp. 299–300). IEEE.
- Williams, R.J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. In *Reinforcement Learning* (pp. 5–32). Springer.

- Xing, E. P., Ho, Q., Dai, W., Kim, J. K., Wei, J., Lee, S., Zheng, X., Xie, P., Kumar, A., & Yu, Y. (2015). Petuum: A new platform for distributed machine learning on big data. *IEEE Transactions on Big Data*.
- Yu, H., Yang, S., & Zhu, S. (2019). Parallel restarted SGD with faster convergence and less communication: Demystifying why model averaging works for deep learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33, 5693–5700.
- Zhang, S., Choromanska, A. E., & LeCun, Y. (2015). Deep learning with elastic averaging sgd. In *Neural Information Processing Systems*.
- Zhao, S.-Y., & Li, W.-J. (2016). Fast asynchronous parallel stochastic gradient descent: A lock-free approach with convergence guarantee. In *AAAI*.
- Zhao, S.-Y., Zhang, G.-D., & Li, W.-J. (2017). Lock-free optimization for non-convex problems. In *AAAI*.
- Zhou, F., & Cong, G. (2017). On the convergence properties of a k -step averaging stochastic gradient descent algorithm for nonconvex optimization. arXiv preprint [arXiv:1708.01012](https://arxiv.org/abs/1708.01012).
- Zinkevich, M., Weimer, M., Li, L., & Smola, A. J. (2010). Parallelized stochastic gradient descent. In *Neural Information Processing Systems*.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.