Distributed Deep Reinforcement Learning: learn how to play Atari games in 21 minutes[*]

Igor Adamski[1][**], Robert Adamski[2,4], Tomasz Grel[1], Adam Jędrych[1], Kamil Kaczmarek[1] and Henryk Michalewski[1,3]

[1] deepsense.ai
[2] Intel
[3] Polish Academy of Sciences
[4] Biz On Sp. z o.o.

Abstract. We present a study in Distributed Deep Reinforcement Learning (DDRL) focused on scalability of a state-of-the-art Deep Reinforcement Learning algorithm known as Batch Asynchronous Advantage Actor-Critic (BA3C). We show that using the Adam optimization algorithm with a batch size of up to 2048 is a viable choice for carrying out large scale machine learning computations. This, combined with careful reexamination of the optimizer's hyperparameters, using synchronous training on the node level (while keeping the local, single node part of the algorithm asynchronous) and minimizing the model's memory footprint, allowed us to achieve linear scaling for up to 64 CPU nodes. This corresponds to a training time of 21 minutes on 768 CPU cores, as opposed to the 10 hours required when using a single node with 24 cores achieved by a baseline single-node implementation.

Keywords: distributed computing, reinforcement learning, deep learning, Atari games, asynchronous computations

1 Introduction

Gradient descent optimization is an indispensable element of solving many real-world problems including but not limited to training deep neural networks [14,19]. Because of its inherent sequentiality it is also particularly difficult to parallelize [17]. Recently a number of advances in developing distributed versions of gradient descent algorithms have been made [15,11,38,39,36]. However, most of them deal with relatively simple variants of the algorithm, for which using larger batch sizes and increasing the learning rate (step size) often yield satisfactory results.

In the case of tasks encountered in Deep Reinforcement Learning these simple optimization procedures are often insufficient and so more advanced algorithms

such as RMSProp [34] and Adam [18] are used more often [22,20,6,2,26]. However, these have not yet been subject to extensive formal analysis or even tests in largely distributed settings. This is crucial since the usual tasks of training models for reinforcement learning are often extremely computationally expensive [22]. Therefore distributed training is gaining more and more traction in the supervised learning community. Devising efficient ways of distributing advanced variants of SGD has the potential to speed up the progress of the entire field.

As our benchmarking task we chose the Atari 2600 emulator [8] provided by the OpenAI Gym framework [10] and the wide variety of games it offers. Atari games are considered a viable benchmark for testing deep reinforcement learning algorithms [6,26,2]. The early attempts to develop agents that would efficiently play Atari games were presented in [21]. This algorithm required as much as 8 days of training on a GPU [20] to reach a level that surpassed a casual human player. Later developments of the Asynchronous Advantage Actor-Critic (A3C) algorithm [20] reduced the learning time to several hours. Because of the work presented in [2] a version of this algorithm optimized for Intel® CPUs was already publicly available. A brief discussion of the single-node version of the algorithm is presented in section 2.1.

In this work we present a distributed version of this algorithm that achieves linear scaling for the tested games for configurations of up to 64 nodes (see figure 8). This allowed us to reduce the training time from roughly 10 hours to around 20 minutes while preserving the original accuracy of the models obtained. For a comparison with other similar implementations we refer the reader to table 3.

Our contribution applies and extends the recent advances [15,11] in distributed supervised learning to the field of reinforcement learning. Sections 2.2, 2.3, 2.4 and 2.5 report on the design choices we made and the results they yielded. We also make our source code available for anyone who would like to reproduce or improve upon our results. Detailed instructions about running the experiments are also provided[5].

1.1 Related work

While distributed machine learning has recently been a topic of extensive research, it has mainly focused on supervised learning. For an in-depth review of scalability of modern supervised learning approaches, we refer the reader to [17]. This work also lists common problems with various approaches to distributing various gradient descent optimization procedures. The pitfalls identified include the communication overhead arising from the necessity to share the weight updates between the nodes. The authors concluded that using larger batches and step sizes had the potential to solve this problem but resulted in less accurate models.

Relation to [11]. Research done in [11] delves more into the architectural aspects of distributed learning, by proposing to abandon the asynchronous design

---

[5] The source code along with game-play videos can be found at: https://github.com/deepsense-ai/Distributed-BA3C.

in favor of a synchronous one. It also makes detailed arguments about the problem of "stale gradients", which prevents the asynchronous paradigm from scaling beyond several nodes. We set out to verify these claims by performing our reinforcement learning experiments using both synchronous and asynchronous training in section 2.3. For a survey of the various asynchronous gradient descent procedures we refer the reader to [9] and [16].

Relation to [15]. Work done in [15] focuses on large-scale supervised learning. It showed that a setup with many machines working concurrently can effectively speed up the training by a large margin. As a result of parallelizing multiple GPUs and using appropriate learning rates for effectively larger batches, training Resnet-50 on Imagenet was completed in 1 hour. The work [15] also showed that using very large batch sizes requires rethinking the optimization algorithms used. In [15] authors focused on the SGD with momentum optimizer which often works very well in supervised learning [33]. Our work attempts to apply similar principles to the Adam optimizer which we found more suitable for reinforcement learning tasks. The details can be found in section 2.4.

The authors of [38] recognized the need to modify the optimization procedures in order to better utilize the distributed settings. The modification proposed a novel procedure called "Layer-wise Adaptive Rate Scaling", which enabled efficient training of supervised learning models with batch size of up to 32768. Deploying this algorithm to the task of training large convolutional nets in [39] yielded extremely competitive training times of 24 minutes, as opposed to 1 hour achieved without these modifications in [15].

An interesting approach to reducing the communication overhead by ternarizing the gradients was recently proposed in [36]. This is a part of larger research aiming at gradient quantization i.e., reducing the precision of the communicated values [4] [27]). A related approach is gradient sparsification, i.e. refraining from the exchange of small gradients, see e.g., [3], [32]). However, both quantization and sparsification drastically change the flow of training of a neural model. Since Reinforcement Learning training is already quite complex, we refrained from employing these methods. Still, they certainly should be considered in future DDRL experiments.

To date, only limited formal research has been done in optimizing and parallelizing targeted strictly at reinforcement learning procedures. Notable works in this domain include [23], which focused on reducing the long training times observed in [22]. A significant speedup (by an order of magnitude [23]) and higher game scores were achieved. This was done using large resources of up to 130 nodes, by applying the Asynchronous SGD paradigm to the model developed in [22] in a manner similar to the work focusing on supervised learning presented in [12].

Further work in [20] applied the asynchronous paradigm to the policy optimization methods, resulting in the Asynchronous Advantage Actor-Critic algorithm (A3C). These experiments used the relatively low computing power of a 16-core CPU. The work in [2] sought to optimize a more efficient batched variant of this algorithm for use with commodity Intel® Xeon CPUs by employing the

Math Kernel Library. A GPU-based version of this algorithm has also been presented in [6]. None of these works explicitly dealt with communication overheads in distributed policy optimization.

Significant computing resources were used in the work presented in [28] to develop AlphaGo – a program for playing the game of Go. This work was based on a combination of reinforcement learning, supervised learning and tree search methods. The authors reported using configurations of up to 1920 CPUs and 280 GPUs for testing the algorithm which provided a significant improvement in the quality of the results achieved. It is also mentioned that the training of the policy network was done using 50 GPUs for one day [28].

Further work in [30] focused on achieving better results without using supervised learning and handcrafted features. Computationally, the training utilized the synchronous paradigm with 64 GPU workers and 19 parameter servers, using a total batch size of 2048. For optimization the Momentum SGD optimizer with learning rate annealing was used. Recently this work has been further extended in [29] where the authors presented a general algorithm able to achieve expert level also in chess and shogi. Notably the training was completed in 24 hours and a relatively large batch size of 4096 was used.

Recently a novel algorithm called Proximal Policy Optimization (PPO) was proposed in [26]. Notably it also uses the Adam optimizer on which we focused in this work and achieved promising scores in Atari games. A distributed version of this algorithm was used in [7], where the authors used 4 GPUs, Adam optimizer and batch size of 5120. Further examination of PPO in a distributed setup appears as a promising area of future research.

A different approach to distributing reinforcement learning has recently been presented in [24]. In this work the parallelization was applied to an evolution strategy (ES), which is a direct search method. This property enables efficient exchange of information between the workers since they only have to communicate scalar values. Because of that the algorithm is especially easy to distribute since the communication costs of sharing the gradient updates don't apply to this scheme. When training a 3D humanoid to walk the authors reported linear scaling for up to 1440 CPU cores [24, p. 8]. After 1 hour of training agents for Atari games on 720 CPUs evolution strategies were able to achieve scores comparable to the ones achieved by A3C (which was trained for 24 hours with a single CPU) [24, p. 7].

Detailed analysis of Deep Reinforcement Learning on a single machine with multiple GPUs was recently published in [31]. The authors reported using a batch size of up to 2048 and utilizing 8 GPUs for various Deep RL algorithms. Impressive training times (under 10 minutes) are also reported.

2 Distributed BA3C implementation

For all our experiments we used the Batch Asynchronous Advantage Actor-Critic algorithm implemented in [37] and later modified in [2]. A similar algorithm was also described in [6]. We will not elaborate on its properties here but rather

focus on explaining the details behind implementing it on multiple parallel machines. A good description of the single-node version of this algorithm can be found in [2]. Here we are using multiple clustered CPUs and each of them individually performs the BA3C algorithm. The individual nodes maintain a shared copy of the model through the use of special nodes called parameter servers, which store the model weights. Table 1 presents the hyperparameters of the algorithms, which were specifically tuned to achieve lowest training times during our research. Some of the hyperparameters are omitted and for those we assume the default values used in [2]. Detailed description of the hyperparameters describing the Adam optimizer is presented in section 2.4.

Table 1: Hyperparameters of the distributed BA3C implementation

Symbol	Default value	Description
n	64	Number of CPU nodes used for the distributed training
c	12	Number of cores in each worker CPU
η	0.001	Learning rate (step-size)
bs	32	Number of data-points in each training batch (on each node)
ps	4	Number of nodes responsible for holding the model parameters
n_sim	10	Number of Atari simulators used simultaneously on one worker
ϵ	10^{-8}	Constant used for numerical stability in the Adam optimizer
β_1, β_2	0.8, 0.75	Decay rates of the running averages used in Adam optimizer

2.1 BA3C background

The BA3C design on a single node focuses on parallel interactions of multiple agents with game environments, that produce experience data later aggregated into mini-batches used for training. We call an "agent" an instance of the model interacting with the outside environment, in our context the Atari 2600 emulator.

The interaction with an Atari game is attained via OpenAI Gym [10] providing the model with input of RGB image pixels and enabling the agents to act upon any state of the game. On a single machine, n_sim simulators of the Atari environment are running concurrently and an agent plays one game on each of them. Each consecutive frame_hist frames from the game count as one state. For each state of the game a policy query is sent to the prediction thread, which then feeds the input image to the neural net, outputting a respective behavioral guideline. The action performed is then gathered together with the state it was acted on and the reward it received from that action. This tuple creates a data-point. Then, bs data-points are assembled into a mini-batch, that is back-propagated through the neural net giving gradients, later used to perform gradient descent on a cost function.

2.2 Architecture

Let us suppose that we are using n workers. Each of the n workers possesses a copy of the BA3C algorithm, which will parallelize the training among its c

cores. As aforementioned, a tuple of state, action and reward constitutes a single data-point, and a mini-batch of `bs` of those data-points is then used to compute the gradients which are synchronously gathered from all the workers, averaged and applied to the weights through the Adam optimizer. All the weights are located in `ps` different parameter servers, each of which stores $1/$`ps` of the model's parameters (4 convolution layers and 3 fully connected layers). The parameter servers then send the updated parameters back to the workers, which then again play the games and the process goes on until a satisfactory model is achieved.
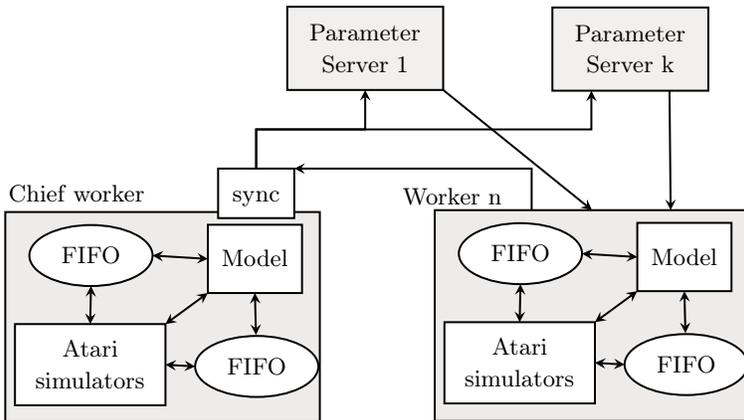

Fig. 1: *Our approach to distributed learning. The figure shows the synchronous training architecture which was our final choice.*

2.3 Synchronous vs Asynchronous training

Background on gradient descent optimization. Training deep neural networks usually involves gradient descent optimization. This is convenient since the gradient of the model with respect to some chosen loss function can be easily obtained by backpropagation [19]. Gradient descent is an iterative algorithm that in each iteration attempts to modify the model's parameters $\boldsymbol{\theta}$ in order to achieve a lower value of the cost function $J(\boldsymbol{\theta}, \boldsymbol{x})$ for some training data $\boldsymbol{x}$. Given the gradient $\boldsymbol{g}_t$ of the cost function w.r.t. the model parameters $\boldsymbol{\theta}$ (which can be obtained from the backpropagation procedure) the basic update rule for obtaining the new values of parameters at time step t given the old values $\boldsymbol{\theta}_{t-1}$ can be written as:

$$\boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} - \lambda \boldsymbol{g}_t, \tag{1}$$

where the λ parameter controls the learning rate (step size).

Numerous improvements to this scheme have been proposed. For a broad overview of different approaches we refer the reader to [19]. Of the recent improvements the RMSProp [34] and Adam [18] procedures are widely used in Reinforcement Learning [22,20,6,2,26]. In the course of our experiments we found that Adam performs better on our task and therefore we will not elaborate on RMSProp further. A brief description of Adam optimizer is given in section 2.4.

An important decision that largely influences the outcome of the model's performance is the way of parallelizing the work of multiple nodes. Data parallelism in gradient descent algorithms can be done in two ways: synchronously or asynchronously. We have found that when using a large number of distributed nodes, these two approaches produce completely different results.

Asynchronous training. In the asynchronous approach, each of the workers, after collecting a mini-batch of data points, computes gradients and then uses them to perform weight updates. The weights of the model reside in parameter servers, which receive gradients from the workers and send the updated copy of the current model to each training instance. Therefore each individual worker updates the commonly shared parameters of the model without delay as soon as it completes computing its gradient. This has several advantages, one is that compared to a single machine implementation our model is guaranteed to perform k times as many updates, if we are using k workers. Another is that because workers do not need to wait for others to finish but rather apply updates continuously, we are utilizing a lock-free paradigm that helps make the most of the processing power at our disposal.

However, the pure asynchronous approach possesses also other characteristics that could impede the learning and prevent convergence. One such disadvantage is called stale gradients [11]. It arises when a worker updates the weights using gradients that are outdated with regard to the current model. This is guaranteed to happen because during the time that the worker was processing the data and computing the gradients, the model has been updated several times by other nodes and now, when the worker applies the gradients, it will do so with respect to the model that is out-of-date. This is shown in the figure 2.
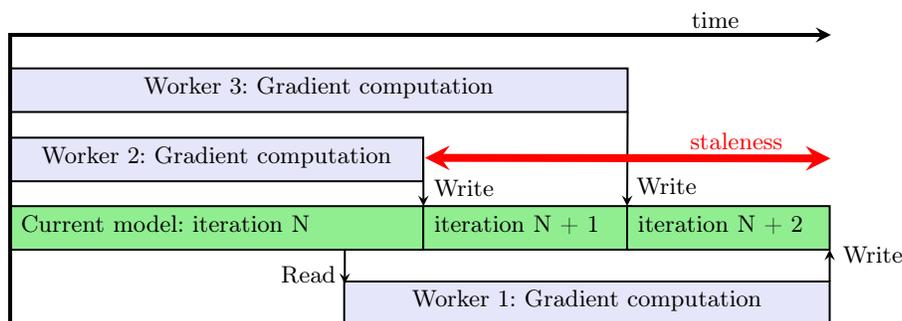

Fig. 2: *A diagram representing gradient staleness - a systematic flaw related to asynchronous training with many workers.*

Synchronous training. The synchronous training architecture is visualized in figure 1. One of the workers (called a *chief worker*) is special in a sense that it's responsible for aggregating the gradients from all the others. A "regular" worker no longer posseses the power to update the model on its own – it can only compute its gradient estimate and send it to the chief worker. Once enough

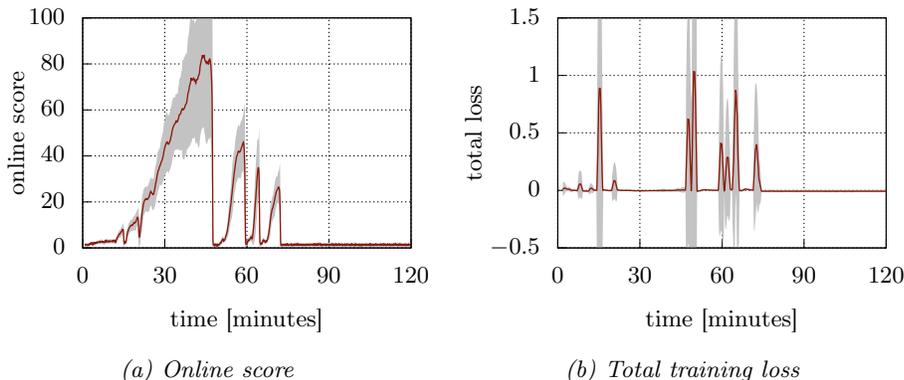(a) Online score

(b) Total training loss

Fig. 3: Typical asynchronous training attempt, 64 workers.

gradients from the workers are accumulated, the chief worker updates the weights and sends the new them to the parameter servers. The new weights are then sent to other workers and a new training iteration begins. This ensures that all the workers always have an up-to-date copy of the model weights, which solves the stale gradients problem.

The side effect of this procedure is the increase of the effective batch size used for performing a single update. Although we are not able to linearly increase the speed of model updates with the number of workers as in the asynchronous design, we expect the updates that are made to be more accurate since the gradient estimate is done using a larger batch size. This might in turn allow for larger step sizes to be used, which can hopefully compensate for the updates being less frequent and provide speed up. Importantly, synchronous training removes the problem of gradient staleness, as no worker computes gradients on an obsolete model, because updates are performed only after all of the workers compute their individual gradients.

The fact that we need to wait for all the workers can cause delays. This arises whenever, for various reasons, some of the workers may be lagging behind others in assembling their batches and computing gradients. We call this the *slow stragglers problem*. Since the synchronous design imposes waiting for all the workers' gradients to perform a weight update, the effective time it takes for an update to occur is the time it takes the slowest worker to assemble his batch and compute the gradients. Therefore, reducing the number of gradients that we need to wait for to make an update could significantly reduce the influence of the slow stragglers. Detailed analysis of this phenomenon presented in [11] confirms that waiting for around 90% of gradients as opposed to all of them significantly improves the training times.

Another key fact that needs to be addressed when discussing synchronous training is the large effective batch size[6] it tends to create. Since we are using

---

[6] We use the term *effective batch size* to denote the number of training samples participating in a single weight update. In synchronous training this is equal to the local

n mini-batches of data-points from every worker and then averaging them, we are virtually using a single batch of size $n \times \mathtt{bs}$ to perform a single update. This may indicate the need to adjust other hyperparameters of the algorithm such as the learning rate. We revisit this issue in section 2.4.

With the slow stragglers problem removed, the synchronous approach is much more intuitive and reasonable – it does not risk gradient staleness and the weight updates that are made are much more accurate and less noisy. Making just one update for all the workers working on a model assures that the nodes are working collectively and efficiently upon a goal, whereas the asynchronous learning strategy seems rather chaotic and unstructured.

We performed a series of experiments to determine which paradigm is better in our use case. We found that asynchronous training causes large instabilities in the learning process. One of such experiments is shown in the figures 3a and 3b. In this experiment the learning was proceeding correctly until after about 50 minutes the online score[7] dropped suddenly to zero. This coincided with a large spike in the total training loss visible in the figure 3b. We suspect it is caused by the stale gradients. We did not encounter this phenomenon when using synchronous training, therefore we have chosen to work with the synchronous architecture.

2.4 Optimizer changes

Background on Adam optimizer. Adam optimizer was first described in [18] and can be thought of as extension of the works presented in [34] and [13]. It maintains the exponentially decaying running averages m_t and v_t of all the previous gradients and squared gradients:

$$m_t = \beta_1 m_{t-1} + (1 - \beta_1)g_t, \tag{2}$$
$$v_t = \beta_2 v_{t-1} + (1 - \beta_1)g_t^2 \tag{3}$$

It then perform bias correction to define $\hat{m}_t = m_t/(1 - \beta_1^t)$ and $\hat{v}_t = v_t/(1 - \beta_2^t)$ and gives the final weight update for parameter θ at timestep t as:

$$\theta_t = \theta_{t-1} - \eta \frac{\hat{m}_t}{\sqrt{\hat{v}_t} + \epsilon} \tag{4}$$

---

batch size on each node multiplied by the number of workers required to perform an update. In asynchronous training effective batch size is equal to the local batch size.

[7] By *online score* we refer to the scores obtained by the agent during training. By contrast an *evaluation score* would be a score obtained during the test phase. These scores can differ substantially, because while training the actions are sampled from the distribution returned by the policy network (this ensures more exploration). On the other hand, during test time the agent always chooses the action that gives the highest expected reward. This usually yields higher scores, but using it while training would prevent exploration.

Instead, many implementations (including TensorFlow [1]) use less clear but more efficient formulation:

$$\eta_t = \eta \frac{\sqrt{1 - \beta_2^t}}{1 - \beta_1^t}, \tag{5}$$

$$\theta_t = \theta_{t-1} - \eta_t \frac{m_t}{\sqrt{v_t} + \hat{\epsilon}} \tag{6}$$

The $\hat{\epsilon}$ in the equation 6 and ϵ in equation 4 are added for numerical stability, not to divide by 0 in the first timestep.

This means that the algorithm in this formulation has 4 hyperparameters that need tuning: the learning rate η, the decay factors for the running averages: β_1 and β_2 and $\hat{\epsilon}$. Next section provides insight into how these might need to be modified when transitioning from a single-node to a multi-node configuration.

Increasing the learning rate. Using very large batches that result from utilizing a large number of workers in the synchronous paradigm poses some challenges on the selection of optimizer hyperparameters. This problem is especially severe when distributing an algorithm that already had its hyperparameters chosen carefully.

Research on large scale distributed SGD by [15] has addressed this problem by deploying the linear scaling rule: when multiplying the mini-batch size by k, multiply the learning rate by k. However this was done using much simpler SGD with momentum optimizer. We on the other hand have experimented with multiple optimizers and have found that only Adam [18] and occasionally RMSProp [34] have brought about positive results in the asynchronous design.

With Adam optimizer, using the linear scaling rule did not yield any positive results. We found that increasing the learning rate made the training highly unstable and often resulted in the model learning how to play well only to later abruptly forget and score 0 until the end (see figure 4). We settled on using $\eta = 0.001$ (the same as in the single-node version), as it was the largest value for which we did not experience large instabilities.

Apart from the learning rate some of the other default optimizer parameters also needed examination. Moving to a synchronous distributed setup requires a re-thinking of how exactly momentum accumulation and learning rate adaptivity are impacted by the batch size.

Modifying the ϵ parameter. Curiously enough, some implementations can be found that manipulate this variable so that it no longer serves the mere purpose of avoiding numerical instability (see e.g. the implementation of BA3C in [37]). Through experiments we found that for some tasks setting the ϵ parameter of the optimizer to much lower values (e.g. 10^{-8} instead of 10^{-3}) can yield much better training times. A comparison of online scores for two similar experiments with different epsilon values is shown in the figure 5. It is important to note that this positive effect when using smaller ϵ was observable only when using large
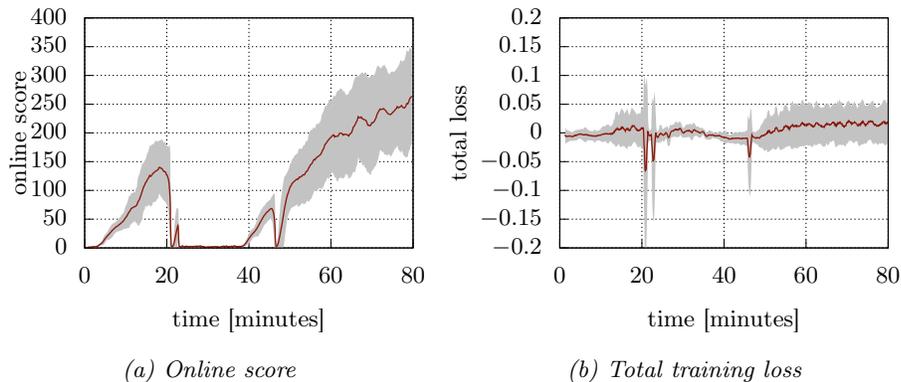
(a) Online score　　　　　*(b) Total training loss*

Fig. 4: Experiment with a learning rate $\eta = 0.002$. 64 nodes, synchronous training, local batch size of 64, total batch size of 4096. Increasing the learning rate for the Adam optimizer from 0.001 to 0.002 causes large instabilities clearly visible on the online score plot (figure 4a) and the total training loss plot (figure 4b)

effective batch sizes (i.e., 512 and more). For smaller effective batch sizes using a lower ϵ did not produce positive results.

This is understandable since a high ϵ value significantly constrains the ability of the Adam optimizer to automatically adapt the learning rate to the variance of the gradients. We suspected that averaging more data points through the use of synchronous data parallelism reduced the variance of the gradient estimate to the point that the algorithm could be allowed more freedom in automatically adapting the learning rate based on the noise estimations.

Based on these experiments we decided to change the default value of the ϵ hyperparameter from 10^{-3} to 10^{-8}. This significantly

Table 2: Number of network parameters when considering different number of neurons in the fully connected layer.

hidden neurons	network weights	% of initial setup
256	538 119	100%
128	332 295	61%
64	229 383	43%
32	177 927	33%
16	152 199	28%

sped up the training for some games (such as Breakout and Boxing). However, we cannot claim this effect is universal, e.g., it made training for Atari Pong slower. The results cited for this game in section 3 were obtained with $\epsilon = 10^{-3}$.

Other hyperparameters. Motivated by the research in [15] we decided that we should optimize Adam's decay factors β_1, β_2 to the very large batch that we are using. This did not turn out to be an easy task - with the Adam optimizer update policy being quite complicated, choosing β_1 and β_2 for the effective batchsize analytically was difficult. The results of our experiments do not support any gains from using different values of these parameters; however we are leaving it to the community to try and find the factors that work best for a distributed setup.
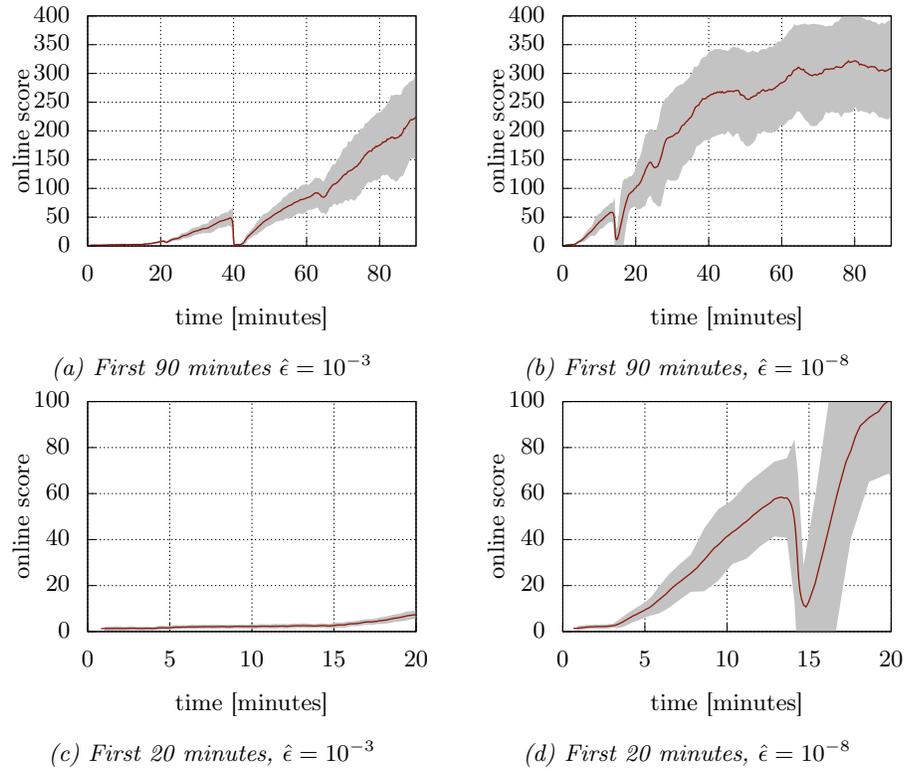
(a) First 90 minutes $\hat{\epsilon} = 10^{-3}$

(b) First 90 minutes, $\hat{\epsilon} = 10^{-8}$

(c) First 20 minutes, $\hat{\epsilon} = 10^{-3}$

(d) First 20 minutes, $\hat{\epsilon} = 10^{-8}$

Fig. 5: Two experiments of training agents for Breakout on 64 nodes with different ϵ parameter values. Figures 5a and 5c show an experiment where $\epsilon = 10^{-3}$ was used. Figures 5b and 5d show training with $\epsilon = 10^{-8}$. The most important difference lies at the beginning of the training. This is visible in the closer views presented in the bottom row. Lower values of ϵ seem to give a significant speedup at this stage. Note that the vertical axis shows online score.

2.5 Communication overhead

In data parallelized synchronous gradient descent procedures the nodes have to transmit roughly $2n$ multiplicities of the size of the model (where n is the number of workers) during a single training step [17]. Our initial model architecture consisted of approximately half a million weights stored as 32-bit floats. Thus, when using 64 workers we needed to send ≈ 263MB of data during every iteration.

Importantly, if all nodes have roughly the same processing speed and synchronous training is being used then all this communication occurs at
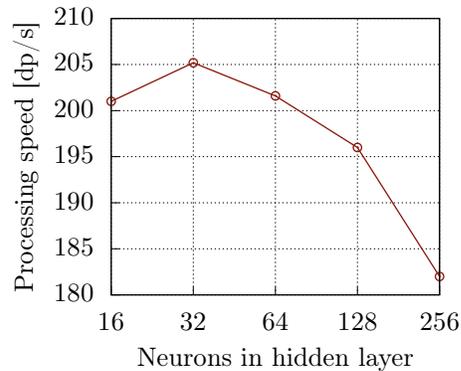

Fig. 6: Data points per second for models with different number of hidden neurons in the fully connected layer. Each experiment was repeated 5 times and results were averaged. Experiments were run with 32 workers and 4 parameter servers.

approximately the same time. This is because in a gradient descent training the communication cannot be easily overlaid with computation to maximally utilize both network bandwidth and compute power (see [17] for details). This further reduces the scaling capabilities.

In our experiments we measured the speed of our algorithm by calculating the number of training examples backpropagated through our model every second. In next section we will refer to this speed as data points per second.

Changing the model. One way to reduce network communication is to shrink the model. In the initial architecture most weights ($\approx 76\%$) were in a single fully connected layer that follows the last convolution layer (see [2] for the details about the exact neural model used). The relation between number of neurons in this layer and the processing speed is shown in figure 6. Although all of the tested architectures were able to achieve decent re-
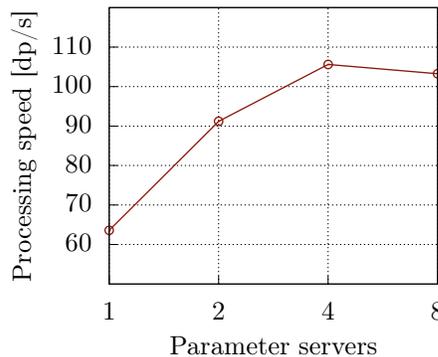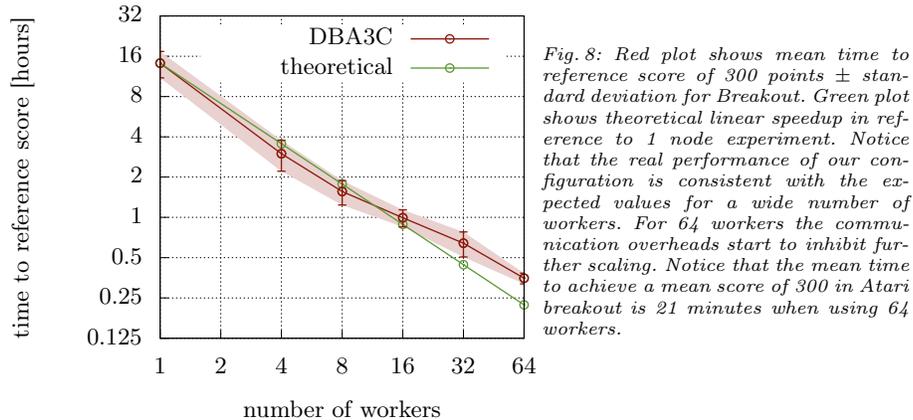

Fig. 7: Data points per second for different numbers of parameter servers. Each experiment was repeated 5 times and the results were averaged. 32 workers, 332k parameters, local; batch size set to 4. The conclusion is that after some point there's no more gains to be achieved by adding more parameter servers.

sults, we decided to use 128 neurons since this setup was able to learn as fast and stable as the initial architecture, while having only $\approx 61\%$ of its weights. Despite the fact that further reduction of the model size accelerated data processing, smaller networks were taking more time to reach corresponding scores.

Adding more parameter servers. Adding more parameter servers, each storing only fraction of model weights, causes data sent through the network to be distributed into more nodes. This leads to more optimized network usage (see figure 7 for details).

3   Results

Synchronous training allowed us to use more workers and avoid instabilities common in the asynchronous paradigm. By reducing model size and adding more parameter servers we could better utilize network communication which led to the possibility of further increase in the number of workers. As a result we were able to train models to reach 300 points in Breakout in 21 ± 2 minutes using 64 workers (each consisting of 12 physical cores, i.e. using 768 cores in total).

3.1   Scaling

We compared times it took to reach a predetermined score in Atari Breakout for different number of workers. The reference was reaching a mean score of 300

points or higher for 50 consecutive games played. This is considered vastly better than a human tester (see [22] for data on detailed human performance for this game). The results are shown in figure 8.

Fig. 8: Red plot shows mean time to reference score of 300 points ± standard deviation for Breakout. Green plot shows theoretical linear speedup in reference to 1 node experiment. Notice that the real performance of our configuration is consistent with the expected values for a wide number of workers. For 64 workers the communication overheads start to inhibit further scaling. Notice that the mean time to achieve a mean score of 300 in Atari breakout is 21 minutes when using 64 workers.

Experiment settings: Each of the 64 workers had 12 CPU cores. We used 4 parameter servers for storing model weights. Model trained had 128 hidden neurons in the fully connected layer described in section 2.3. Every experiment was repeated 10 times and results were averaged. Additionally we have plotted a theoretical linear speedup. This line represents the theoretical time that should be achieved when using n times more computing power in reference to a single node experiment.

Learning rate: All experiments used the learning rate of 10^{-3}.

Optimizer's hyperparameters: In all experiments optimizer's hyperparameters were: $\epsilon = 10^{-8}$, $\beta_1 = 0.8$, $\beta_2 = 0.75$.

Batch size: In experiments with 32 and 64 workers batch size was set to 32, because smaller batches caused too much network communication overhead. For the rest of experiments the per-worker batch size was set, so that the effective batch size equaled $n \times \mathtt{bs} = 512$.

Evaluation: Every 1000 steps the model played 50 games and the mean score was saved. During the games the model parameters were frozen. By step we mean single global update performed by the chief worker.

Baseline: As a baseline we have chosen the single node setup (i.e. using a single 12-core CPU). To be comparable with effective batch sizes on multiple nodes, a relatively large batch size of 512 was chosen. This baseline achieves the solving score in mean time of 14.2 hours.

3.2 Training times

In this section we present example learning curves for various Atari games. The plots show mean and max scores from evaluation games. Each game was played on the 64 worker setup.
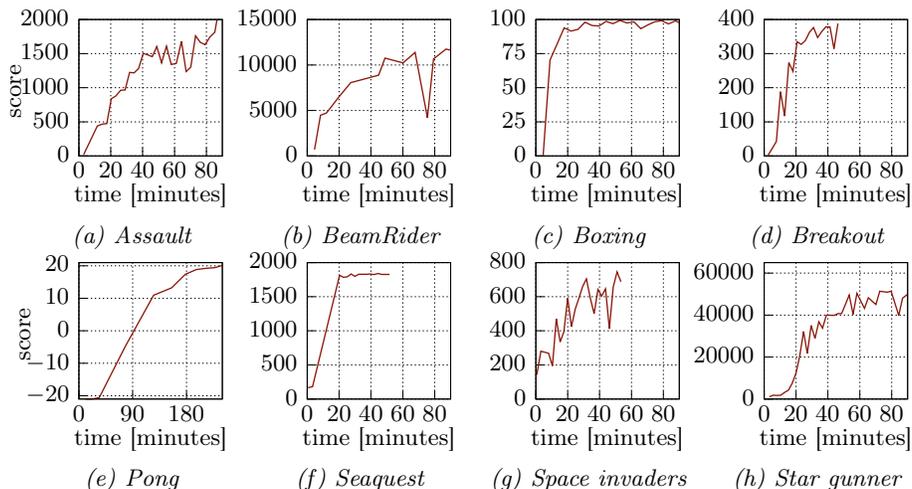
Fig. 9: Score vs time plots for different games in the final setup.

3.3 Comparison with other solutions

The most notable similar work in optimizing Atari games training for speed is presented in [6]. The results presented there were achieved by a hybrid GPU-CPU algorithm called GA3C which is a flavor of A3C focusing on batching the data points in order to better utilize the massively parallel nature of GPU computations. This is similar to the single node algorithm called BA3C [2] which we used as a starting point for this work.

Comparing the training curves included in [6] for 3 common games tested in both works (Boxing, Breakout, SpaceInvaders) shows that our implementation is much faster and achieves as good or better scores[8] (see table 3). Importantly our experiments used 64 CPU nodes of 12 cores each, while the experiments presented in [5] were all single node. However the results show that using distributed computations on CPU clusters is a viable alternative for using GPUs, even when training convolutional neural networks.

4 Conclusions and future work

We presented a detailed description of our experiments with large scale Distributed Deep Reinforcement Learning (DDRL). Detailed motivation behind all the important design choices was given in sections 2.2, 2.3, 2.4 and 2.5. We also provided some empirical information about tuning the Adam optimizer to

---

[8] It is important to note that the scores achieved by different implementations are not directly comparable and should interpreted cautiously. For future comparisons we'd like to state that the evaluation scores presented by us in this work are always mean scores of 50 consecutive games played by the agent. Unless otherwise stated they're evaluation scores achieved by choosing the action giving the highest future expected reward.

Table 3: Algorithm performance in 3 games tested in both papers. Best stable score and time (in hours) to achieve it are given. The data are based on the best reported results found in the training plots in [6,2,20].

Game	DDRL A3C	GA3C [6]	BA3C [2]	A3C [20]
BeamRider	14900 (2.7h)	3000 (24h)	–	15000 (15h)
Breakout	350 (0.5h)	350 (21h)	400 (15h)	500 (11h)
Boxing	98 (0.5h)	92 (2h)	–	–
Pong	20 (4h)	18 (1h)	17 (24h)	20 (8h)
Seaquest	1832 (0.5h)	1706 (24h)	1840 (24h)	2300 (24h)
SpaceInvaders	650 (0.5h)	600 (24h)	700 (24h)	1400 (15h)

perform well when using large training batches that arise in synchronous data parallelism. Our key experimental results described in section 3 involve being able to train agents for playing Atari games in minutes rather than hours on clusters of commodity CPUs.

Extending this work to other RL algorithms, most notably those presented in [25,26,35] would provide a natural extension to this work. Also developing a framework for distributed RL training that is independent of the algorithm itself would certainly be a valuable contribution.

Given the results reported in [39] testing the Intel® Xeon Phi™ architecture on distributed RL training would also be an interesting experiment.

On a wider scale, further research on adaptive optimization algorithms, most notably those presented in [18,13,34] in the context of training with large batch sizes seems to be necessary to further reduce training times both in supervised and reinforcement learning.

5   Acknowledgments

References

1. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng,

X.: TensorFlow: Large-scale machine learning on heterogeneous systems (2015), https://www.tensorflow.org/, software available from tensorflow.org

2. Adamski, R., Grel, T., Klimek, M., Michalewski, H.: Atari games and intel processors. CoRR abs/1705.06936 (2017)
3. Aji, A.F., Heafield, K.: Sparse communication for distributed gradient descent. CoRR abs/1704.05021 (2017)
4. Alistarh, D., Li, J., Tomioka, R., Vojnovic, M.: QSGD: Randomized quantization for communication-optimal stochastic gradient descent. CoRR abs/1610.02132 (2016)
5. Babaeizadeh, M., Frosio, I., Tyree, S., Clemons, J., Kautz, J.: GA3C: GPU-based A3C for deep reinforcement learning. CoRR abs/1611.06256 (2016)
6. Babaeizadeh, M., Frosio, I., Tyree, S., Clemons, J., Kautz, J.: Reinforcement learning through asynchronous advantage actor-critic on a GPU. In: ICLR (2017)
7. Bansal, T., Pachocki, J., Sidor, S., Sutskever, I., Mordatch, I.: Emergent complexity via multi-agent competition. CoRR abs/1710.03748 (2017)
8. Bellemare, M.G., Naddaf, Y., Veness, J., Bowling, M.: The arcade learning environment: An evaluation platform for general agents. CoRR abs/1207.4708 (2012)
9. Bhardwaj, O., Cong, G.: Practical efficiency of asynchronous stochastic gradient descent. In: 2016 2nd Workshop on Machine Learning in HPC Environments (MLHPC). vol. 00, pp. 56–62 (Nov 2016)
10. Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., Zaremba, W.: OpenAI Gym. CoRR abs/1606.01540 (2016)
11. Chen, J., Monga, R., Bengio, S., Jozefowicz, R.: Revisiting distributed synchronous SGD. In: International Conference on Learning Representations Workshop Track (2016)
12. Dean, J., Corrado, G.S., Monga, R., Chen, K., Devin, M., Le, Q.V., Mao, M.Z., Ranzato, M., Senior, A., Tucker, P., Yang, K., Ng, A.Y.: Large scale distributed deep networks. In: Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1. pp. 1223–1231. NIPS'12, Curran Associates Inc., USA (2012)
13. Duchi, J., Hazan, E., Singer, Y.: Adaptive subgradient methods for online learning and stochastic optimization. J. Mach. Learn. Res. 12, 2121–2159 (Jul 2011)
14. Goodfellow, I., Bengio, Y., Courville, A.: Deep Learning. MIT Press (2016), http://www.deeplearningbook.org
15. Goyal, P., Dollár, P., Girshick, R.B., Noordhuis, P., Wesolowski, L., Kyrola, A., Tulloch, A., Jia, Y., He, K.: Accurate, large minibatch SGD: training ImageNet in 1 hour. CoRR abs/1706.02677 (2017)
16. Keuper, J., Pfreundt, F.: Asynchronous parallel stochastic gradient descent - A numeric core for scalable distributed machine learning algorithms. CoRR abs/1505.04956 (2015)
17. Keuper, J., Preundt, F.J.: Distributed training of deep neural networks: Theoretical and practical limits of parallel scalability. In: Proceedings of the Workshop on Machine Learning in High Performance Computing Environments. pp. 19–26. MLHPC '16, IEEE Press, Piscataway, NJ, USA (2016)
18. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. CoRR abs/1412.6980 (2014)
19. Le, Q.V., Ngiam, J., Coates, A., Lahiri, A., Prochnow, B., Ng, A.Y.: On optimization methods for deep learning. In: Proceedings of the 28th International Conference on International Conference on Machine Learning. pp. 265–272. ICML'11, Omnipress, USA (2011)

20. Mnih, V., Badia, A.P., Mirza, M., Graves, A., Lillicrap, T.P., Harley, T., Silver, D., Kavukcuoglu, K.: Asynchronous methods for deep reinforcement learning. CoRR abs/1602.01783 (2016)
21. Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., Riedmiller, M.: Playing atari with deep reinforcement learning. In: NIPS Deep Learning Workshop (2013)
22. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., Hassabis, D.: Human-level control through deep reinforcement learning. Nature 518(7540), 529–533 (Feb 2015)
23. Nair, A., Srinivasan, P., Blackwell, S., Alcicek, C., Fearon, R., Maria, A.D., Panneershelvam, V., Suleyman, M., Beattie, C., Petersen, S., Legg, S., Mnih, V., Kavukcuoglu, K., Silver, D.: Massively parallel methods for deep reinforcement learning. CoRR abs/1507.04296 (2015)
24. Salimans, T., Ho, J., Chen, X., Sutskever, I.: Evolution strategies as a scalable alternative to reinforcement learning. CoRR abs/1703.03864 (2017)
25. Schulman, J., Levine, S., Moritz, P., Jordan, M.I., Abbeel, P.: Trust region policy optimization. CoRR abs/1502.05477 (2015)
26. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal policy optimization algorithms. CoRR abs/1707.06347 (2017)
27. Seide, F., Fu, H., Droppo, J., Li, G., Yu, D.: 1-bit stochastic gradient descent and application to data-parallel distributed training of speech DNNs. In: Interspeech 2014 (September 2014)
28. Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., Hassabis, D.: Mastering the game of Go with deep neural networks and tree search. Nature 529(7587), 484–489 (Jan 2016)
29. Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T., Lillicrap, T., Simonyan, K., Hassabis, D.: Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm (Dec 2017)
30. Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., Chen, Y., Lillicrap, T., Hui, F., Sifre, L., van den Driessche, G., Graepel, T., Hassabis, D.: Mastering the game of Go without human knowledge 550, 354–359 (10 2017)
31. Stooke, A., Abbeel, P.: Accelerated methods for deep reinforcement learning. CoRR abs/1803.02811 (March 2018)
32. Strom, N.: Scalable distributed DNN training using commodity GPU cloud computing. In: INTERSPEECH. pp. 1488–1492. ISCA (2015)
33. Sutskever, I., Martens, J., Dahl, G., Hinton, G.: On the importance of initialization and momentum in deep learning. In: Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28. pp. III–1139–III–1147. ICML'13, JMLR.org (2013)
34. Tieleman, T., Hinton, G.: Lecture 6.5—RmsProp: Divide the gradient by a running average of its recent magnitude. COURSERA: Neural Networks for Machine Learning (2012)
35. Wang, Z., Bapst, V., Heess, N., Mnih, V., Munos, R., Kavukcuoglu, K., de Freitas, N.: Sample efficient actor-critic with experience replay. CoRR abs/1611.01224 (2016)

36. Wen, W., Xu, C., Yan, F., Wu, C., Wang, Y., Chen, Y., Li, H.: TernGrad: Ternary gradients to reduce communication in distributed deep learning. CoRR abs/1705.07878 (2017)
37. Wu, Y.: Tensorpack. https://github.com/ppwwyyxx/tensorpack (2016)
38. You, Y., Gitman, I., Ginsburg, B.: Scaling SGD batch size to 32K for ImageNet training. CoRR abs/1708.03888 (2017)
39. You, Y., Zhang, Z., Hsieh, C.J., Demmel, J.: 100-epoch ImageNet training with AlexNet in 24 minutes. arXiv preprint arXiv:1709.05011 (2017)