

# Distributed stochastic optimization for deep learning

by

Sixin Zhang

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Computer Science

New York University

May, 2016

---

Yann LeCun

## Dedication

*To my parents and grandparents*

## Acknowledgements

Almost all the thesis start with a thank you to the thesis advisor, why this is so. Maybe it is the time. The time one spends most with during the PhD. The time the impact from whom will last for long on shaping one's flavor of the research. My advisor, Yann LeCun, is of no double having a big influence on me. To some extent, the way we interact is like the *Elastic Averaging SGD (EASGD)* method that we have named together. During my whole PhD, I am very grateful to have the freedom to explore various research subjects. He always has patience to wait, even though sometimes there is no interesting results. It is an experiment. Sometimes, the good result is only one step away, and he has a good feeling to sense that. Sometimes, he is also very strict. Remember when I was preparing the thesis proposal, but was proposing too many directions. Everyone was worried, as the deadline was approaching. To focus and be quick. A lot of the time, the way how we think decides the question we ask and where we go. My advisor's conceptual way of thinking is still to me a mystery of arts and science.

I would also like to thank Anna Choromanska for being my collaborator and for helping me at many critical points. When we started to work on the *EASGD* method, it was not very clear why we gave it a new name. The answer only became sharp when Anna questioned me countlessly why we do what we do. Luckily, we have discovered interesting answers. Chapter 2, Chapter 3 and Chapter 4 of this thesis are based on our joint work.

Chapter 5 is inspired by the discussion with Professor Weinan E, and his students Qianxiao Li and Cheng Tai during my visit to Princeton University. We were trying to use stochastic differential equation to analyze *EASGD* method. The idea is to get the maximum insight with minimum effort. The results in Chapter 5 share a similar flavor.

Chapter 6 would not come into shape without the helpful discussions with Professor Jinyang Li, and her students Russell Power, Minjie Wang, Zhaoguo Wang, Christopher Mitchell and Yang Zhang. The design and the implementation of the *EASGD* and *EASGD Tree* are guided by their valuable experience. Also the numerical results would

not be obtained without their feedback, as well as the consistent support of the NYU HPC team, in particular Shenglong Wang.

I would also like to give special thanks to Professor Rob Fergus, Margaret Wright, and David Sontag for serving in my thesis committee, as well as being my teacher during my PhD studies. Margaret was so eager to read my thesis and asked when do you expect your thesis to converge whenever I have a new version. Rob was so glad to see me when I first came to the CBLI lab, while David is always sharing his knowledge with us during the CBLI talks.

I also learn a lot from my earlier collaborators, in particular Marco Scoffier, Tom Schaul and Wan Li. Arthur Szlam, Camille Couprie, Clement Farabet, David Eigen, Dilip Krishnan, Joan Bruna, Koray Kavukcuoglu, Matt Zeiler, Olivier Henaff, Pablo Sprechmann, Pierre Sermanet, Ross Goroshin, Xiang Zhang and Y-Lan Boureau are always very mind-refreshing to be around.

I also appreciate the feedback from many senior researchers on the *EASGD* method. In particular, Mark Tygert, Samy Bengio, Patrick Combettes, Peter Richtarik, Yoshua Bengio and Yuri Bakhtin.

The coach of the NYU running club, Mike Galvan, kicks me up to practice at 6:30am every Monday morning during my last year of the PhD. It gives me a consistency energy to keep running and to finish my PhD.

Almost half of my PhD time is spent at the building of Courant Institute. I have learned a lot of mathematics which is still so tremendous to me. The more you write, the less you see. There are a lot of inspiring stories that I would like to write, but maybe I should keep them in mind for the moment.

## Abstract

We study the problem of how to distribute the training of large-scale deep learning models in the parallel computing environment. We propose a new distributed stochastic optimization method called Elastic Averaging SGD (*EASGD*). We analyze the convergence rate of the *EASGD* method in the synchronous scenario and compare its stability condition with the existing *ADMM* method in the round-robin scheme. An asynchronous and momentum variant of the *EASGD* method is applied to train deep convolutional neural networks for image classification on the *CIFAR* and *ImageNet* datasets. Our approach accelerates the training and furthermore achieves better test accuracy. It also requires a much smaller amount of communication than other common baseline approaches such as the *DOWNPOUR* method.

We then investigate the limit in speedup of the initial and the asymptotic phase of the mini-batch *SGD*, the momentum *SGD*, and the *EASGD* methods. We find that the spread of the input data distribution has a big impact on their initial convergence rate and stability region. We also find a surprising connection between the momentum *SGD* and the *EASGD* method with a negative moving average rate. A non-convex case is also studied to understand when *EASGD* can get trapped by a saddle point.

Finally, we scale up the *EASGD* method by using a tree structured network topology. We show empirically its advantage and challenge. We also establish a connection between the *EASGD* and the *DOWNPOUR* method with the classical *Jacobi* and the *Gauss-Seidel* method, thus unifying a class of distributed stochastic optimization methods.

# Contents

<b>Dedication</b>	<b>ii</b>
<b>Acknowledgements</b>	<b>iv</b>
<b>Abstract</b>	<b>v</b>
<b>List of Figures</b>	<b>viii</b>
<b>List of Tables</b>	<b>xvi</b>
<b>1 Introduction</b>	<b>1</b>
1.1 What is the problem . . . . .	1
1.2 Formalizing the problem . . . . .	4
1.3 An overview . . . . .	5
<b>2 Elastic Averaging SGD (EASGD)</b>	<b>9</b>
2.1 Synchronous <i>EASGD</i> . . . . .	10
2.2 Asynchronous <i>EASGD</i> . . . . .	12
2.3 Momentum <i>EASGD</i> . . . . .	14
<b>3 Convergence Analysis of EASGD</b>	<b>17</b>
3.1 Quadratic case . . . . .	17
3.1.1 One-dimensional case . . . . .	18
3.1.2 Generalization to multidimensional case . . . . .	26
3.2 Strongly convex case . . . . .	29
3.3 Stability of <i>EASGD</i> and <i>ADMM</i> . . . . .	37

<b>4</b>	<b>Performance in Deep Learning</b>	<b>43</b>
4.1	Experimental setup . . . . .	44
4.2	Experimental results . . . . .	47
4.3	Further discussion and understanding . . . . .	59
4.3.1	Comparison of <i>SGD</i> , <i>ASGD</i> , <i>MVASGD</i> and <i>MSGD</i> . . . . .	59
4.3.2	Dependence of the learning rate . . . . .	62
4.3.3	Dependence of the communication period . . . . .	62
4.3.4	The tradeoff between data and parameter communication . . . . .	64
4.3.5	Time speed-up . . . . .	66
4.4	Additional pseudo-codes of the algorithms . . . . .	67
<b>5</b>	<b>The Limit in Speedup</b>	<b>71</b>
5.1	Additive noise . . . . .	71
5.1.1	<i>SGD</i> with mini-batch . . . . .	72
5.1.2	Momentum <i>SGD</i> . . . . .	73
5.1.3	<i>EASGD</i> and <i>EAMSGD</i> . . . . .	78
5.2	Multiplicative noise . . . . .	89
5.2.1	<i>SGD</i> with mini-batch . . . . .	89
5.2.2	Momentum <i>SGD</i> . . . . .	92
5.2.3	<i>EASGD</i> and <i>EAMSGD</i> . . . . .	95
5.3	A non-convex case . . . . .	109
<b>6</b>	<b>Scaling up Elastic Averaging SGD</b>	<b>112</b>
6.1	<i>EASGD Tree</i> . . . . .	112
6.1.1	The algorithm . . . . .	113
6.1.2	The result . . . . .	117
6.2	Unifying <i>EASGD</i> and <i>DOWNPOUR</i> . . . . .	130
<b>7</b>	<b>Conclusion</b>	<b>134</b>
	<b>Bibliography</b>	<b>138</b>

# List of Figures

2.1	The big picture of <i>EASGD</i> . . . . .	10
3.1	Theoretical mean squared error (MSE) of the center $\tilde{x}$ in the quadratic case, with various choices of the learning rate $\eta$ (horizontal within each block), and the moving rate $\beta = p\alpha$ (vertical within each block), the number of processors $p = \{1, 10, 100, 1000, 10000\}$ (vertical across blocks), and the time steps $t = \{1, 2, 10, 100, \infty\}$ (horizontal across blocks). The MSE is plotted in log scale, ranging from $10^{-3}$ to $10^3$ (from deep blue to red). The dark red (i.e. on the upper-right corners) indicates divergence. . . . .	22
3.2	The largest absolute eigenvalue of the linear map $\mathcal{F} = F_3^p \circ F_2^p \circ F_1^p \circ \dots \circ F_3^1 \circ F_2^1 \circ F_1^1$ as a function of $\eta \in (0, 10^{-2})$ and $\rho \in (0, 10)$ when $p = 3$ and $p = 8$ . To simulate the chaotic behavior of the <i>ADMM</i> algorithm, one may pick $\eta = 0.001$ and $\rho = 2.5$ and initialize the state $s_0$ either randomly or with $\lambda_0^i = 0, x_0^i = \tilde{x}_0 = 1000, \forall i = 1, \dots, p$ . <i>Figure should be read in color.</i> . . . . .	41
3.3	Instability of <i>ADMM</i> in the round-robin scheme. Pick $p = 3, \eta = 0.001, \rho = 2.5$ and initialize the state $s_0$ with $\lambda_0^i = 0, x_0^i = \tilde{x}_0 = 1000, \forall i = 1, \dots, p$ . The x-axis is the time step $t$ , the y-axis is the (one-dimensional) value of the center variable $\tilde{x}_t$ . . . . .	42
4.1	Training and test loss and the test error for the center variable versus a wallclock time for communication period $\tau = 1$ on <i>CIFAR</i> dataset with the 7-layer convolutional neural network. $p = 4$ . . . . .	49



- 4.2 Training and test loss and the test error for the center variable versus a wallclock time for communication period  $\tau = 4$  on *CIFAR* dataset with the 7-layer convolutional neural network.  $p = 4$ . . . . . 50
- 4.3 Training and test loss and the test error for the center variable versus a wallclock time for communication period  $\tau = 16$  on *CIFAR* dataset with the 7-layer convolutional neural network.  $p = 4$ . . . . . 51
- 4.4 Training and test loss and the test error for the center variable versus a wallclock time for communication period  $\tau = 64$  on *CIFAR* dataset with the 7-layer convolutional neural network.  $p = 4$ . . . . . 52
- 4.5 Training and test loss and the test error for the center variable versus a wallclock time with the number of local workers  $p = 4$  for parallel methods on *CIFAR* with the 7-layer convolutional neural network. . . . . 53
- 4.6 Training and test loss and the test error for the center variable versus a wallclock time with the number of local workers  $p = 8$  for parallel methods on *CIFAR* with the 7-layer convolutional neural network. . . . . 54
- 4.7 Training and test loss and the test error for the center variable versus a wallclock time with the number of local workers  $p = 16$  for parallel methods on *CIFAR* with the 7-layer convolutional neural network. . . . . 55
- 4.8 Training and test loss and the test error for the center variable versus a wallclock time with the number of local workers  $p = 4$  on *ImageNet* with the 11-layer convolutional neural network. . . . . 56
- 4.9 Training and test loss and the test error for the center variable versus a wallclock time with the number of local workers  $p = 8$  on *ImageNet* with the 11-layer convolutional neural network. . . . . 57
- 4.10 Convergence of the training and test loss (negative log-likelihood) and the test error (original and zoomed) computed for the center variable as a function of wallclock time for *SGD*, *ASGD*, *MVASGD* and *MSGD* ( $p = 1$ ) on the *CIFAR* experiment. . . . . 60

4.11 Convergence of the training and test loss (negative log-likelihood) and the test error (original and zoomed) computed for the center variable as a function of wallclock time for *SGD*, *ASGD*, *MVASGD* and *MSGD* ( $p = 1$ ) on the *ImageNet* experiment. . . . . 61

4.12 Convergence of the training loss (negative log-likelihood, original) and the test error (zoomed) computed for the center variable as a function of wallclock time for *EAMSGD* and *EASGD* run with different values of  $\eta$  on the *CIFAR* experiment.  $p = 16, \tau = 10$ . . . . . 63

4.13 Convergence of the training loss (negative log-likelihood, original) and the test error (zoomed) computed for the center variable as a function of wallclock time for *EASGD* and *EAMSGD* ( $p = 16, \eta = 0.01, \beta = 0.9, \delta = 0.99$ ) on the *CIFAR* experiment with various communication period  $\tau$  and learning rate decay  $\gamma$ . The learning rate is decreased gradually over time based each local worker’s own clock  $t$  with  $\eta_t = \eta / (1 + \gamma t)^{0.5}$ . . . . . 65

4.14 The wall clock time needed to achieve the same level of the test error *thr* as a function of the number of local workers  $p$  on the *CIFAR* dataset. From left to right:  $thr = \{21\%, 20\%, 19\%, 18\%\}$ . Missing bars denote that the method never achieved specified level of test error. . . . . 68

4.15 The wall clock time needed to achieve the same level of the test error *thr* as a function of the number of local workers  $p$  on the *ImageNet* dataset. From left to right:  $thr = \{49\%, 47\%, 45\%, 43\%\}$ . Missing bars denote that the method never achieved specified level of test error. . . . . 69

5.1 The largest absolute eigenvalue of the matrix  $M$  ( $\text{sp}(M)$ ) in Equation 5.6 as a function of the learning rate  $\eta \in (0, 2)$  and the momentum rate  $\delta \in (-1, 1)$ .  $h = 1$ . . . . . 77

5.2 The largest absolute eigenvalue of the matrix  $M$  in Equation 5.12 as a function of the learning rate  $\eta \in (0, 2)$  and the moving rate  $\alpha \in (-1, 1)$ .  $h = 1$  and  $\beta = 0.9$ . . . . . 82

5.3 Three independent simulations of *EASGD* using the elastic averaging  $\alpha = \beta/p$  and the optimal  $\alpha$  given in Equation 5.17. The x-axis is the time step  $t$  in the *EASGD* updates of Equation 5.9. The y-axis is the squared distance of the center variable to the optimum zero, i.e.  $\tilde{x}_t^2$ . We have chosen  $h = 1, \sigma = 10^{-2}, p = 4, \eta = 0.1$  and  $\beta = 0.9$ . . . . . 83

5.4 The absolute value of the eigenvalues of  $z_1, z_2$  and  $z_3$  in Equation 5.19.  $\eta_h = 0.1$  and  $\beta = 0.9$ . . . . . 85

5.5 The absolute value of the eigenvalues of  $z_1, z_2$  and  $z_3$  in Equation 5.19.  $\eta_h = 1.5$  and  $\beta = 0.9$ . . . . . 85

5.6 The largest absolute eigenvalue of the matrix  $M_p$  in Equation 5.18 as a function of the learning rate  $\eta \in (0, 2)$  and the moving rate  $\alpha \in (-1, 1)$ .  $h = 1$  and  $\beta = 0.9$ . Note that we computed this spectrum using  $p = 2$ , as we have discussed in the text, it is independent of the choice of  $p$  for  $p > 1$ . 86

5.7 Three independent simulations of *EASGD* using the elastic averaging  $\alpha = \beta/p$  and the optimal  $\alpha$  given in Equation 5.17. The x-axis is the time step  $t$  in the *EASGD* updates of Equation 5.9. The y-axis is the squared distance of the center variable to the optimum zero, i.e.  $\tilde{x}_t^2$ . We have chosen  $h = 1, \sigma = 10^{-2}, p = 4, \eta = 1.5$  and  $\beta = 0.9$ . . . . . 87

5.8 The largest absolute eigenvalue of the matrix  $M_p$  in Equation 5.20 as a function of the learning rate  $\eta \in (0, 2)$  and the moving rate  $\alpha \in (-1, 1)$ .  $h = 1, \beta = 0.9$  and  $\delta = 0.99$ . Note that we computed this spectrum again using  $p = 2$ , as we have discussed in the text, it is independent of the choice of  $p$  for  $p > 1$ . . . . . 90

5.9 The probability density function of the Gamma distribution  $\Gamma(\lambda, \omega)$ . The x-axis and y-axis are both in log-scale to enlarge the singularity at zero and the decay of the tail toward infinity. . . . . 93

5.10 The largest absolute eigenvalue of the matrix  $M$  in Equation 5.30 as a function of the learning rate  $\eta \in (0, 1)$  and the momentum rate  $\delta \in (-1, 1)$ .  $\lambda = 0.5, \omega = 0.5$ . . . . . 96

5.11 The largest absolute eigenvalue of the matrix  $M$  in Equation 5.30 as a function of the learning rate  $\eta \in (0, 1)$  and the momentum rate  $\delta \in (-1, 1)$ .  $\lambda = 1, \omega = 1$ . . . . . 97

5.12 The largest absolute eigenvalue of the matrix  $M$  in Equation 5.30 as a function of the learning rate  $\eta \in (0, 1)$  and the momentum rate  $\delta \in (-1, 1)$ .  $\lambda = 2, \omega = 2$ . . . . . 98

5.13 The largest absolute eigenvalue of the matrix  $M$  in Equation 5.30 as a function of the momentum rate  $\delta \in (-1, 1)$  at  $\eta = \frac{\lambda}{\omega+1}$ .  $(\lambda, \omega) \in \{(0.5, 0.5), (1, 1), (2, 2)\}$ . . . . . 99

5.14 The largest absolute eigenvalue of the matrix  $M$  in Equation 5.30 as a function of the input Gamma distribution  $\Gamma(\lambda, \omega)$  in the range  $\omega \in (0, 100)$  and the  $\lambda \in (0, 100)$ .  $(\eta, \delta) \in \{(1, 0), (0.1, 0), (0.1, 0.9)\}$ . . . . . 100

5.15 The largest absolute eigenvalue of the matrix  $M$  in Equation 5.34 as a function of the learning rate  $\eta \in (0, 1)$  and the number of workers  $p \in [1, 64]$ .  $\lambda = 0.5, \omega = 0.5, \beta = 0.9, \alpha = \beta/p$ . . . . . 103

5.16 The largest absolute eigenvalue of the matrix  $M$  in Equation 5.34 as a function of the learning rate  $\eta \in (0, 1)$  and the number of workers  $p \in [1, 64]$ .  $\lambda = 1, \omega = 1, \beta = 0.9, \alpha = \beta/p$ . . . . . 104

5.17 The largest absolute eigenvalue of the matrix  $M$  in Equation 5.34 as a function of the learning rate  $\eta \in (0, 1)$  and the number of workers  $p \in [1, 64]$ .  $\lambda = 2, \omega = 2, \beta = 0.9, \alpha = \beta/p$ . . . . . 105

5.18 The largest absolute eigenvalue of the matrix  $M$  in Equation 5.34 as a function of the learning rate  $\eta \in (0, 2)$  and the number of workers  $p \in [1, 64]$ .  $\lambda = 10, \omega = 10, \beta = 0.9, \alpha = \beta/p$ . The minimal  $\text{sp}(M) = 0.0868$  is achieved at  $p = 29$  and  $\eta = 0.8929$ . . . . . 106

5.19 The largest absolute eigenvalue of the matrix  $M$  in Equation 5.34 as a function of the learning rate  $\eta \in (0, 1)$  and the moving rate  $\alpha \in (-1, 1)$ .  $\lambda = 0.5, \omega = 0.5, \beta = 0.9, p = 100$ . The minimal  $\text{sp}(M) = 0.5024$  is achieved at  $\eta = 0.4343$  and  $\alpha = 0.2525$ . . . . . 108

5.20 The smallest eigenvalue of the Hessian matrix  $H$  in Equation 5.38 as a function of the penalty term  $\rho$ , evaluated at the critical point  $x = \sqrt{1 - \rho}$ ,  $y = -\sqrt{1 - \rho}$ ,  $z = 0$ . . . . . 111

6.1 The behavior of *EASGD Tree* in Algorithm 6. . . . . 115

6.2 The two communication schemes of the *EASGD Tree*. . . . . 116

6.3 *EASGD Tree* on *CIFAR-lowrank* using the first communication scheme with  $\tau_1 = 10$  and  $\tau_2 = 100$ . Training loss and error, test loss and error of the root node versus a wallclock time. We run this experiment six times independently (labeled with a,b,c,d,e,f) with a same random initialization.  $p = 256, d = 16, \eta = 5e - 3, \alpha = 0.9/(d + 1)$ . . . . . 120

6.4 *EASGD Tree* on *CIFAR-lowrank* using the second communication scheme with  $\tau_u = 8$  and  $\tau_d = 80$ . Training loss and error, test loss and error of the root node versus a wallclock time. We run this experiment six times independently (labeled with a,b,c,d,e,f) with a same random initialization.  $p = 256, d = 16, \eta = 5e - 3, \alpha = 0.9/(d + 1)$ . . . . . 121

6.5 *EASGD Tree* on *CIFAR-lowrank* using the first communication scheme with  $\tau_1 = 1$  and  $\tau_2 = 10$ . Training loss and error, test loss and error of the root node versus a wallclock time. We run this experiment six times independently (labeled with a,b,c,d,e,f) with a same random initialization.  $p = 256, d = 16, \eta = 5e - 2, \alpha = 0.9/(d + 1)$ . The momentum rate  $\delta = 0$ . . . . . 122

6.6 *EASGD Tree* on *CIFAR-lowrank* using the first communication scheme with  $\tau_1 = 1$  and  $\tau_2 = 10$ . Training loss and error, test loss and error of the root node versus a wallclock time. We run this experiment six times independently (labeled with a,b,c,d,e,f) with a same random initialization.  $p = 256, d = 16, \eta = 5e - 3, \alpha = 0.9/(d + 1)$ . The momentum rate  $\delta = 0.9$ . 123

6.7 *EASGD Tree* on *CIFAR-lowrank* using the first communication scheme with  $\tau_1 = 1$  and  $\tau_2 = 10$ . Training loss and error, test loss and error of the root node versus a wallclock time. We run this experiment six times independently (labeled with a,b,c,d,e,f) with a same random initialization.  $p = 256, d = 16, \eta = 5e - 4, \alpha = 0.9/(d + 1)$ . The momentum rate  $\delta = 0.99$ . 124

6.8 *EASGD Tree* on *CIFAR-lowrank* using the second communication scheme with  $\tau_u = 1$  and  $\tau_d = 10$ . Training loss and error, test loss and error of the root node versus a wallclock time. We run this experiment six times independently (labeled with a,b,c,d,e,f) with a same random initialization.  $p = 256, d = 16, \eta = 5e - 2, \alpha = 0.9/(d + 1)$ . The momentum rate  $\delta = 0$ . . 125

6.9 *EASGD Tree* on *CIFAR-lowrank* using the second communication scheme with  $\tau_u = 1$  and  $\tau_d = 10$ . Training loss and error, test loss and error of the root node versus a wallclock time. We run this experiment six times independently (labeled with a,b,c,d,e,f) with a same random initialization.  $p = 256, d = 16, \eta = 5e - 3, \alpha = 0.9/(d + 1)$ . The momentum rate  $\delta = 0.9$ . 126

6.10 *EASGD Tree* on *CIFAR-lowrank* using the second communication scheme with  $\tau_u = 1$  and  $\tau_d = 10$ . Training loss and error, test loss and error of the root node versus a wallclock time. We run this experiment six times independently (labeled with a,b,c,d,e,f) with a same random initialization.  $p = 256, d = 16, \eta = 5e - 4, \alpha = 0.9/(d + 1)$ . The momentum rate  $\delta = 0.99$ . 127

6.11 *EASGD Tree* on *CIFAR-lowrank* using the first and second communication scheme with momentum. Training loss and error, test loss and error of the root node versus a wallclock time. The curve e in Figure 6.5, the curve e in Figure 6.6, the curve a in Figure 6.7, the curve b in Figure 6.8, the curve b in Figure 6.9 and the curve d in Figure 6.10 are selected and plotted. . . . . 128

6.12 Best test performance of *DOWNPOUR* (p=16), *EASGD* (p=16), and *EASGD Tree* (p=256) on *CIFAR-lowrank*. Training loss and error, test loss and error (of the center variable or the root node) versus a wallclock time. The momentum is not applied ( $\delta = 0$ ). . . . . 129

# List of Tables

4.1	Learning rates explored for each method shown in Figure 4.1, 4.2, 4.3 and 4.4 ( <i>CIFAR</i> experiment). . . . .	48
4.2	Learning rates explored for each method shown in Figure 4.5, 4.6 and 4.7 ( <i>CIFAR</i> experiment). . . . .	48
4.3	Learning rates explored for each method shown in Figure 4.8 and 4.9 ( <i>ImageNet</i> experiment). . . . .	48
4.4	Approximate computation time, data loading time and parameter communication time [sec] for <i>DOWNPOUR</i> (top line for $\tau = 1$ ) and <i>EASGD</i> (the time breakdown for <i>EAMSGD</i> is almost identical) (bottom line for $\tau = 10$ ). Left time corresponds to <i>CIFAR</i> experiment and right table corresponds to <i>ImageNet</i> experiment. The computation and the communication time may have some overlap, due to the MPI implementation. . .	66



# Chapter 1

## Introduction

### 1.1 What is the problem

The subject of this thesis is on how to parallelize the training of large deep learning models that use a form of stochastic gradient descent (*SGD*) [9].

The classical *SGD* method processes each data point sequentially on a single processor. As an optimization method, it often exhibits fast initial convergence toward the local optimum as compared to the batch gradient method [30]. As a learning algorithm, it often leads to better solutions in terms of the test accuracy [30, 23].

However, as the size of the dataset explodes [22], the amount of time taken to go through each data point sequentially becomes prohibitive. To meet this challenge, many distributed stochastic optimization methods have been proposed in literature, e.g. the mini-batch *SGD*, the asynchronous *SGD*, and the *ADMM*-based methods.

Meanwhile, the scale and the complexity of the deep learning models is also growing to adapt to the growth of the data. There have been attempts to parallelize the training for large-scale deep learning models on thousands of CPUs, including the Google's Distbelief system [16]. But practical image recognition systems consist of large-scale convolutional

neural networks trained on a few GPU cards sitting in a single computer [27, 49]. The main challenge is to devise parallel *SGD* algorithms to train large-scale deep learning models that yield a significant speedup when run on multiple GPU cards across multiple machines. To date, the AlphaGo system is trained using 50 GPUs for a few weeks [52].

To solve such large-scale optimization problem, one line of research is to fill-in the gap between the stochastic gradient descent method (*SGD*) and the batch gradient method [40]. The batch gradient method evaluates the gradient (in a batch) using all the data points while the stochastic gradient descent method uses only a single data point to estimate the gradient of the objective function. The mini-batch *SGD*, i.e. sampling a subset of the data points to estimate the full gradient, is one possibility to fill-in this gap [14, 17, 25, 31]. Larger mini-batch size reduces the variance of the stochastic gradient. Consequently, one can use a larger learning rate to gain a faster convergence in training. However, to implement it efficiently requires very skillful engineering efforts in order to minimize the communication overhead, which is difficult in particular for training large deep learning models on GPUs [25, 15]. It is even observed that in deep learning problems, using too large mini-batch size may lead to solutions of very poor test accuracy [62].

Another possibility is to use the asynchronous stochastic gradient descent methods [8, 29, 63, 2, 46, 16]. The idea is similar to the mini-batch *SGD*, i.e. to distribute the computation of the gradients to the local workers (processors) and collect them by the master (a parameter server), but asynchronously. The advantage of using asynchronous communication is that it allows local workers to communicate with the master at different time intervals, and thus it can significantly reduce the waiting time spent on the synchronization. The tradeoff is that asynchronous behavior results in large communication delay, which can in turn slow down the convergence rate [62].

The *DOWNPOUR* method belongs to the above class of asynchronous *SGD* methods, and is proposed for training deep learning models [16]. The main ingredient of the *DOWNPOUR* method is to reduce the (gradient) communication overhead by running

the *SGD* method on each local worker for multiple steps instead of just one single step. This idea resembles the incremental gradient method [6]. The merit is that each local worker can spend more time on the computation than on the communication. The disadvantage, however, is that the method is not very stable when the (gradient) communication period is large. We shall discuss this phenomenon further in Chapter 4.

The mini-batch *SGD* and the asynchronous *SGD* methods that we have discussed so far can be implemented in a distributed computing environment [34, 26]. However, conceptually it is still centralized as there is a master server which is needed to collect all the gradients and store the latest model parameter. There's a class of the distributed optimization methods which is conceptually decentralized. It is based on the idea of consensus averaging [57]. A classical problem is to compute the average value of the local clocks in a sensor network (aka. clock synchronization) [41]. In such setting, one needs to consider how to optimize the design of the averaging network [59] and to analyze the convergence rate on various networks subject to link failure [37, 19].

The *ADMM* (Alternating Direction Method of Multipliers) [11] method can also be used to solve the consensus averaging problem above. The basic idea is to decompose the objective function into several smaller ones so that each one can be solved separately in parallel and then be combined into one solution. In some sense, it is more effective because the dual Lagrangian update is used to close the primal-dual gap associated with the consensus constraints (i.e. to reach the consensus). *ADMM* is also generalized to the stochastic and the asynchronous setting for solving large-scale machine learning problems [42, 61]. Nevertheless, the consensus can be harder to reach, as the oscillations from the stochastic sampling and the asynchronous behavior need to be absorbed (averaged out) by the whole system (for the constraints to be satisfied).

In this thesis, we explore another dimension of such possibility. Unlike the mini-batch *SGD* and the asynchronous *SGD* method, we would like to maintain the stochastic nature (oscillation) of the *SGD* method as the number of workers grows. To accelerate

the training of the large-scale deep learning models, in particular under communication constraints, we study instead a weak consensus reaching problem. We use a central variable to average out the noise, but we only maintain a weak coupling (consensus) between the local variables and the central variable. The goal is thus very different to the consensus averaging method and the *ADMM* method.

## 1.2 Formalizing the problem

Consider minimizing a function  $F(x)$  in a parallel computing environment [7] with  $p \in \mathbb{N}$  workers and a master. In this thesis we focus on the stochastic optimization problem of the following form

$$\min_x F(x) := \mathbb{E}[f(x, \xi)], \quad (1.1)$$

where  $x$  is the model parameter to be estimated and  $\xi$  is a random variable that follows the probability distribution  $\mathbb{P}$  over  $\Omega$  such that  $F(x) = \int_{\Omega} f(x, \xi) \mathbb{P}(d\xi)$ . The optimization problem in Equation 1.1 can be reformulated as follows

$$\min_{x^1, \dots, x^p, \tilde{x}} \sum_{i=1}^p \mathbb{E}[f(x^i, \xi^i)] + \frac{\rho}{2} \|x^i - \tilde{x}\|^2, \quad (1.2)$$

where each  $\xi^i$  follows the same distribution  $\mathbb{P}$  (thus we assume each worker can sample the entire dataset). In this thesis we refer to  $x^i$ 's as local variables and we refer to  $\tilde{x}$  as a center variable.

The problem of the equivalence of these two objectives is studied in the literature and is known as the *augmentability* or the *global variable consensus* problem [24, 11]. The quadratic penalty term  $\rho$  in Equation 1.2 is expected to ensure that local workers will not fall into different attractors that are far away from the center variable.

We will focus on the problem of reducing the parameter communication overhead between the master and the local workers [50, 16, 60, 43, 48]. The problem of data communication

when the data is distributed among the workers [7, 5] is a more general problem and is not addressed in this work. We however emphasize that our problem setting is still highly non-trivial under the communication constraints due to the existence of many local optima [13].

We remark further a connection between the quadratic penalty term in Equation 1.2 and the *Moreau-Yosida* regularization [33] in convex optimization used to smooth the non-smooth part of the objective function. When using the rectified linear units [35] in deep learning models, the objective function also becomes non-smooth. This suggests that it may indeed be a good idea to use such a quadratic penalty term.

### 1.3 An overview

We now give an overview of the main focus and results of the following chapters.

Chapter 2 introduces the Elastic Averaging SGD (*EASGD*) method. We first discuss the basic idea and motivate *EASGD*. We then propose synchronous *EASGD* and discuss its connection with the classical *Jacobi* method in the numerical analysis. We further propose an asynchronous extension of *EASGD* and discuss how the asynchronous *EASGD* can be thought of as an perturbation of the synchronous *EASGD*. We end up with combining *EASGD* with the classical momentum method, notably the *Nesterov's* momentum, giving the *EAMSGD* algorithm.

Chapter 3 provides the convergence analysis of the synchronous *EASGD* algorithm. The analysis is focused on the convergence of the center variable. We discuss a one-dimensional quadratic objective function first. We show that the variance of the center variable tends to zero as the number of workers grows. Meanwhile, the variance of the local variables keeps increasing. We also introduce a double averaging sequence computing the time average of the center variable and show that it is asymptotically optimal. We further extend our analysis to the strongly convex case, and discusses the

tightness of the bound we have obtained compared to the quadratic case. Finally, we study the stability condition of *EASGD* and *ADMM* method in the round-robin scheme. We compute numerically the stability region for the *ADMM* method and show that it can be unstable when the quadratic penalty term is relatively small. This is in contrast to the stability region of the *EASGD* method.

Chapter 4 is an empirical study on the performance of various serial and parallel stochastic optimization methods for training deep convolutional neural networks on the *CIFAR* and *ImageNet* dataset. We first describe the experimental setup, in particular how we preprocess and sample the dataset by the local workers in parallel. We then compare the performance of *EASGD* and its momentum variant *EAMSGD* with the *SGD* and *DOWNPOUR* methods (including their averaging and momentum variants). We find that *EASGD* is very robust when the communication period is large, which is not the case for *DOWNPOUR*. Furthermore, *EAMSGD* (the combination of *EASGD* and Nesterov’s momentum method) achieves the best performance measured both in the wall-clock time and in the smallest achievable test error. The smallest achievable test error is further improved as we gradually increase the communication period and the number of processors. We provide further empirical results on the dependency of the learning rate and the communication period. To choose a proper communication period, we also give an explicit analysis in terms of the bandwidth requirement for the data communication and the parameter communication in the *ImageNet* case.

Chapter 5 studies the limit in speedup of various stochastic optimization methods. We would like to know what would be the limit in theory if we were given an infinite amount of processors. We study first the asymptotic phase of these methods using an additive noise model (a one-dimensional quadratic objective with Gaussian noise). The asymptotic variance can be obtained explicitly, and it can be reduced by either the mini-batch *SGD* or the *EASGD* method. However, the *SGD* method using momentum can strictly increase this asymptotic variance. On the other hand, we seek the optimal momentum rate such that the momentum *SGD* method converges the fastest. We find that this op-

timal momentum rate can either be positive or negative. We ask a similar question for the *EASGD* method by fixing the moving rate of the center variable and then optimize the moving rate of the local variable. We find that the optimal moving (average) rate of the *EASGD* method is either zero or negative. The surprising connection between these two results is that they are obtained based on a nearly same proof. We perform a similar moment analysis on the *EAMSGD* method, and find (but only numerically) that the optimal moving rate can be either positive or negative, depending on the choice of the learning rate.

We then move on to study the initial phase of these methods based on a multiplicative noise model. We introduce the Gamma distribution to parametrize the spread the input data distribution. We obtain the optimal learning rate for the mini-batch *SGD* method, and discuss how fast the optimal rate of convergence varies with the mini-batch size. We find that if the input data distribution has a large spread (to be made precise in Chapter 5), then one can gain more effective speedup by using mini-batch *SGD*. For the momentum *SGD* method, we observe that using momentum can slow down the optimal convergence rate, but it can accelerate the convergence when the learning rate is chosen to be sub-optimal. For the *EASGD* method, we observe a quite different picture to the mini-batch *SGD*. There is an optimal number of workers that *EASGD* method will achieve the best convergence rate. We also perform an asymptotic analysis when the number of workers is infinite, and show that the stability region can still be enlarged if the spread of the input data distribution is large.

We finally discuss a non-convex case to understand when *EASGD* can get trapped by a saddle point. This is the phenomenon that we have observed in Chapter 4 when the communication period is too large. We find that if the quadratic penalty term is smaller than a critical value, then the local variables can stay on both sides of a saddle point, and it is a stable configuration. This suggests that *EASGD* can spend a lot of time in such configuration if the coupling between the master and the workers is too weak.

Chapter 6 attempts to scale up the *EASGD* method to a larger number of processors. Due to the communication constraints, we propose a tree extension of the *EASGD* method. The leaf node of the tree performs the gradient descent locally and from time to time performs the elastic averaging with their parent. The root node of the tree tracks the spatial average of the variable of its children, and in turn the average of all the leaf nodes. We perform an empirical study with two different communication schemes. The first scheme exploits the fact that faster communication can be achieved at the bottom layer (between the leaf nodes and their parent) than the upper layers. The second scheme uses a faster upward communication rate and a slower downward communication rate so that the root node can be informed of the latest information from the bottom as quick as possible. We observe that the first scheme gives better training speedup, while the second scheme gives better test accuracy. One difference compared to the asynchronous *EASGD* experiment in Chapter 4 is that the communication protocol between the tree nodes is fully asynchronous so as to maximize the I/O throughput.

We end up the Chapter 6 by establishing a connection between the *DONWPOUR* and *EASGD* methods. For clarity, we focus on the synchronous scenario. These two methods can be unified by a same equation once we transform *EASGD* from the Jacobi form into the Gauss-Seidel form. The difference between the two is the choice of the moving rates. A further stability analysis shows that *DONWPOUR* has a very singular region for these rates which is separated from *EASGD* when the number of processors is large.

The last chapter concludes the thesis with a reprise of this overview, together with some open questions and directions to follow in the future. We have also made an open source project named mpiT on github to facilitate the communication using MPI under Torch, including our implementation of *DOWNPOUR* and *EASGD*.



## Chapter 2

# Elastic Averaging SGD (EASGD)

In this chapter we introduce the *Elastic Averaging SGD* method (*EASGD*) and its variants. *EASGD* is motivated by the quadratic penalty method [40], but is re-interpreted as a parallelized extension of the averaging *SGD* method [45]. The basic idea is to let each worker maintain its own local parameter, and the communication and coordination of work among the local workers is based on an elastic force which links the parameters they compute with a center variable stored by the master. The center variable is updated as a moving average where the average is taken in time and also in space over the parameters computed by local workers. The local variables are updated with the *SGD*-based methods, so as to keep the oscillations during the training process.

We discuss first the synchronous *EASGD* in Section 2.1. Then in Section 2.2 we discuss how to extend the synchronous *EASGD* to the asynchronous scenario. Finally in Section 2.3 we combine *EASGD* with two classical momentum methods in the first-order convex optimization literature, i.e. the heavy-ball method and the *Nesterov*'s momentum method, to accelerate *EASGD*. The big picture is illustrated in Figure 2.1.

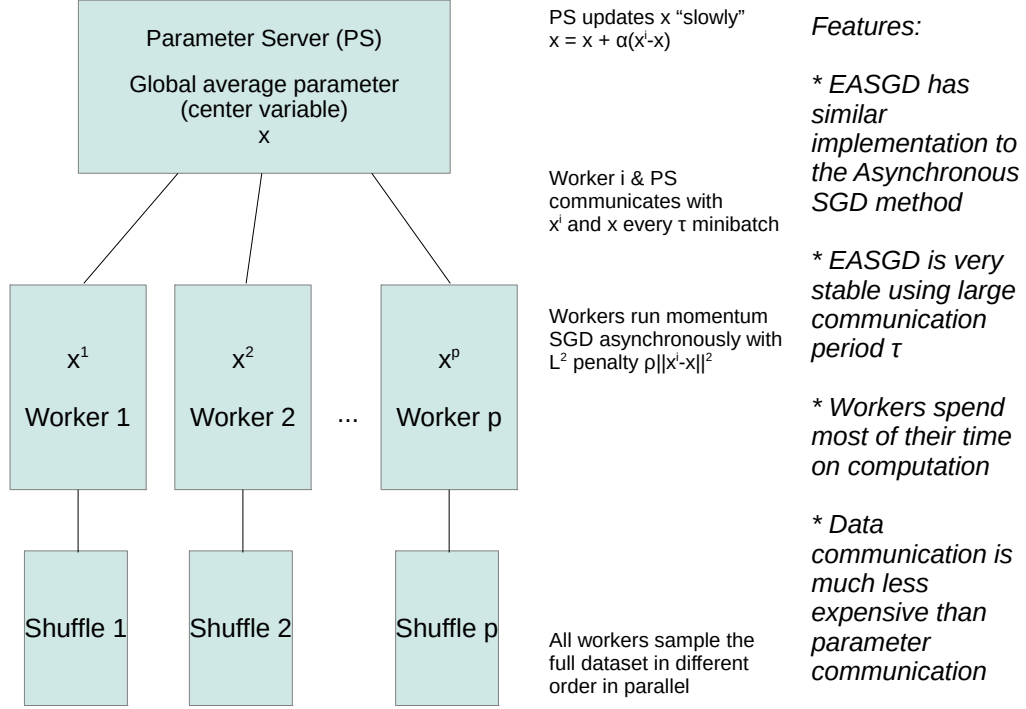


Figure 2.1: The big picture of *EASGD*.

## 2.1 Synchronous *EASGD*

The *EASGD* updates captured in resp. Equation 2.1 and 2.2 are obtained by taking the gradient descent step on the objective in Equation 1.2 with respect to resp. variable  $x^i$  and  $\tilde{x}$ ,

$$x_{t+1}^i = x_t^i - \eta(g_t^i(x_t^i) + \rho(x_t^i - \tilde{x}_t)) \quad (2.1)$$

$$\tilde{x}_{t+1} = \tilde{x}_t + \eta \sum_{i=1}^p \rho(x_t^i - \tilde{x}_t), \quad (2.2)$$

where  $g_t^i(x_t^i)$  denotes the stochastic gradient of  $F$  with respect to  $x^i$  evaluated at iteration  $t$ ,  $x_t^i$  and  $\tilde{x}_t$  denote respectively the value of variables  $x^i$  and  $\tilde{x}$  at iteration  $t$ , and  $\eta$  is the learning rate.

The update rule for the center variable  $\tilde{x}$  takes the form of moving average where the average is taken over both space and time. Denote  $\alpha = \eta\rho$  and  $\beta = p\alpha$ , then Equation 2.1 and 2.2 become

$$x_{t+1}^i = x_t^i - \eta g_t^i(x_t^i) - \alpha(x_t^i - \tilde{x}_t) \quad (2.3)$$

$$\tilde{x}_{t+1} = (1 - \beta)\tilde{x}_t + \beta \left( \frac{1}{p} \sum_{i=1}^p x_t^i \right). \quad (2.4)$$

Note that choosing  $\beta = p\alpha$  leads to an elastic symmetry in the update rule, i.e. there exists a symmetric force equal to  $\alpha(x_t^i - \tilde{x}_t)$  between the update of each  $x^i$  and  $\tilde{x}$ . It gives us a simple and intuitive reason for the the algorithm's stability over the *ADMM* method as will be explained in Section 3.3. However, this relation ( $\beta = p\alpha$ ) is by no means optimal as we shall see through our analysis in Chapter 5.

We interpret our synchronous *EASGD* as an approximate model for the asynchronous *EASGD*. Thus in order to minimize the staleness [25] of the difference  $x_t^i - \tilde{x}_t$  between the center and the local variable for the asynchronous *EASGD* (described in Algorithm 1), the update for the master in Equation 2.4 involves  $x_t^i$  instead of  $x_{t+1}^i$ . On the other hand, we can also think of our synchronous *EASGD* update rules (defined by Equation 2.3 and 2.4) as a *Jacobi* method [47]. We could have proposed a *Gauss-Seidel* version of the synchronous *EASGD* by successively updating the local and center variables with local averaging, local gradient descent, and then the global averaging. This possibility will be made precise in Chapter 6.

Note also that  $\alpha = \eta\rho$ , where the magnitude of  $\rho$  (the quadratic penalty term in Equation 1.2) represents the amount of exploration we allow in the model. In particular, small  $\rho$  allows for more exploration as it allows  $x^i$ 's to fluctuate further from the center  $\tilde{x}$ .

The distinctive idea of *EASGD* is to allow the local workers to perform more exploration (small  $\rho$ ) and the master to perform exploitation. This approach differs from other settings explored in the literature [16, 4, 8, 36, 29, 2, 46, 63], and focus on how fast the

center variable converges.

## 2.2 Asynchronous *EASGD*

We discussed the synchronous update of *EASGD* algorithm in the previous section, where the workers update the local variables in parallel such that the  $i^{\text{th}}$  worker reads the current value of the center variable and use it to update local variable  $x^i$  using Equation 2.1. All workers share the same global clock. The master has to wait for the  $x^i$  updates from all  $p$  workers before being allowed to update the value of the center variable  $\tilde{x}$  according to Equation 2.2.

In this section we propose its asynchronous variant. The local workers are still responsible for updating the local variables  $x^i$ 's, whereas the master is updating the center variable  $\tilde{x}$ . Each worker maintains its own clock  $t^i$ , which starts from 0 and is incremented by 1 after each stochastic gradient update of  $x^i$  as shown in Algorithm 1. The master performs an update whenever the local workers finished  $\tau$  steps of their gradient updates, where we refer to  $\tau$  as the *communication period*. As can be seen in Algorithm 1, whenever  $\tau$  divides the local clock of the  $i^{\text{th}}$  worker, the  $i^{\text{th}}$  worker communicates with the master and requests the current value of the center variable  $\tilde{x}$ . The worker then waits until the master sends back the requested parameter value, and computes the elastic difference  $\alpha(x - \tilde{x})$  (this entire procedure is captured in step a) in Algorithm 1). The elastic difference is then sent back to the master (step b) in Algorithm 1) who then updates  $\tilde{x}$ .

Note that the asynchronous behavior described above is partially asynchronous [7]. As in the beginning of each communication period, each worker needs to read the latest parameter from the master (blocking) and then sends the elastic difference back (blocking). Although this only involves local synchronization, one can avoid this synchronization cost by using a fully asynchronous protocol such that no waiting is necessary. The fully asynchronous protocol may however increase the network traffic and the delay in the parameter communication. We shall be more precise on this fully asynchronous protocol

when we discuss the *EASGD Tree* algorithm in Chapter 6 (Section 6.1).

Recall that we have chosen the Jacobi form in our synchronous EASGD update rules (Equation 2.3 and 2.4) as an approximate model for the asynchronous behavior. It suggests another more efficient way to realize the partially asynchronous protocol as follows. At the beginning of each communication period, each local worker sends (non-blocking) its parameter to the master, and the master will send (non-blocking) back the elastic difference once having received that local worker's parameter. During that period of time, the local worker's computation can still make progress. At the end of that communication period (i.e. all the  $\tau$  gradient updates have completed), each local worker will read (blocking) the elastic difference sent from the master, and then apply it. On the master side, it can either sum the  $p$  elastic differences altogether in one step as in the synchronous case or make an update whenever sending an elastic difference.

The communication period  $\tau$  controls the frequency of the communication between every local worker and the master, and thus the trade-off between exploration and exploitation. We show demonstrate empirically in Chapter 4 (Section 4.3.3) that in deep learning problems, too large or too small communication period can both hurt the performance.

<b>Algorithm 1:</b> Asynchronous EASGD:
Processing by worker $i$ and the master
<p><b>Input:</b> learning rate <math>\eta</math>, moving rate <math>\alpha</math>, communication period <math>\tau \in \mathbb{N}</math></p> <p><b>Initialize:</b> <math>\tilde{x}</math> is initialized randomly, <math>x^i = \tilde{x}</math>, <math>t^i = 0</math></p>
<p><b>Repeat</b></p> <p><math>x \leftarrow x^i</math></p> <p><b>if</b> (<math>\tau</math> divides <math>t^i</math>) <b>then</b></p> <p style="padding-left: 2em;"><b>a)</b> <math>x^i \leftarrow x^i - \alpha(x - \tilde{x})</math></p> <p style="padding-left: 2em;"><b>b)</b> <math>\tilde{x} \leftarrow \tilde{x} + \alpha(x - \tilde{x})</math></p> <p><b>end</b></p> <p><math>x^i \leftarrow x^i - \eta g_{t^i}^i(x)</math></p> <p><math>t^i \leftarrow t^i + 1</math></p> <p><b>Until forever</b></p>

## 2.3 Momentum *EASGD*

The momentum EASGD (*EAMSGD*) is a variant of our Algorithm 1 and is captured in Algorithm 2. It is based on the *Nesterov's* momentum scheme [39, 28, 55], where the update of the local worker of the form captured in Equation 2.1 is replaced by the following update

$$\begin{aligned}
 v_{t+1}^i &= \delta v_t^i - \eta g_t^i(x_t^i + \delta v_t^i) \\
 x_{t+1}^i &= x_t^i + v_{t+1}^i - \eta \rho(x_t^i - \tilde{x}_t),
 \end{aligned}
 \tag{2.5}$$

where  $\delta$  is the momentum rate. Note that when  $\delta = 0$  we recover the original *EASGD* algorithm.

The idea of momentum is to accelerate the slow components in the gradient descent method. The tradeoff is that we may slow down the components which were originally fast. We shall give an explicit example to illustrate this tradeoff in Chapter 5 (Section 5.2.2).

In literature, there's another well-known momentum variant called heavy-ball method (aka *Polyak's method*) [44]. The analysis of its global convergence property is still a very challenging problem in convex optimization literature [21]. If we were to combine it with *EASGD*, we would have the following update

$$\begin{aligned} v_{t+1}^i &= \delta v_t^i - \eta g_t^i(x_t^i) \\ x_{t+1}^i &= x_t^i + v_{t+1}^i - \eta \rho(x_t^i - \tilde{x}_t). \end{aligned} \tag{2.6}$$

Note that in both cases, we do not add the momentum to the center variable. One reason is that the momentum method has an error accumulation effect [18]. Due to the stochastic noise in the gradient, using momentum can actually result in higher asymptotic variance (see [32] and our discussion in Section 5.1.2). The role of the center variable is indeed to reduce the asymptotic variance.

**Algorithm 2:** Asynchronous EAMSGD:Processing by worker  $i$  and the master

**Input:** learning rate  $\eta$ , moving rate  $\alpha$ ,  
communication period  $\tau \in \mathbb{N}$ ,  
momentum term  $\delta$

**Initialize:**  $\tilde{x}$  is initialized randomly,  $x^i = \tilde{x}$ ,  
 $v^i = 0$ ,  $t^i = 0$

**Repeat** $x \leftarrow x^i$ **if** ( $\tau$  divides  $t^i$ ) **then****a)**  $x^i \leftarrow x^i - \alpha(x - \tilde{x})$ **b)**  $\tilde{x} \leftarrow \tilde{x} + \alpha(x - \tilde{x})$ **end** $v^i \leftarrow \delta v^i - \eta g_{t^i}^i(x + \delta v^i)$  $x^i \leftarrow x^i + v^i$  $t^i \leftarrow t^i + 1$ **Until forever**



## Chapter 3

# Convergence Analysis of EASGD

In this chapter, we provide the convergence analysis of the synchronous *EASGD* algorithm with constant learning rate. The analysis is focused on the convergence of the center variable to the optimum. We discuss one-dimensional quadratic case first (Lemma 3.1.1), then we introduce a double averaging sequence and prove that it is asymptotically optimal. For this, we provide two distinct proofs (one in Lemma 3.1.2 for the one-dimensional case, and the other in Lemma 3.1.3 for the multidimensional case). We extend the analysis to the strongly convex case as stated in Theorem 3.2.1. Finally, we provide stability analysis of the asynchronous *EASGD* and *ADMM* methods in the round-robin scheme in Section 3.3.

### 3.1 Quadratic case

Our analysis in the quadratic case extends the analysis of ASGD in [45]. Assume each of the  $p$  local workers  $x_t^i \in \mathbb{R}^n$  observes a noisy gradient at time  $t \geq 0$  of the linear form given in Equation 3.1.

$$g_t^i(x_t^i) = Ax_t^i - b - \xi_t^i, \quad i \in \{1, \dots, p\}, \quad (3.1)$$

where the matrix  $A$  is positive-definite (each eigenvalue is strictly positive) and  $\{\xi_t^i\}$ 's are i.i.d. random variables, with zero mean and positive-definite covariance matrix  $\Sigma$ . Let  $x^*$  denote the optimum solution, where  $x^* = A^{-1}b \in \mathbb{R}^n$ .

### 3.1.1 One-dimensional case

In this section we analyze the behavior of the mean squared error (MSE) of the center variable  $\tilde{x}_t$ , where this error is denoted as  $\mathbb{E}[\|\tilde{x}_t - x^*\|^2]$ , as a function of  $t$ ,  $p$ ,  $\eta$ ,  $\alpha$  and  $\beta$ , where  $\beta = p\alpha$ . Note that the MSE error can be decomposed as (squared) bias and variance<sup>1</sup>:  $\mathbb{E}[\|\tilde{x}_t - x^*\|^2] = \|\mathbb{E}[\tilde{x}_t - x^*]\|^2 + \mathbb{V}[\tilde{x}_t - x^*]$ . For one-dimensional case ( $n = 1$ ), we assume  $A = h > 0$  and  $\Sigma = \sigma^2 > 0$ .

**Lemma 3.1.1.** *Let  $\tilde{x}_0$  and  $\{x_0^i\}_{i=1,\dots,p}$  be arbitrary constants, then*

$$\mathbb{E}[\tilde{x}_t - x^*] = \gamma^t(\tilde{x}_0 - x^*) + \frac{\gamma^t - \phi^t}{\gamma - \phi}\alpha u_0, \quad (3.2)$$

$$\mathbb{V}[\tilde{x}_t - x^*] = \frac{p^2\alpha^2\eta^2}{(\gamma - \phi)^2} \left( \frac{\gamma^2 - \gamma^{2t}}{1 - \gamma^2} + \frac{\phi^2 - \phi^{2t}}{1 - \phi^2} - 2\frac{\gamma\phi - (\gamma\phi)^t}{1 - \gamma\phi} \right) \frac{\sigma^2}{p}, \quad (3.3)$$

where  $u_0 = \sum_{i=1}^p (x_0^i - x^* - \frac{\alpha}{1-p\alpha-\phi}(\tilde{x}_0 - x^*))$ ,  $a = \eta h + (p+1)\alpha$ ,  $c^2 = \eta h p \alpha$ ,  $\gamma = 1 - \frac{a - \sqrt{a^2 - 4c^2}}{2}$ , and  $\phi = 1 - \frac{a + \sqrt{a^2 - 4c^2}}{2}$ .

It follows from Lemma 3.1.1 that for the center variable to be stable the following has to hold

$$-1 < \phi < \gamma < 1. \quad (3.4)$$

It can be verified that  $\phi$  and  $\gamma$  are the two zero-roots of the polynomial in  $\lambda$ :  $\lambda^2 - (2 - a)\lambda + (1 - a + c^2)$ . Recall that  $\phi$  and  $\lambda$  are the functions of  $\eta$  and  $\alpha$ . Thus (see proof in Section 3.1.1)

- $\gamma < 1$  iff  $c^2 > 0$  (i.e.  $\eta > 0$  and  $\alpha > 0$ ).
- $\phi > -1$  iff  $(2 - \eta h)(2 - p\alpha) > 2\alpha$  and  $(2 - \eta h) + (2 - p\alpha) > \alpha$ .

---

<sup>1</sup>In our notation,  $\mathbb{V}$  denotes the variance.

- $\phi = \gamma$  iff  $a^2 = 4c^2$  (i.e.  $\eta h = \alpha = 0$ ).

The proof the above Lemma is based on the diagonalization of the linear gradient map (this map is symmetric due to the relation  $\beta = p\alpha$ ). The stability analysis of the asynchronous *EASGD* algorithm in the round-robin scheme is similar due to this elastic symmetry.

*Proof.* Substituting the gradient from Equation 3.1 into the update rule used by each local worker in the synchronous *EASGD* algorithm (Equation 2.3 and 2.4) we obtain

$$x_{t+1}^i = x_t^i - \eta(Ax_t^i - b - \xi_t^i) - \alpha(x_t^i - \tilde{x}_t), \quad (3.5)$$

$$\tilde{x}_{t+1} = \tilde{x}_t + \sum_{i=1}^p \alpha(x_t^i - \tilde{x}_t), \quad (3.6)$$

where  $\eta$  is the learning rate, and  $\alpha$  is the moving rate. Recall that  $\alpha = \eta\rho$  and  $A = h$ .

For the ease of notation we redefine  $\tilde{x}_t$  and  $x_t^i$  as follows:

$$\tilde{x}_t \triangleq \tilde{x}_t - x^* \quad \text{and} \quad x_t^i \triangleq x_t^i - x^*.$$

We prove the lemma by explicitly solving the linear equations 3.5 and 3.6. Let  $x_t = (x_t^1, \dots, x_t^p, \tilde{x}_t)^T$ . We rewrite the recursive relation captured in Equation 3.5 and 3.6 as simply

$$x_{t+1} = Mx_t + b_t,$$

where the drift matrix  $M$  is defined as

$$M = \begin{bmatrix} 1 - \alpha - \eta h & 0 & \dots & 0 & \alpha \\ 0 & 1 - \alpha - \eta h & 0 & \dots & \alpha \\ \dots & 0 & \dots & 0 & \dots \\ 0 & \dots & 0 & 1 - \alpha - \eta h & \alpha \\ \alpha & \alpha & \dots & \alpha & 1 - p\alpha \end{bmatrix},$$

and the (diffusion) vector  $b_t = (\eta\xi_t^1, \dots, \eta\xi_t^p, 0)^T$ .

Note that one of the eigenvalues of matrix  $M$ , that we call  $\phi$ , satisfies  $(1 - \alpha - \eta h - \phi)(1 - p\alpha - \phi) = p\alpha^2$ . The corresponding eigenvector is  $(1, 1, \dots, 1, -\frac{p\alpha}{1-p\alpha-\phi})^T$ . Let  $u_t$  be the projection of  $x_t$  onto this eigenvector. Thus  $u_t = \sum_{i=1}^p (x_t^i - \frac{\alpha}{1-p\alpha-\phi} \tilde{x}_t)$ . Let furthermore  $\xi_t = \sum_{i=1}^p \xi_t^i$ . Therefore we have

$$u_{t+1} = \phi u_t + \eta \xi_t. \quad (3.7)$$

By combining Equation 3.6 and 3.7 as follows

$$\begin{aligned} \tilde{x}_{t+1} &= \tilde{x}_t + \sum_{i=1}^p \alpha (x_t^i - \tilde{x}_t) = (1 - p\alpha) \tilde{x}_t + \alpha (u_t + \frac{p\alpha}{1-p\alpha-\phi} \tilde{x}_t) \\ &= (1 - p\alpha + \frac{p\alpha^2}{1-p\alpha-\phi}) \tilde{x}_t + \alpha u_t = \gamma \tilde{x}_t + \alpha u_t, \end{aligned}$$

where the last step results from the following relations:  $\frac{p\alpha^2}{1-p\alpha-\phi} = 1 - \alpha - \eta h - \phi$  and  $\phi + \gamma = 1 - \alpha - \eta h + 1 - p\alpha$ . Thus we obtained

$$\tilde{x}_{t+1} = \gamma \tilde{x}_t + \alpha u_t. \quad (3.8)$$

Based on Equation 3.7 and 3.8, we can then expand  $u_t$  and  $\tilde{x}_t$  recursively,

$$u_{t+1} = \phi^{t+1} u_0 + \phi^t (\eta \xi_0) + \dots + \phi^0 (\eta \xi_t), \quad (3.9)$$

$$\tilde{x}_{t+1} = \gamma^{t+1} \tilde{x}_0 + \gamma^t (\alpha u_0) + \dots + \gamma^0 (\alpha u_t). \quad (3.10)$$

Substituting  $u_0, u_1, \dots, u_t$ , each given through Equation 3.9, into Equation 3.10 we obtain

$$\tilde{x}_t = \gamma^t \tilde{x}_0 + \frac{\gamma^t - \phi^t}{\gamma - \phi} \alpha u_0 + \alpha \eta \sum_{l=1}^{t-1} \frac{\gamma^{t-l} - \phi^{t-l}}{\gamma - \phi} \xi_{l-1}. \quad (3.11)$$

To be more specific, the Equation 3.11 is obtained by interchanging the order of sum-

mation,

$$\begin{aligned}
\tilde{x}_{t+1} &= \gamma^{t+1}\tilde{x}_0 + \sum_{i=0}^t \gamma^{t-i}(\alpha u_i) \\
&= \gamma^{t+1}\tilde{x}_0 + \sum_{i=0}^t \gamma^{t-i}(\alpha(\phi^i u_0 + \sum_{l=0}^{i-1} \phi^{i-1-l}\eta\xi_l)) \\
&= \gamma^{t+1}\tilde{x}_0 + \sum_{i=0}^t \gamma^{t-i}\phi^i(\alpha u_0) + \sum_{l=0}^{t-1} \sum_{i=l+1}^t \gamma^{t-i}\phi^{i-1-l}(\alpha\eta\xi_l) \\
&= \gamma^{t+1}\tilde{x}_0 + \frac{\gamma^{t+1} - \phi^{t+1}}{\gamma - \phi}(\alpha u_0) + \sum_{l=0}^{t-1} \frac{\gamma^{t-l} - \phi^{t-l}}{\gamma - \phi}(\alpha\eta\xi_l).
\end{aligned}$$

Since the random variables  $\xi_l$  are i.i.d, we may sum the variance term by term as follows

$$\begin{aligned}
\sum_{l=0}^{t-1} \left( \frac{\gamma^{t-l} - \phi^{t-l}}{\gamma - \phi} \right)^2 &= \sum_{l=0}^{t-1} \frac{\gamma^{2(t-l)} - 2\gamma^{t-l}\phi^{t-l} + \phi^{2(t-l)}}{(\gamma - \phi)^2} \\
&= \frac{1}{(\gamma - \phi)^2} \left( \frac{\gamma^2 - \gamma^{2(t+1)}}{1 - \gamma^2} - 2\frac{\gamma\phi - (\gamma\phi)^{t+1}}{1 - \gamma\phi} + \frac{\phi^2 - \phi^{2(t+1)}}{1 - \phi^2} \right).
\end{aligned} \tag{3.12}$$

Note that  $\mathbb{E}[\xi_t] = \sum_{i=1}^p \mathbb{E}[\xi_t^i] = 0$  and  $\mathbb{V}[\xi_t] = \sum_{i=1}^p \mathbb{V}[\xi_t^i] = p\sigma^2$ . These two facts, the equality in Equation 3.11 and Equation 3.12 can then be used to compute  $\mathbb{E}[\tilde{x}_t]$  and  $\mathbb{V}[\tilde{x}_t]$  as given in Equation 3.2 and 3.3 in Lemma 3.1.1.  $\square$

### Visualizing Lemma 3.1.1

In Figure 3.1, we illustrate the dependence of MSE on  $\beta$ ,  $\eta$  and the number of processors  $p$  over time  $t$ . We consider the large-noise setting where  $\tilde{x}_0 = x_0^i = 1$ ,  $h = 1$  and  $\sigma = 10$ . The MSE error is color-coded such that the deep blue color corresponds to the MSE equal to  $10^{-3}$ , the green color corresponds to the MSE equal to 1, the red color corresponds to MSE equal to  $10^3$  and the dark red color corresponds to the divergence of algorithm *EASGD* (condition in Equation 3.4 is then violated). The plot shows that we can achieve significant variance reduction by increasing the number of local workers

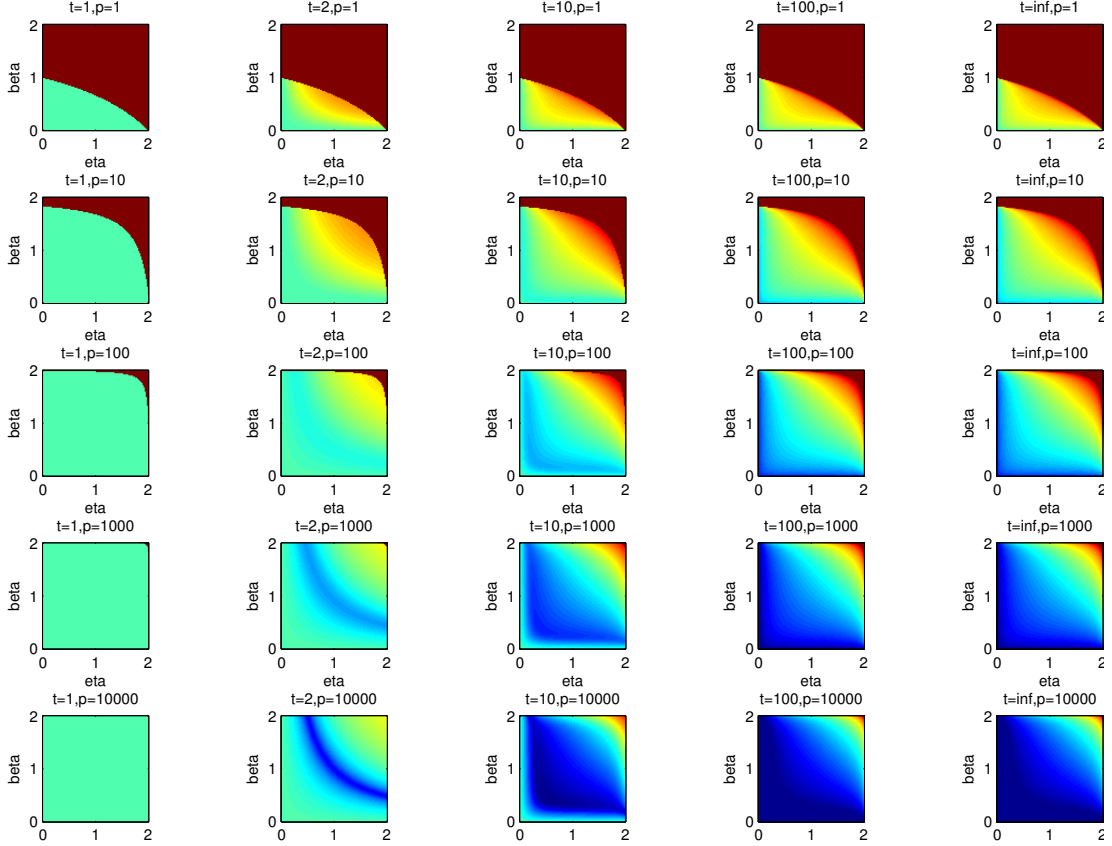


Figure 3.1: Theoretical mean squared error (MSE) of the center  $\tilde{x}$  in the quadratic case, with various choices of the learning rate  $\eta$  (horizontal within each block), and the moving rate  $\beta = p\alpha$  (vertical within each block), the number of processors  $p = \{1, 10, 100, 1000, 10000\}$  (vertical across blocks), and the time steps  $t = \{1, 2, 10, 100, \infty\}$  (horizontal across blocks). The MSE is plotted in log scale, ranging from  $10^{-3}$  to  $10^3$  (from deep blue to red). The dark red (i.e. on the upper-right corners) indicates divergence.

$p$ . This effect is less sensitive to the choice of  $\beta$  and  $\eta$  for large  $p$ .

### Condition in Equation 3.4

We are going to show that

- $\gamma < 1$  iff  $c^2 > 0$  (i.e.  $\eta > 0$  and  $\beta > 0$ ).
- $\phi > -1$  iff  $(2 - \eta h)(2 - \beta) > 2\beta/p$  and  $(2 - \eta h) + (2 - \beta) > \beta/p$ .
- $\phi = \gamma$  iff  $a^2 = 4c^2$  (i.e.  $\eta h = \beta = 0$ ).

Recall that  $a = \eta h + (p + 1)\alpha$ ,  $c^2 = \eta h p \alpha$ ,  $\gamma = 1 - \frac{a - \sqrt{a^2 - 4c^2}}{2}$ ,  $\phi = 1 - \frac{a + \sqrt{a^2 - 4c^2}}{2}$ , and  $\beta = p\alpha$ . We have

- $\gamma < 1 \Leftrightarrow \frac{a - \sqrt{a^2 - 4c^2}}{2} > 0 \Leftrightarrow a > \sqrt{a^2 - 4c^2} \Leftrightarrow a^2 > a^2 - 4c^2 \Leftrightarrow c^2 > 0$ .
- $\phi > -1 \Leftrightarrow 2 > \frac{a + \sqrt{a^2 - 4c^2}}{2} \Leftrightarrow 4 - a > \sqrt{a^2 - 4c^2} \Leftrightarrow 4 - a > 0, (4 - a)^2 > a^2 - 4c^2 \Leftrightarrow 4 - a > 0, 4 - 2a + c^2 > 0 \Leftrightarrow 4 > \eta h + \beta + \alpha, 4 - 2(\eta h + \beta + \alpha) + \eta h \beta > 0$ .
- $\phi = \gamma \Leftrightarrow \sqrt{a^2 - 4c^2} = 0 \Leftrightarrow a^2 = 4c^2$ .

The next corollary is a consequence of Lemma 3.1.1. As the number of workers  $p$  grows, the averaging property of the *EASGD* can be characterized as follows

**Corollary 3.1.1.** *Let the Elastic Averaging relation  $\beta = p\alpha$  and the condition 3.4 hold, then*

$$\lim_{p \rightarrow \infty} \lim_{t \rightarrow \infty} p \mathbb{E}[(\tilde{x}_t - x^*)^2] = \frac{\beta \eta h}{(2 - \beta)(2 - \eta h)} \cdot \frac{2 - \beta - \eta h + \beta \eta h}{\beta + \eta h - \beta \eta h} \cdot \frac{\sigma^2}{h^2}.$$

*Proof.* Note that when  $\beta$  is fixed,  $\lim_{p \rightarrow \infty} a = \eta h + \beta$  and  $c^2 = \eta h \beta$ . Then  $\lim_{p \rightarrow \infty} \phi = \min(1 - \beta, 1 - \eta h)$  and  $\lim_{p \rightarrow \infty} \gamma = \max(1 - \beta, 1 - \eta h)$ . Also note that using Lemma 3.1.1

we obtain

$$\begin{aligned}
\lim_{t \rightarrow \infty} \mathbb{E}[(\tilde{x}_t - x^*)^2] &= \frac{\beta^2 \eta^2}{(\gamma - \phi)^2} \left( \frac{\gamma^2}{1 - \gamma^2} + \frac{\phi^2}{1 - \phi^2} - \frac{2\gamma\phi}{1 - \gamma\phi} \right) \frac{\sigma^2}{p} \\
&= \frac{\beta^2 \eta^2}{(\gamma - \phi)^2} \left( \frac{\gamma^2(1 - \phi^2)(1 - \phi\gamma) + \phi^2(1 - \gamma^2)(1 - \phi\gamma) - 2\gamma\phi(1 - \gamma^2)(1 - \phi^2)}{(1 - \gamma^2)(1 - \phi^2)(1 - \gamma\phi)} \right) \frac{\sigma^2}{p} \\
&= \frac{\beta^2 \eta^2}{(\gamma - \phi)^2} \left( \frac{(\gamma - \phi)^2(1 + \gamma\phi)}{(1 - \gamma^2)(1 - \phi^2)(1 - \gamma\phi)} \right) \frac{\sigma^2}{p} \\
&= \frac{\beta^2 \eta^2}{(1 - \gamma^2)(1 - \phi^2)} \cdot \frac{1 + \gamma\phi}{1 - \gamma\phi} \cdot \frac{\sigma^2}{p}.
\end{aligned}$$

Corollary 3.1.1 is obtained by plugging in the limiting values of  $\phi$  and  $\gamma$ .  $\square$

The crucial point of Corollary 3.1.1 is that the MSE in the limit  $t \rightarrow \infty$  is in the order of  $1/p$  which implies that as the number of processors  $p$  grows, the MSE will decrease for the *EASGD* algorithm. Also note that the smaller the  $\beta$  is (recall that  $\beta = p\alpha = p\eta\rho$ ), the more exploration is allowed (small  $\rho$ ) and simultaneously the smaller the MSE is.

The next lemma (Lemma 3.1.2) shows that *EASGD* algorithm achieves the highest possible rate of convergence when we consider the double averaging sequence (similarly to [45])  $\{z_1, z_2, \dots\}$  defined as below

$$z_{t+1} = \frac{1}{t+1} \sum_{k=0}^t \tilde{x}_k. \quad (3.13)$$

**Lemma 3.1.2** (Weak convergence). *If the condition in Equation 3.4 holds, then the normalized double averaging sequence defined in Equation 3.13 converges weakly to the normal distribution with zero mean and variance  $\sigma^2/ph^2$ ,*

$$\sqrt{t}(z_t - x^*) \rightarrow \mathcal{N}\left(0, \frac{\sigma^2}{ph^2}\right), \quad t \rightarrow \infty. \quad (3.14)$$

*Proof.* As in the proof of Lemma 3.1.1, for the ease of notation we redefine  $\tilde{x}_t$  and  $x_t^i$  as follows:

$$\tilde{x}_t \triangleq \tilde{x}_t - x^* \quad \text{and} \quad x_t^i \triangleq x_t^i - x^*.$$



Also recall that  $\{\xi_t^i\}$ 's are *i.i.d.* random variables (noise) with zero mean and the same positive definite covariance matrix  $\Sigma \succ 0$ . We are interested in the asymptotic behavior of the double averaging sequence  $\{z_1, z_2, \dots\}$  defined as

$$z_{t+1} = \frac{1}{t+1} \sum_{k=0}^t \tilde{x}_k. \quad (3.15)$$

Recall the Equation 3.11 from the proof of Lemma 3.1.1 (for the convenience it is provided below):

$$\tilde{x}_k = \gamma^k \tilde{x}_0 + \alpha u_0 \frac{\gamma^k - \phi^k}{\gamma - \phi} + \alpha \eta \sum_{l=1}^{k-1} \frac{\gamma^{k-l} - \phi^{k-l}}{\gamma - \phi} \xi_{l-1},$$

where  $\xi_t = \sum_{i=1}^p \xi_t^i$ . Therefore

$$\begin{aligned} \sum_{k=0}^t \tilde{x}_k &= \frac{1 - \gamma^{t+1}}{1 - \gamma} \tilde{x}_0 + \alpha u_0 \frac{1}{\gamma - \mu} \left( \frac{1 - \gamma^{t+1}}{1 - \gamma} - \frac{1 - \phi^{t+1}}{1 - \phi} \right) + \alpha \eta \sum_{l=1}^{t-1} \sum_{k=l+1}^t \frac{\gamma^{k-l} - \phi^{k-l}}{\gamma - \phi} \xi_{l-1} \\ &= O(1) + \alpha \eta \sum_{l=1}^{t-1} \frac{1}{\gamma - \phi} \left( \gamma \frac{1 - \gamma^{t-l}}{1 - \gamma} - \phi \frac{1 - \phi^{t-l}}{1 - \phi} \right) \xi_{l-1} \end{aligned}$$

Note that the only non-vanishing term (in weak convergence) of  $1/\sqrt{t} \sum_{k=0}^t \tilde{x}_k$  as  $t \rightarrow \infty$  is

$$\frac{1}{\sqrt{t}} \alpha \eta \sum_{l=1}^{t-1} \frac{1}{\gamma - \phi} \left( \frac{\gamma}{1 - \gamma} - \frac{\phi}{1 - \phi} \right) \xi_{l-1}. \quad (3.16)$$

Also recall that  $\mathbb{V}[\xi_{l-1}] = p\sigma^2$  and

$$\frac{1}{\gamma - \phi} \left( \frac{\gamma}{1 - \gamma} - \frac{\phi}{1 - \phi} \right) = \frac{1}{(1 - \gamma)(1 - \phi)} = \frac{1}{\eta h p \alpha}.$$

Therefore the expression in Equation 3.16 is asymptotically normal with zero mean and variance  $\sigma^2/p h^2$ .  $\square$

### 3.1.2 Generalization to multidimensional case

The asymptotic variance in the Lemma 3.1.2 is optimal with any fixed  $\eta$  and  $\beta$  for which Equation 3.4 holds. The next lemma (Lemma 3.1.3) extends the result in Lemma 3.1.2 to the multi-dimensional setting.

**Lemma 3.1.3** (Weak convergence). *Let  $h$  denotes the largest eigenvalue of  $A$ . If  $(2 - \eta h)(2 - \beta) > 2\beta/p$ ,  $(2 - \eta h) + (2 - \beta) > \beta/p$ ,  $\eta > 0$  and  $\beta > 0$ , then the normalized double averaging sequence converges weakly to the normal distribution with zero mean and the covariance matrix  $V = A^{-1}\Sigma(A^{-1})^T$ ,*

$$\sqrt{tp}(z_t - x^*) \rightharpoonup \mathcal{N}(0, V), \quad t \rightarrow \infty. \quad (3.17)$$

*Proof.* Since  $A$  is symmetric, one can use the proof technique of Lemma 3.1.2 to prove Lemma 3.1.3 by diagonalizing the matrix  $A$ . This diagonalization essentially generalizes Lemma 3.1.1 to the multidimensional case. We will not go into the details of this proof as we will provide a simpler way to look at the system. As in the proof of Lemma 3.1.1 and Lemma 3.1.2, for the ease of notation we redefine  $\tilde{x}_t$  and  $x_t^i$  as follows:

$$\tilde{x}_t \triangleq \tilde{x}_t - x^* \quad \text{and} \quad x_t^i \triangleq x_t^i - x^*.$$

Let the spatial average of the local parameters at time  $t$  be denoted as  $y_t$  where  $y_t = \frac{1}{p} \sum_{i=1}^p x_t^i$ , and let the average noise be denoted as  $\xi_t$ , where  $\xi_t = \frac{1}{p} \sum_{i=1}^p \xi_t^i$ . Equations 3.5 and 3.6 can then be reduced to the following

$$y_{t+1} = y_t - \eta(Ay_t - \xi_t) + \alpha(\tilde{x}_t - y_t), \quad (3.18)$$

$$\tilde{x}_{t+1} = \tilde{x}_t + \beta(y_t - \tilde{x}_t). \quad (3.19)$$

We focus on the case where the learning rate  $\eta$  and the moving rate  $\alpha$  are kept constant over time<sup>2</sup>. Recall  $\beta = p\alpha$  and  $\alpha = \eta\rho$ .

---

<sup>2</sup>As a side note, notice that the center parameter  $\tilde{x}_t$  is tracking the spatial average  $y_t$  of the local

Let's introduce the block notation  $U_t = (y_t, \tilde{x}_t)$ ,  $\Xi_t = (\eta\xi_t, 0)$ ,  $M = I - \eta L$  and

$$L = \begin{pmatrix} A + \frac{\alpha}{\eta}I & -\frac{\alpha}{\eta}I \\ -\frac{\beta}{\eta}I & \frac{\beta}{\eta}I \end{pmatrix}.$$

From Equations 3.18 and 3.19 it follows that  $U_{t+1} = MU_t + \Xi_t$ . Note that this linear system has a degenerate noise  $\Xi_t$  which prevents us from directly applying results of [45].

Expanding this recursive relation and summing by parts, we have

$$\begin{aligned} \sum_{k=0}^t U_k &= M^0 U_0 + \\ &M^1 U_0 + M^0 \Xi_0 + \\ &M^2 U_0 + M^1 \Xi_0 + M^0 \Xi_1 + \\ &\dots \\ &M^t U_0 + M^{t-1} \Xi_0 + \dots + M^0 \Xi_{t-1}. \end{aligned}$$

By Lemma 3.1.4,  $\|M\|_2 < 1$  and thus

$$M^0 + M^1 + \dots + M^t + \dots = (I - M)^{-1} = \eta^{-1} L^{-1}.$$

Since  $A$  is invertible, we get

$$L^{-1} = \begin{pmatrix} A^{-1} & \frac{\alpha}{\beta} A^{-1} \\ A^{-1} & \frac{\eta}{\beta} + \frac{\alpha}{\beta} A^{-1} \end{pmatrix},$$

thus

$$\frac{1}{\sqrt{t}} \sum_{k=0}^t U_k = \frac{1}{\sqrt{t}} U_0 + \frac{1}{\sqrt{t}} \eta L^{-1} \sum_{k=1}^t \Xi_{k-1} - \frac{1}{\sqrt{t}} \sum_{k=1}^t M^{k+1} \Xi_{k-1}.$$

---

parameters with a non-symmetric spring in Equation 3.18 and 3.19. To be more precise note that the update on  $y_{t+1}$  contains  $(\tilde{x}_t - y_t)$  scaled by  $\alpha$ , whereas the update on  $\tilde{x}_{t+1}$  contains  $-(\tilde{x}_t - y_t)$  scaled by  $\beta$ . Since  $\alpha = \beta/p$  the impact of the center  $\tilde{x}_{t+1}$  on the spatial local average  $y_{t+1}$  becomes more negligible as  $p$  grows.

Note that the only non-vanishing term of  $\frac{1}{\sqrt{t}} \sum_{k=0}^t U_k$  is  $\frac{1}{\sqrt{t}}(\eta L)^{-1} \sum_{k=1}^t \Xi_{k-1}$ , thus in weak convergence we have

$$\frac{1}{\sqrt{t}}(\eta L)^{-1} \sum_{k=1}^t \Xi_{k-1} \rightharpoonup \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} V & V \\ V & V \end{pmatrix}\right), \quad (3.20)$$

where  $V = A^{-1}\Sigma(A^{-1})^T$ . □

**Lemma 3.1.4.** *If the following conditions hold:*

$$\begin{aligned} (2 - \eta h)(2 - p\alpha) &> 2\alpha \\ (2 - \eta h) + (2 - p\alpha) &> \alpha \\ \eta &> 0 \\ \alpha &> 0 \end{aligned}$$

then  $\|M\|_2 < 1$ .

*Proof.* The eigenvalue  $\lambda$  of  $M$  and the (non-zero) eigenvector  $(y, z)$  of  $M$  satisfy

$$M \begin{pmatrix} y \\ z \end{pmatrix} = \lambda \begin{pmatrix} y \\ z \end{pmatrix}. \quad (3.21)$$

Recall that

$$M = I - \eta L = \begin{pmatrix} I - \eta A - \alpha I & \alpha I \\ \beta I & I - \beta I \end{pmatrix}. \quad (3.22)$$

From the Equations 3.21 and 3.22 we obtain

$$\begin{cases} y - \eta Ay - \alpha y + \alpha z = \lambda y \\ \beta y + (1 - \beta)z = \lambda z \end{cases}. \quad (3.23)$$

Since  $(y, z)$  is assumed to be non-zero, we can write  $z = \beta y / (\lambda + \beta - 1)$ . Then the Equation 3.23 can be reduced to

$$\eta Ay = (1 - \alpha - \lambda)y + \frac{\alpha\beta}{\lambda + \beta - 1}y. \quad (3.24)$$

Thus  $y$  is the eigenvector of  $A$ . Let  $\lambda_A$  be the eigenvalue of matrix  $A$  such that  $Ay = \lambda_A y$ . Thus based on Equation 3.24 it follows that

$$\eta\lambda_A = (1 - \alpha - \lambda) + \frac{\alpha\beta}{\lambda + \beta - 1}. \quad (3.25)$$

Equation 3.25 is equivalent to

$$\lambda^2 - (2 - a)\lambda + (1 - a + c^2) = 0, \quad (3.26)$$

where  $a = \eta\lambda_A + (p + 1)\alpha$ ,  $c^2 = \eta\lambda_A p \alpha$ . It follows from the condition in Equation 3.4 that  $-1 < \lambda < 1$  iff  $\eta > 0$ ,  $\beta > 0$ ,  $(2 - \eta\lambda_A)(2 - \beta) > 2\beta/p$  and  $(2 - \eta\lambda_A) + (2 - \beta) > \beta/p$ . Let  $h$  denote the maximum eigenvalue of  $A$  and note that  $2 - \eta\lambda_A \geq 2 - \eta h$ . This implies that the condition of our lemma is sufficient.  $\square$

As in Lemma 3.1.2, the asymptotic covariance in the Lemma 3.1.3 is optimal, i.e. meets the Fisher information lower-bound. The fact that this asymptotic covariance matrix  $V$  does not contain any term involving  $\rho$  is quite remarkable, since the penalty term  $\rho$  does have an impact on the condition number of the Hessian in Equation 1.2.

## 3.2 Strongly convex case

We now extend the above proof ideas to analyze the strongly convex case, in which the noisy gradient  $g_t^i(x) = \nabla F(x) - \xi_t^i$  has the regularity that there exists some  $0 < \mu \leq L$ , for which  $\mu \|x - y\|^2 \leq \langle \nabla F(x) - \nabla F(y), x - y \rangle \leq L \|x - y\|^2$  holds uniformly for any  $x \in \mathbb{R}^d, y \in \mathbb{R}^d$ . The noise  $\{\xi_t^i\}$ 's is assumed to be i.i.d. with zero mean and bounded

variance  $\mathbb{E}[\|\xi_t^i\|^2] \leq \sigma^2$ .

**Theorem 3.2.1.** Let  $a_t = \mathbb{E} \left\| \frac{1}{p} \sum_{i=1}^p x_t^i - x^* \right\|^2$ ,  $b_t = \frac{1}{p} \sum_{i=1}^p \mathbb{E} \|x_t^i - x^*\|^2$ ,  $c_t = \mathbb{E} \|\tilde{x}_t - x^*\|^2$ ,  $\gamma_1 = 2\eta \frac{\mu L}{\mu + L}$  and  $\gamma_2 = 2\eta L(1 - \frac{2\sqrt{\mu L}}{\mu + L})$ . If  $0 \leq \eta \leq \frac{2}{\mu + L}(1 - \alpha)$ ,  $0 \leq \alpha < 1$  and  $0 \leq \beta \leq 1$  then

$$\begin{pmatrix} a_{t+1} \\ b_{t+1} \\ c_{t+1} \end{pmatrix} \leq \begin{pmatrix} 1 - \gamma_1 - \gamma_2 - \alpha & \gamma_2 & \alpha \\ 0 & 1 - \gamma_1 - \alpha & \alpha \\ \beta & 0 & 1 - \beta \end{pmatrix} \begin{pmatrix} a_t \\ b_t \\ c_t \end{pmatrix} + \begin{pmatrix} \eta^2 \frac{\sigma^2}{p} \\ \eta^2 \sigma^2 \\ 0 \end{pmatrix}.$$

*Proof.* The idea of the proof is based on the point of view in Lemma 3.1.3, i.e. how close the center variable  $\tilde{x}_t$  is to the spatial average of the local variables  $y_t = \frac{1}{p} \sum_{i=1}^p x_t^i$ . To further simplify the notation, let the noisy gradient be  $\nabla f_{t,\xi}^i = g_t^i(x_t^i) = \nabla F(x_t^i) - \xi_t^i$ , and  $\nabla f_t^i = \nabla F(x_t^i)$  be its deterministic part. Then *EASGD* updates can be rewritten as follows,

$$x_{t+1}^i = x_t^i - \eta \nabla f_{t,\xi}^i - \alpha(x_t^i - \tilde{x}_t), \quad (3.27)$$

$$\tilde{x}_{t+1} = \tilde{x}_t + \beta(y_t - \tilde{x}_t). \quad (3.28)$$

We have thus the update for the spatial average,

$$y_{t+1} = y_t - \eta \frac{1}{p} \sum_{i=1}^p \nabla f_{t,\xi}^i - \alpha(y_t - \tilde{x}_t). \quad (3.29)$$

The idea of the proof is to bound the distance  $\|\tilde{x}_t - x^*\|^2$  through  $\|y_t - x^*\|^2$  and  $\frac{1}{p} \sum_i \|x_t^i - x^*\|^2$ . We start from the following estimate for the strongly convex function [38],

$$\langle \nabla F(x) - \nabla F(y), x - y \rangle \geq \frac{\mu L}{\mu + L} \|x - y\|^2 + \frac{1}{\mu + L} \|\nabla F(x) - \nabla F(y)\|^2.$$

Since  $\nabla f(x^*) = 0$ , we have

$$\langle \nabla f_t^i, x_t^i - x^* \rangle \geq \frac{\mu L}{\mu + L} \|x_t^i - x^*\|^2 + \frac{1}{\mu + L} \|\nabla f_t^i\|^2. \quad (3.30)$$

From Equation 3.27 the following relation holds,

$$\begin{aligned} \|x_{t+1}^i - x^*\|^2 &= \|x_t^i - x^*\|^2 + \eta^2 \|\nabla f_{t,\xi}^i\|^2 + \alpha^2 \|x_t^i - \tilde{x}_t\|^2 \\ &\quad - 2\eta \langle \nabla f_{t,\xi}^i, x_t^i - x^* \rangle - 2\alpha \langle x_t^i - \tilde{x}_t, x_t^i - x^* \rangle \\ &\quad + 2\eta\alpha \langle \nabla f_{t,\xi}^i, x_t^i - \tilde{x}_t \rangle. \end{aligned} \quad (3.31)$$

By the cosine rule ( $2 \langle a - b, c - d \rangle = \|a - d\|^2 - \|a - c\|^2 + \|c - b\|^2 - \|d - b\|^2$ ), we have

$$2 \langle x_t^i - \tilde{x}_t, x_t^i - x^* \rangle = \|x_t^i - x^*\|^2 + \|x_t^i - \tilde{x}_t\|^2 - \|\tilde{x}_t - x^*\|^2. \quad (3.32)$$

By the Cauchy-Schwarz inequality, we have

$$\langle \nabla f_t^i, x_t^i - \tilde{x}_t \rangle \leq \|\nabla f_t^i\| \|x_t^i - \tilde{x}_t\|. \quad (3.33)$$

Combining the above estimates in Equations 3.30, 3.31, 3.32, 3.33, we obtain

$$\begin{aligned} \|x_{t+1}^i - x^*\|^2 &\leq \|x_t^i - x^*\|^2 + \eta^2 \|\nabla f_t^i - \xi_t^i\|^2 + \alpha^2 \|x_t^i - \tilde{x}_t\|^2 \\ &\quad - 2\eta \left( \frac{\mu L}{\mu + L} \|x_t^i - x^*\|^2 + \frac{1}{\mu + L} \|\nabla f_t^i\|^2 \right) + 2\eta \langle \xi_t^i, x_t^i - x^* \rangle \\ &\quad - \alpha (\|x_t^i - x^*\|^2 + \|x_t^i - \tilde{x}_t\|^2 - \|\tilde{x}_t - x^*\|^2) \\ &\quad + 2\eta\alpha \|\nabla f_t^i\| \|x_t^i - \tilde{x}_t\| - 2\eta\alpha \langle \xi_t^i, x_t^i - \tilde{x}_t \rangle. \end{aligned} \quad (3.34)$$

Choosing  $0 \leq \alpha < 1$ , we can have this upper-bound for the terms  $\alpha^2 \|x_t^i - \tilde{x}_t\|^2 - \alpha \|x_t^i - \tilde{x}_t\|^2 + 2\eta\alpha \|\nabla f_t^i\| \|x_t^i - \tilde{x}_t\| = -\alpha(1 - \alpha) \|x_t^i - \tilde{x}_t\|^2 + 2\eta\alpha \|\nabla f_t^i\| \|x_t^i - \tilde{x}_t\| \leq \frac{\eta^2\alpha}{1-\alpha} \|\nabla f_t^i\|^2$  by applying  $-ax^2 + bx \leq \frac{b^2}{4a}$  with  $x = \|x_t^i - \tilde{x}_t\|$ . Thus we can further bound

Equation 3.34 with

$$\begin{aligned} \|x_{t+1}^i - x^*\|^2 &\leq (1 - 2\eta \frac{\mu L}{\mu + L} - \alpha) \|x_t^i - x^*\|^2 + (\eta^2 + \frac{\eta^2 \alpha}{1 - \alpha} - \frac{2\eta}{\mu + L}) \|\nabla f_t^i\|^2 \\ &\quad - 2\eta^2 \langle \nabla f_t^i, \xi_t^i \rangle + 2\eta \langle \xi_t^i, x_t^i - x^* \rangle - 2\eta \alpha \langle \xi_t^i, x_t^i - \tilde{x}_t \rangle \end{aligned} \quad (3.35)$$

$$+ \eta^2 \|\xi_t^i\|^2 + \alpha \|\tilde{x}_t - x^*\|^2 \quad (3.36)$$

As in Equation 3.35 and 3.36, the noise  $\xi_t^i$  is zero mean ( $\mathbb{E}\xi_t^i = 0$ ) and the variance of the noise  $\xi_t^i$  is bounded ( $\mathbb{E}\|\xi_t^i\|^2 \leq \sigma^2$ ), if  $\eta$  is chosen small enough such that  $\eta^2 + \frac{\eta^2 \alpha}{1 - \alpha} - \frac{2\eta}{\mu + L} \leq 0$ , then

$$\mathbb{E} \|x_{t+1}^i - x^*\|^2 \leq (1 - 2\eta \frac{\mu L}{\mu + L} - \alpha) \mathbb{E} \|x_t^i - x^*\|^2 + \eta^2 \sigma^2 + \alpha \mathbb{E} \|\tilde{x}_t - x^*\|^2 \quad (3.37)$$

Now we apply similar idea to estimate  $\|y_t - x^*\|^2$ . From Equation 3.29 the following relation holds,

$$\begin{aligned} \|y_{t+1} - x^*\|^2 &= \|y_t - x^*\|^2 + \eta^2 \left\| \frac{1}{p} \sum_{i=1}^p \nabla f_{t,\xi}^i \right\|^2 + \alpha^2 \|y_t - \tilde{x}_t\|^2 \\ &\quad - 2\eta \left\langle \frac{1}{p} \sum_{i=1}^p \nabla f_{t,\xi}^i, y_t - x^* \right\rangle - 2\alpha \langle y_t - \tilde{x}_t, y_t - x^* \rangle \\ &\quad + 2\eta \alpha \left\langle \frac{1}{p} \sum_{i=1}^p \nabla f_{t,\xi}^i, y_t - \tilde{x}_t \right\rangle. \end{aligned} \quad (3.38)$$

By  $\left\langle \frac{1}{p} \sum_{i=1}^p a_i, \frac{1}{p} \sum_{j=1}^p b_j \right\rangle = \frac{1}{p} \sum_{i=1}^p \langle a_i, b_i \rangle - \frac{1}{p^2} \sum_{i>j} \langle a_i - a_j, b_i - b_j \rangle$ , we have

$$\left\langle \frac{1}{p} \sum_{i=1}^p \nabla f_t^i, y_t - x^* \right\rangle = \frac{1}{p} \sum_{i=1}^p \langle \nabla f_t^i, x_t^i - x^* \rangle - \frac{1}{p^2} \sum_{i>j} \langle \nabla f_t^i - \nabla f_t^j, x_t^i - x_t^j \rangle. \quad (3.39)$$

By the cosine rule, we have

$$2 \langle y_t - \tilde{x}_t, y_t - x^* \rangle = \|y_t - x^*\|^2 + \|y_t - \tilde{x}_t\|^2 - \|\tilde{x}_t - x^*\|^2. \quad (3.40)$$



Denote  $\xi_t = \frac{1}{p} \sum_{i=1}^p \xi_t^i$ , we can rewrite Equation 3.38 as

$$\begin{aligned}
\|y_{t+1} - x^*\|^2 &= \|y_t - x^*\|^2 + \eta^2 \left\| \frac{1}{p} \sum_{i=1}^p \nabla f_t^i - \xi_t \right\|^2 + \alpha^2 \|y_t - \tilde{x}_t\|^2 \\
&\quad - 2\eta \left\langle \frac{1}{p} \sum_{i=1}^p \nabla f_t^i - \xi_t, y_t - x^* \right\rangle - 2\alpha \langle y_t - \tilde{x}_t, y_t - x^* \rangle \\
&\quad + 2\eta\alpha \left\langle \frac{1}{p} \sum_{i=1}^p \nabla f_t^i - \xi_t, y_t - \tilde{x}_t \right\rangle. \tag{3.41}
\end{aligned}$$

By combining the above Equations 3.39, 3.40 with 3.41, we obtain

$$\begin{aligned}
\|y_{t+1} - x^*\|^2 &= \|y_t - x^*\|^2 + \eta^2 \left\| \frac{1}{p} \sum_{i=1}^p \nabla f_t^i - \xi_t \right\|^2 + \alpha^2 \|y_t - \tilde{x}_t\|^2 \\
&\quad - 2\eta \left( \frac{1}{p} \sum_{i=1}^p \langle \nabla f_t^i, x_t^i - x^* \rangle - \frac{1}{p^2} \sum_{i>j} \langle \nabla f_t^i - \nabla f_t^j, x_t^i - x_t^j \rangle \right) \tag{3.42}
\end{aligned}$$

$$\begin{aligned}
&\quad + 2\eta \langle \xi_t, y_t - x^* \rangle - \alpha (\|y_t - x^*\|^2 + \|y_t - \tilde{x}_t\|^2 - \|\tilde{x}_t - x^*\|^2) \\
&\quad + 2\eta\alpha \left\langle \frac{1}{p} \sum_{i=1}^p \nabla f_t^i - \xi_t, y_t - \tilde{x}_t \right\rangle. \tag{3.43}
\end{aligned}$$

Thus it follows from Equation 3.30 and 3.43 that

$$\begin{aligned}
\|y_{t+1} - x^*\|^2 &\leq \|y_t - x^*\|^2 + \eta^2 \left\| \frac{1}{p} \sum_{i=1}^p \nabla f_t^i - \xi_t \right\|^2 + \alpha^2 \|y_t - \tilde{x}_t\|^2 \\
&\quad - 2\eta \frac{1}{p} \sum_{i=1}^p \left( \frac{\mu L}{\mu + L} \|x_t^i - x^*\|^2 + \frac{1}{\mu + L} \|\nabla f_t^i\|^2 \right) \\
&\quad + 2\eta \frac{1}{p^2} \sum_{i>j} \langle \nabla f_t^i - \nabla f_t^j, x_t^i - x_t^j \rangle \\
&\quad + 2\eta \langle \xi_t, y_t - x^* \rangle - \alpha (\|y_t - x^*\|^2 + \|y_t - \tilde{x}_t\|^2 - \|\tilde{x}_t - x^*\|^2) \\
&\quad + 2\eta\alpha \left\langle \frac{1}{p} \sum_{i=1}^p \nabla f_t^i - \xi_t, y_t - \tilde{x}_t \right\rangle. \tag{3.44}
\end{aligned}$$

Recall  $y_t = \frac{1}{p} \sum_{i=1}^p x_t^i$ , we have the following bias-variance relation,

$$\begin{aligned} \frac{1}{p} \sum_{i=1}^p \|x_t^i - x^*\|^2 &= \frac{1}{p} \sum_{i=1}^p \|x_t^i - y_t\|^2 + \|y_t - x^*\|^2 = \frac{1}{p^2} \sum_{i>j} \|x_t^i - x_t^j\|^2 + \|y_t - x^*\|^2, \\ \frac{1}{p} \sum_{i=1}^p \|\nabla f_t^i\|^2 &= \frac{1}{p^2} \sum_{i>j} \|\nabla f_t^i - \nabla f_t^j\|^2 + \left\| \frac{1}{p} \sum_{i=1}^p \nabla f_t^i \right\|^2. \end{aligned} \quad (3.45)$$

By the Cauchy-Schwarz inequality, we have

$$\frac{\mu L}{\mu + L} \|x_t^i - x_t^j\|^2 + \frac{1}{\mu + L} \|\nabla f_t^i - \nabla f_t^j\|^2 \geq \frac{2\sqrt{\mu L}}{\mu + L} \langle \nabla f_t^i - \nabla f_t^j, x_t^i - x_t^j \rangle. \quad (3.46)$$

Combining the above estimates in Equations 3.44, 3.45, 3.46, we obtain

$$\begin{aligned} \|y_{t+1} - x^*\|^2 &\leq \|y_t - x^*\|^2 + \eta^2 \left\| \frac{1}{p} \sum_{i=1}^p \nabla f_t^i - \xi_t \right\|^2 + \alpha^2 \|y_t - \tilde{x}_t\|^2 \\ &\quad - 2\eta \left( \frac{\mu L}{\mu + L} \|y_t - x^*\|^2 + \frac{1}{\mu + L} \left\| \frac{1}{p} \sum_{i=1}^p \nabla f_t^i \right\|^2 \right) \\ &\quad + 2\eta \left( 1 - \frac{2\sqrt{\mu L}}{\mu + L} \right) \frac{1}{p^2} \sum_{i>j} \langle \nabla f_t^i - \nabla f_t^j, x_t^i - x_t^j \rangle \\ &\quad + 2\eta \langle \xi_t, y_t - x^* \rangle - \alpha (\|y_t - x^*\|^2 + \|y_t - \tilde{x}_t\|^2 - \|\tilde{x}_t - x^*\|^2) \\ &\quad + 2\eta\alpha \left\langle \frac{1}{p} \sum_{i=1}^p \nabla f_t^i - \xi_t, y_t - \tilde{x}_t \right\rangle. \end{aligned} \quad (3.47)$$

Similarly if  $0 \leq \alpha < 1$ , we can have this upper-bound for the terms  $\alpha^2 \|y_t - \tilde{x}_t\|^2 - \alpha \|y_t - \tilde{x}_t\|^2 + 2\eta\alpha \left\| \frac{1}{p} \sum_{i=1}^p \nabla f_t^i \right\| \|y_t - \tilde{x}_t\| \leq \frac{\eta^2\alpha}{1-\alpha} \left\| \frac{1}{p} \sum_{i=1}^p \nabla f_t^i \right\|^2$  by applying  $-ax^2 +$

$bx \leq \frac{b^2}{4a}$  with  $x = \|y_t - \tilde{x}_t\|$ . Thus we have the following bound for the Equation 3.47

$$\begin{aligned}
\|y_{t+1} - x^*\|^2 &\leq (1 - 2\eta \frac{\mu L}{\mu + L} - \alpha) \|y_t - x^*\|^2 + (\eta^2 + \frac{\eta^2 \alpha}{1 - \alpha} - \frac{2\eta}{\mu + L}) \left\| \frac{1}{p} \sum_{i=1}^p \nabla f_t^i \right\|^2 \\
&\quad - 2\eta^2 \left\langle \frac{1}{p} \sum_{i=1}^p \nabla f_t^i, \xi_t \right\rangle + 2\eta \langle \xi_t, y_t - x^* \rangle - 2\eta \alpha \langle \xi_t, y_t - \tilde{x}_t \rangle \\
&\quad + 2\eta \left( 1 - \frac{2\sqrt{\mu L}}{\mu + L} \right) \frac{1}{p^2} \sum_{i>j} \langle \nabla f_t^i - \nabla f_t^j, x_t^i - x_t^j \rangle \\
&\quad + \eta^2 \|\xi_t\|^2 + \alpha \|\tilde{x}_t - x^*\|^2.
\end{aligned} \tag{3.48}$$

Since  $\frac{2\sqrt{\mu L}}{\mu + L} \leq 1$ , we need also bound the nonlinear term  $\langle \nabla f_t^i - \nabla f_t^j, x_t^i - x_t^j \rangle \leq L \|x_t^i - x_t^j\|^2$ . Recall the bias-variance relation  $\frac{1}{p} \sum_{i=1}^p \|x_t^i - x^*\|^2 = \frac{1}{p^2} \sum_{i>j} \|x_t^i - x_t^j\|^2 + \|y_t - x^*\|^2$ . The key observation is that if  $\frac{1}{p} \sum_{i=1}^p \|x_t^i - x^*\|^2$  remains bounded, then larger variance  $\sum_{i>j} \|x_t^i - x_t^j\|^2$  implies smaller bias  $\|y_t - x^*\|^2$ . Thus this nonlinear term can be compensated.

Again choose  $\eta$  small enough such that  $\eta^2 + \frac{\eta^2 \alpha}{1 - \alpha} - \frac{2\eta}{\mu + L} \leq 0$  and take expectation in Equation 3.48,

$$\begin{aligned}
\mathbb{E} \|y_{t+1} - x^*\|^2 &\leq (1 - 2\eta \frac{\mu L}{\mu + L} - \alpha) \mathbb{E} \|y_t - x^*\|^2 \\
&\quad + 2\eta L \left( 1 - \frac{2\sqrt{\mu L}}{\mu + L} \right) \left( \frac{1}{p} \sum_{i=1}^p \mathbb{E} \|x_t^i - x^*\|^2 - \mathbb{E} \|y_t - x^*\|^2 \right) \\
&\quad + \eta^2 \frac{\sigma^2}{p} + \alpha \mathbb{E} \|\tilde{x}_t - x^*\|^2.
\end{aligned} \tag{3.49}$$

As for the center variable in Equation 3.28, we apply simply the convexity of the norm  $\|\cdot\|^2$  to obtain

$$\|\tilde{x}_{t+1} - x^*\|^2 \leq (1 - \beta) \|\tilde{x}_t - x^*\|^2 + \beta \|y_t - x^*\|^2. \tag{3.50}$$

Combining the estimates from Equations 3.37, 3.49, 3.50, and denote  $a_t = \mathbb{E} \|y_t - x^*\|^2$ ,

$b_t = \frac{1}{p} \sum_{i=1}^p \mathbb{E} \|x_t^i - x^*\|^2$ ,  $c_t = \mathbb{E} \|\tilde{x}_t - x^*\|^2$ ,  $\gamma_1 = 2\eta \frac{\mu L}{\mu+L}$ ,  $\gamma_2 = 2\eta L(1 - \frac{2\sqrt{\mu L}}{\mu+L})$ , then

$$\begin{pmatrix} a_{t+1} \\ b_{t+1} \\ c_{t+1} \end{pmatrix} \leq \begin{pmatrix} 1 - \gamma_1 - \gamma_2 - \alpha & \gamma_2 & \alpha \\ 0 & 1 - \gamma_1 - \alpha & \alpha \\ \beta & 0 & 1 - \beta \end{pmatrix} \begin{pmatrix} a_t \\ b_t \\ c_t \end{pmatrix} + \begin{pmatrix} \frac{\eta^2 \sigma^2}{p} \\ \eta^2 \sigma^2 \\ 0 \end{pmatrix},$$

as long as  $0 \leq \beta \leq 1$ ,  $0 \leq \alpha < 1$  and  $\eta^2 + \frac{\eta^2 \alpha}{1-\alpha} - \frac{2\eta}{\mu+L} \leq 0$ , i.e.  $0 \leq \eta \leq \frac{2}{\mu+L}(1-\alpha)$ .  $\square$

The above theorem captures the bias-variance tradeoff of the spatial average  $\frac{1}{p} \sum_{i=1}^p x_t^i$  of the local variables (the  $a_t$ ), with respect to the averaged mean squared error of each local variable (the  $b_t$ ). The center variable  $\tilde{x}_t$  is tracking  $\frac{1}{p} \sum_{i=1}^p x_t^i$  over time (the  $c_t$ ).

To get an upper bound on the rate of convergence for  $\tilde{x}_t$ , we need to assume the matrix  $M$  to be positive, and its spectral norm to be smaller than one. Here

$$M = \begin{pmatrix} 1 - \gamma_1 - \gamma_2 - \alpha & \gamma_2 & \alpha \\ 0 & 1 - \gamma_1 - \alpha & \alpha \\ \beta & 0 & 1 - \beta \end{pmatrix}.$$

We have three eigenvalues of  $M$  as follows:

$$\begin{aligned} \lambda_1 &= 1 - \alpha - \gamma_1 - \gamma_2, \\ \lambda_2 &= 1 + \frac{1}{2}(-\alpha - \beta - \gamma_1 + \sqrt{(\alpha + \beta + \gamma_1)^2 - 4\beta\gamma_1}), \\ \lambda_3 &= 1 + \frac{1}{2}(-\alpha - \beta - \gamma_1 - \sqrt{(\alpha + \beta + \gamma_1)^2 - 4\beta\gamma_1}). \end{aligned}$$

Under the conditions of the above theorem ( $0 \leq \eta \leq \frac{2}{\mu+L}(1-\alpha)$ ,  $0 \leq \alpha < 1$  and  $0 \leq \beta \leq 1$ ), we still need to assume  $\lambda_1 \geq 0$  so that  $M$  is positive. Since  $\gamma_1 = 2\eta \frac{\mu L}{\mu+L} \geq 0$  and  $\gamma_2 = 2\eta L(1 - \frac{2\sqrt{\mu L}}{\mu+L}) \geq 0$ , we deduce that  $\lambda_1 \leq 1$ . We can also verify that  $\lambda_3 \leq \lambda_2 \leq 1$ . Thus for the stability we only need  $\lambda_3 \geq -1$ .

For  $\lambda_1 > 0$ , we get the condition  $0 < \eta < \frac{1-\alpha}{\frac{2\mu L}{\mu+L} + 2L(1 - \frac{2\sqrt{\mu L}}{\mu+L})}$ . For  $\lambda_3 > -1$ , we have the

condition  $\eta < \frac{\mu+L}{\mu L}(1 - \frac{\alpha}{2-\beta})$ . When  $\mu = L$ , these two conditions mean  $0 < \eta < \frac{1-\alpha}{L}$  and  $0 < \eta < \frac{2}{L}(1 - \frac{\alpha}{2-\beta})$ . On the other hand, when  $\mu = 0$ , we have  $0 < \eta < \frac{1-\alpha}{2L}$ . In either case, our method operates in the under-damping (no oscillations) region.

With the above conditions, we can now ask what is the asymptotic variance of  $a_t$ ,  $b_t$  and  $c_t$ . By solving the fixed point equation  $(a_\infty, b_\infty, c_\infty)' = M(a_\infty, b_\infty, c_\infty)' + (\eta^2 \frac{\sigma^2}{p}, \eta^2 \sigma^2, 0)'$ , we obtain

$$\begin{aligned} a_\infty = c_\infty &= \frac{\alpha/p + \gamma_1/p + \gamma_2}{\gamma_1(\alpha + \gamma_1 + \gamma_2)} \eta^2 \sigma^2, \\ b_\infty &= \frac{\alpha/p + \gamma_1 + \gamma_2}{\gamma_1(\alpha + \gamma_1 + \gamma_2)} \eta^2 \sigma^2. \end{aligned}$$

If  $\mu = L$ , then  $\gamma_2 = 0$ , we indeed get the asymptotic variance  $c_\infty$  of order  $\sigma^2/p$ . This order matches our quadratic case analysis above. However, if  $\mu \ll L$ , then  $\gamma_1$  will be close to zero, and we don't see in this upper bound the benefit of variance reduction by increasing  $p$  (number of workers). It would be interesting to find a non-quadratic example such that this can actually happen.

### 3.3 Stability of *EASGD* and *ADMM*

In this section we study the stability of the asynchronous *EASGD* and *ADMM* methods in the round-robin scheme [29]. We first state the updates of both algorithms in this setting, and then we study their stability. We will show that in the one-dimensional quadratic case, *ADMM* algorithm can exhibit chaotic behavior, leading to exponential divergence. The analytic condition for the *ADMM* algorithm to be stable is still unknown, while for the *EASGD* algorithm it is very simple.

In our setting, the *ADMM* method [11, 61, 42] involves solving the following minimax

problem<sup>3</sup>,

$$\max_{\lambda^1, \dots, \lambda^p} \min_{x^1, \dots, x^p, \tilde{x}} \sum_{i=1}^p F(x^i) - \lambda^i(x^i - \tilde{x}) + \frac{\rho}{2} \|x^i - \tilde{x}\|^2, \quad (3.51)$$

where  $\lambda^i$ 's are the Lagrangian multipliers. The resulting updates of the *ADMM* algorithm in the round-robin scheme are given next. Let  $t \geq 0$  be a global clock. At each  $t$ , we linearize the function  $F(x^i)$  with  $F(x_t^i) + \langle \nabla F(x_t^i), x^i - x_t^i \rangle + \frac{1}{2\eta} \|x^i - x_t^i\|^2$  as in [42].

The updates become

$$\lambda_{t+1}^i = \begin{cases} \lambda_t^i - (x_t^i - \tilde{x}_t) & \text{if } \text{mod}(t, p) = i - 1; \\ \lambda_t^i & \text{if } \text{mod}(t, p) \neq i - 1. \end{cases} \quad (3.52)$$

$$x_{t+1}^i = \begin{cases} \frac{x_t^i - \eta \nabla F(x_t^i) + \eta \rho (\lambda_{t+1}^i + \tilde{x}_t)}{1 + \eta \rho} & \text{if } \text{mod}(t, p) = i - 1; \\ x_t^i & \text{if } \text{mod}(t, p) \neq i - 1. \end{cases} \quad (3.53)$$

$$\tilde{x}_{t+1} = \frac{1}{p} \sum_{i=1}^p (x_{t+1}^i - \lambda_{t+1}^i). \quad (3.54)$$

Each local variable  $x^i$  is periodically updated (with period  $p$ ). First, the Lagrangian multiplier  $\lambda^i$  is updated with the dual ascent update as in Equation 3.52. It is followed by the gradient descent update of the local variable as given in Equation 3.53. Then the center variable  $\tilde{x}$  is updated with the most recent values of all the local variables and Lagrangian multipliers as in Equation 3.54. Note that since the step size for the dual ascent update is chosen to be  $\rho$  by convention [11, 61, 42], we have re-parametrized the Lagrangian multiplier to be  $\lambda_t^i \leftarrow \lambda_t^i / \rho$  in the above updates.

The *EASGD* algorithm in the round-robin scheme is defined similarly and is given below

$$x_{t+1}^i = \begin{cases} x_t^i - \eta \nabla F(x_t^i) - \alpha(x_t^i - \tilde{x}_t) & \text{if } \text{mod}(t, p) = i - 1; \\ x_t^i & \text{if } \text{mod}(t, p) \neq i - 1. \end{cases} \quad (3.55)$$

$$\tilde{x}_{t+1} = \tilde{x}_t + \sum_{i: \text{mod}(t, p) = i - 1} \alpha(x_t^i - \tilde{x}_t). \quad (3.56)$$

---

<sup>3</sup>The convergence analysis in [61] is based on the assumption that ‘‘At any master iteration, updates from the workers have the same probability of arriving at the master.’’, which is not satisfied in the round-robin scheme.

At time  $t$ , only the  $i$ -th local worker (whose index  $i - 1$  equals  $t$  modulo  $p$ ) is activated, and performs the update in Equations 3.55 which is followed by the master update given in Equation 3.56.

We will now focus on the one-dimensional quadratic case without noise, i.e.

$$F(x) = \frac{x^2}{2}, \quad x \in \mathbb{R}.$$

For the *ADMM* algorithm, let the state of the (dynamical) system at time  $t$  be  $s_t = (\lambda_t^1, x_t^1, \dots, \lambda_t^p, x_t^p, \tilde{x}_t) \in \mathbb{R}^{2p+1}$ . The local worker  $i$ 's updates in Equations 3.52, 3.53, and 3.54 are composed of three linear maps which can be written as  $s_{t+1} = (F_3^i \circ F_2^i \circ F_1^i)(s_t)$ . For simplicity, we will only write them out below for the case when  $i = 1$  and  $p = 2$ :

$$F_1^1 = \begin{pmatrix} 1 & -1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}, F_2^1 = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ \frac{\eta\rho}{1+\eta\rho} & \frac{1-\eta}{1+\eta\rho} & 0 & 0 & \frac{\eta\rho}{1+\eta\rho} \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}, F_3^1 = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ -\frac{1}{p} & \frac{1}{p} & -\frac{1}{p} & \frac{1}{p} & 0 \end{pmatrix}.$$

For each of the  $p$  linear maps, it's possible to find a simple condition such that each map, where the  $i^{\text{th}}$  map has the form  $F_3^i \circ F_2^i \circ F_1^i$ , is stable (the absolute value of the eigenvalues of the map are smaller or equal to one). However, when these non-symmetric maps are composed one after another as follows  $\mathcal{F} = F_3^p \circ F_2^p \circ F_1^p \circ \dots \circ F_3^1 \circ F_2^1 \circ F_1^1$ , the resulting map  $\mathcal{F}$  can become unstable! (more precisely, some eigenvalues of the map can sit outside the unit circle in the complex plane).

We now present the numerical conditions for which the *ADMM* algorithm becomes unstable in the round-robin scheme for  $p = 3$  and  $p = 8$ , by computing the largest absolute eigenvalue of the map  $\mathcal{F}$ . Figure 3.2 summarizes the obtained result. We also illustrate this unstable behavior in Figure 3.3

On the other hand, the *EASGD* algorithm involves composing only symmetric linear maps due to the elasticity. Let the state of the (dynamical) system at time  $t$  be  $s_t = (x_t^1, \dots, x_t^p, \tilde{x}_t) \in \mathbb{R}^{p+1}$ . The activated local worker  $i$ 's update in Equation 3.55 and the master update in Equation 3.56 can be written as  $s_{t+1} = F^i(s_t)$ . In case of  $p = 2$ , the map  $F^1$  and  $F^2$  are defined as follows

$$F^1 = \begin{pmatrix} 1 - \eta - \alpha & 0 & \alpha \\ 0 & 1 & 0 \\ \alpha & 0 & 1 - \alpha \end{pmatrix}, \quad F^2 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 - \eta - \alpha & \alpha \\ 0 & \alpha & 1 - \alpha \end{pmatrix}$$

For the composite map  $F^p \circ \dots \circ F^1$  to be stable, the condition that needs to be satisfied is actually the same for each  $i$ , and is furthermore independent of  $p$  (since each linear map  $F^i$  is symmetric). It essentially involves the stability of the  $2 \times 2$  matrix

$$\begin{pmatrix} 1 - \eta - \alpha & \alpha \\ \alpha & 1 - \alpha \end{pmatrix},$$

whose two (real) eigenvalues  $\lambda$  satisfy  $(1 - \eta - \alpha - \lambda)(1 - \alpha - \lambda) = \alpha^2$ . The resulting stability condition ( $|\lambda| \leq 1$ ) is simple and given as

$$0 \leq \eta \leq 2, \quad 0 \leq \alpha \leq \frac{4 - 2\eta}{4 - \eta}.$$






Figure 3.2: The largest absolute eigenvalue of the linear map  $\mathcal{F} = F_3^p \circ F_2^p \circ F_1^p \circ \dots \circ F_3^1 \circ F_2^1 \circ F_1^1$  as a function of  $\eta \in (0, 10^{-2})$  and  $\rho \in (0, 10)$  when  $p = 3$  and  $p = 8$ . To simulate the chaotic behavior of the *ADMM* algorithm, one may pick  $\eta = 0.001$  and  $\rho = 2.5$  and initialize the state  $s_0$  either randomly or with  $\lambda_0^i = 0, x_0^i = \tilde{x}_0 = 1000, \forall i = 1, \dots, p$ . *Figure should be read in color.*

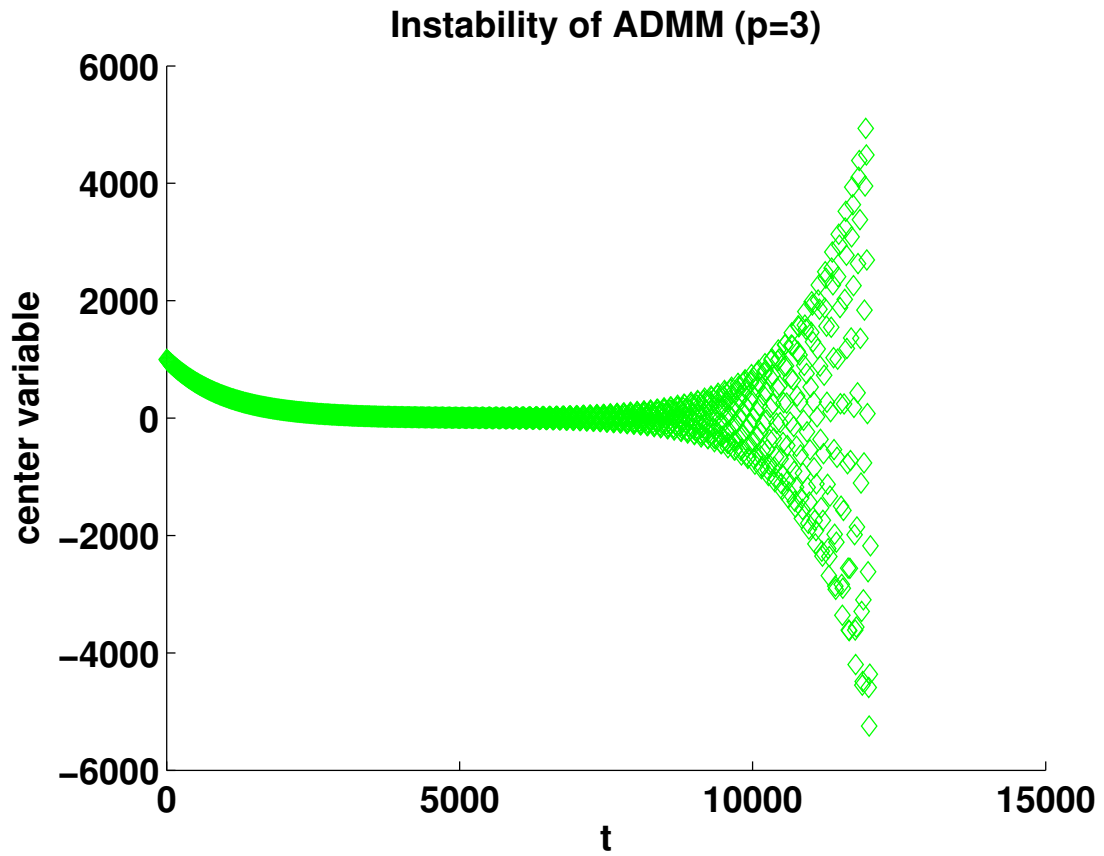


Figure 3.3: Instability of *ADMM* in the round-robin scheme. Pick  $p = 3$ ,  $\eta = 0.001$ ,  $\rho = 2.5$  and initialize the state  $s_0$  with  $\lambda_0^i = 0$ ,  $x_0^i = \tilde{x}_0 = 1000$ ,  $\forall i = 1, \dots, p$ . The x-axis is the time step  $t$ , the y-axis is the (one-dimensional) value of the center variable  $\tilde{x}_t$ .

## Chapter 4

# Performance in Deep Learning

In this chapter, we compare empirically the performance in deep learning of asynchronous *EASGD* and *EAMSGD* with the parallel method *DOWNPOUR* and the sequential method *SGD*, as well as their averaging and momentum variants.

All the parallel comparator methods are listed below<sup>1</sup>:

- *DOWNPOUR* [16], the detail and the pseudo-code of the implementation are described in Section 4.4 (Algorithm 3).
- *Momentum DOWNPOUR* (*MDOWNPOUR*), where the Nesterov’s momentum scheme is applied to the master’s update (note it is unclear how to apply it to the local workers or for the case when  $\tau > 1$ ). The pseudo-code is described in Section 4.4 (Algorithms 4 and 5).
- A method that we call *ADOWNPOUR*, where we compute the average over time of the center variable  $\tilde{x}$  as follows:  $z_{t+1} = (1 - \alpha_{t+1})z_t + \alpha_{t+1}\tilde{x}_t$ , and  $\alpha_{t+1} = \frac{1}{t+1}$  is a moving rate, and  $z_0 = \tilde{x}_0$ . The  $t$  denotes the master clock, which is initialized to 0 and incremented every time the center variable  $\tilde{x}$  is updated.
- A method that we call *MVADOWNPOUR*, where we compute the moving average

---

<sup>1</sup>We have compared asynchronous *ADMM* [61] with *EASGD* in our setting as well, the performance is nearly the same. However, *ADMM*’s momentum variant is not as stable when using large communication period  $\tau$ .

of the center variable  $\tilde{x}$  as follows:  $z_{t+1} = (1 - \alpha)z_t + \alpha\tilde{x}_t$ , and the moving rate  $\alpha$  was chosen to be constant, and  $z_0 = \tilde{x}_0$ . The  $t$  denotes the master clock and is defined in the same way as for the *ADOWNPOUR* method.

All the sequential comparator methods ( $p = 1$ ) are listed below:

- *SGD* [9] with constant learning rate  $\eta$ .
- *Momentum SGD (MSGD)* [55] with constant (Nesterov’s) momentum rate  $\delta$ .
- *ASGD* [45] with moving rate  $\alpha_{t+1} = \frac{1}{t+1}$ .
- *MVASGD* [45] with moving rate  $\alpha$  set to a constant.

We perform experiments on two benchmark datasets: CIFAR-10 (we refer to it as *CIFAR*)<sup>2</sup> and ImageNet ILSVRC 2013 (we refer to it as *ImageNet*)<sup>3</sup>. We focus on the image classification task with deep convolutional neural networks. We first explain the experimental setup in Section 4.1 and then present the main experimental results in Section 4.2. We present further experimental results in Section 4.3 and discuss the effect of the averaging, the momentum, the learning rate, the communication period, the data and parameter communication tradeoff, and finally the speedup.

## 4.1 Experimental setup

For all our experiments we use a GPU-cluster interconnected with InfiniBand. Each node has 4 Titan GPU processors where each local worker corresponds to one GPU processor. The center variable of the master is stored and updated on the centralized parameter server [16]. Our implementation is available at <https://github.com/sixin-zh/mpiT>.

To describe the architecture of the convolutional neural network, we will first introduce a notation. Let  $(c, x, y)$  denotes the size of the input image to each layer, where  $c$  is the number of color channels and  $(x, y)$  denotes the horizontal and the vertical dimension of

<sup>2</sup>Downloaded from <http://www.cs.toronto.edu/~kriz/cifar.html>.

<sup>3</sup>Downloaded from <http://image-net.org/challenges/LSVRC/2013>.

the input. Let  $C$  denotes the fully-connected convolutional operator and let  $R$  denotes the rectified linear non-linearity (relu, c.f. [35]),  $P$  denotes the max pooling operator,  $L$  denotes the linear operator and  $D$  denotes the dropout operator with rate equal to 0.5 and  $S$  denotes the the softmax nonlinearity. We use the cross-entropy loss for the classification.

For the *ImageNet* experiment we use the similar approach to [49] with the following 11-layer convolutional neural network:  $(3, 221, 221) \xrightarrow[(7,7,2,2)]{C,R} (96, 108, 108) \xrightarrow[(3,3,3,3)]{P} (96, 36, 36) \xrightarrow[(5,5,1,1)]{C,R} (256, 32, 32) \xrightarrow[(2,2,2,2)]{P} (256, 16, 16) \xrightarrow[(3,3,1,1)]{C,R} (384, 14, 14) \xrightarrow[(2,2,1,1)]{C,R} (384, 13, 13) \xrightarrow[(2,2,1,1)]{C,R} (256, 12, 12) \xrightarrow[(2,2,2,2)]{P} (256, 6, 6) \xrightarrow[0.5]{L,R,D} (4096, 1, 1) \xrightarrow[0.5]{L,R,D} (4096, 1, 1) \xrightarrow{L,S} (1000, 1, 1)$ .

For the *CIFAR* experiment we use the similar approach to [58] with the following 7-layer convolutional neural network:  $(3, 28, 28) \xrightarrow[(5,5,1,1)]{C,R} (64, 24, 24) \xrightarrow[(2,2,2,2)]{P} (64, 12, 12) \xrightarrow[(5,5,1,1)]{C,R} (128, 8, 8) \xrightarrow[(2,2,2,2)]{P} (128, 4, 4) \xrightarrow[(3,3,1,1)]{C,R} (64, 2, 2) \xrightarrow[0.5]{L,R,D} (256, 1, 1) \xrightarrow{L,S} (10, 1, 1)$ .

Note that the numbers below the rightarrow of the C and P operator represent the kernel size (first horizontal and then vertical), the stride size (first horizontal and then vertical) and the padding size (if exists, first horizontal and then vertical) on each of the two sides of the image. The number below the rightarrow of the D operator emphasizes the dropout rate 0.5 [54].

In our experiments, all the methods we run use the same initial parameter chosen randomly, except that we set all the biases to zero for *CIFAR* case and to 0.1 for *ImageNet* case. This parameter is used to initialize the master and all the local workers<sup>4</sup>. We add  $l_2$ -regularization  $\frac{\lambda}{2} \|x\|^2$  to the loss function  $F(x)$ . For *ImageNet* we use  $\lambda = 10^{-5}$  and for *CIFAR* we use  $\lambda = 10^{-4}$ . We also compute the stochastic gradient using mini-batches of sample size 128.

---

<sup>4</sup>On the contrary, initializing the local workers and the master with different random seeds 'traps' the algorithm in the symmetry breaking phase.

## Data preprocessing

For the *ImageNet* experiment, we re-size each RGB image so that the smallest dimension is 256 pixels. We also re-scale each pixel value to the interval  $[0, 1]$ . We then extract random crops (and their horizontal flips) of size  $3 \times 221 \times 221$  pixels and present these to the network in mini-batches of size 128.

For the *CIFAR* experiment, we use the original RGB image of size  $3 \times 32 \times 32$ . As before, we re-scale each pixel value to the interval  $[0, 1]$ . We then extract random crops (and their horizontal flips) of size  $3 \times 28 \times 28$  pixels and present these to the network in mini-batches of size 128.

The training and test loss and the test error are only computed from the center patch ( $3 \times 28 \times 28$ ) for the *CIFAR* experiment and the center patch ( $3 \times 221 \times 221$ ) for the *ImageNet* experiment.

## Data prefetching (Sampling the dataset by the local workers in parallel)

We will now explain precisely how the dataset is sampled by each local worker as uniformly and efficiently as possible. The general parallel data loading scheme on a single machine is as follows: we use  $k$  CPUs, where  $k = 8$ , to load the data in parallel. Each data loader reads from the memory-mapped (mmap) file a chunk of  $c$  raw images (preprocessing was described in the previous subsection) and their labels (for *CIFAR*  $c = 512$  and for *ImageNet*  $c = 64$ ). For the *CIFAR*, the mmap file of each data loader contains the entire dataset whereas for *ImageNet*, each mmap file of each data loader contains different  $1/k$  fractions of the entire dataset. A chunk of data is always sent by one of the data loaders to the first worker who requests the data. The next worker requesting the data from the same data loader will get the next chunk. Each worker requests in total  $k$  data chunks from  $k$  different data loaders and then process them before asking for new data chunks. Notice that each data loader cycles<sup>5</sup> through the

---

<sup>5</sup>Its advantage is observed in [10].

data in the mmap file, sending consecutive chunks to the workers in order in which it receives requests from them. When the data loader reaches the end of the mmap file, it selects the address in memory uniformly at random from the interval  $[0, s]$ , where  $s = (\text{number of images in the mmap file modulo mini-batch size})$ , and uses this address to start cycling again through the data in the mmap file. After the local worker receives the  $k$  data chunks from the data loaders, it shuffles them and divides it into mini-batches of size 128.

## 4.2 Experimental results

For all experiments in this section we use *EASGD* with  $\beta = 0.9$  and  $\alpha = \beta/p$ , for all momentum-based methods we set the momentum term  $\delta = 0.99$  and finally for *MVADOWNPOUR* we set the moving rate to  $\alpha = 0.001$ . We start with the experiment on *CIFAR* dataset with  $p = 4$  local workers running on a single computing node.

For all the methods, we examined the communication periods from the following set  $\tau = \{1, 4, 16, 64\}$ . For each method we examined a wide range of learning rates. The learning rates explored in all experiments are summarized in Table 4.1, 4.2 and 4.3. The *CIFAR* experiment was run 3 times independently from the same random initialization and for each method we report its best performance measured by the smallest achievable test error.

From the results in Figure 4.1, 4.2, 4.3 and 4.4, we conclude that all *DOWNPOUR*-based methods achieve their best performance (test error) for small  $\tau$  ( $\tau \in \{1, 4\}$ ), and become highly unstable for  $\tau \in \{16, 64\}$ . While *EAMSGD* significantly outperforms comparator methods for all values of  $\tau$  by having faster convergence. It also finds better-quality solution measured by the test error and this advantage becomes more significant for  $\tau \in \{16, 64\}$ . Note that the tendency to achieve better test performance with larger  $\tau$  is also characteristic for the *EASGD* algorithm. We remark that if the stochastic gradient is sparse, *DOWNPOUR* empirically performs well with large communication period [20].

Table 4.1: Learning rates explored for each method shown in Figure 4.1, 4.2, 4.3 and 4.4 (*CIFAR* experiment).

	$\eta$
EASGD	{0.05, 0.01, 0.005}
EAMSGD	{0.01, 0.005, 0.001}
DOWNPOUR ADOWNPOUR MVADOWNPOUR	{0.005, 0.001, 0.0005}
MDOWNPOUR	{0.00005, 0.00001, 0.000005}
SGD, ASGD, MVASGD	{0.05, 0.01, 0.005}
MSGD	{0.001, 0.0005, 0.0001}

Table 4.2: Learning rates explored for each method shown in Figure 4.5, 4.6 and 4.7 (*CIFAR* experiment).

	$\eta$
EASGD	{0.05, 0.01, 0.005}
EAMSGD	{0.01, 0.005, 0.001}
DOWNPOUR	{0.005, 0.001, 0.0005}
MDOWNPOUR	{0.00005, 0.00001, 0.000005}
SGD, ASGD, MVASGD	{0.05, 0.01, 0.005}
MSGD	{0.001, 0.0005, 0.0001}

Table 4.3: Learning rates explored for each method shown in Figure 4.8 and 4.9 (*ImageNet* experiment).

	$\eta$
EASGD	0.1
EAMSGD	0.001
DOWNPOUR	for $p = 4$ : 0.02 for $p = 8$ : 0.01
SGD, ASGD, MVASGD	0.05
MSGD	0.0005



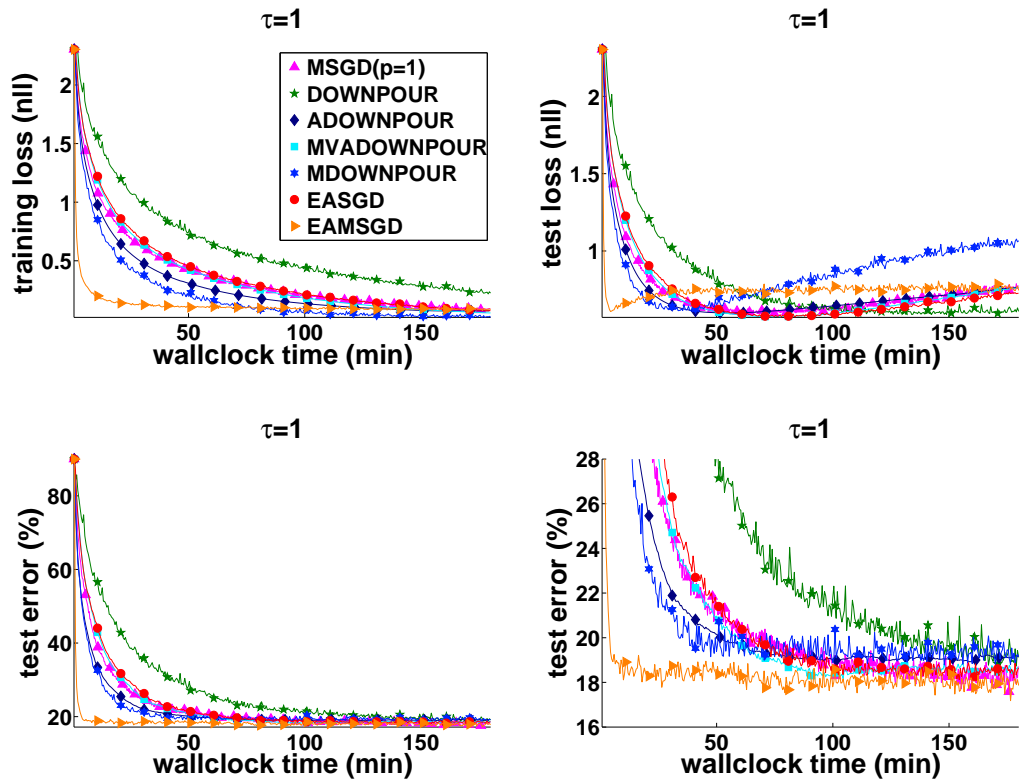


Figure 4.1: Training and test loss and the test error for the center variable versus a wallclock time for communication period  $\tau = 1$  on *CIFAR* dataset with the 7-layer convolutional neural network.  $p = 4$ .

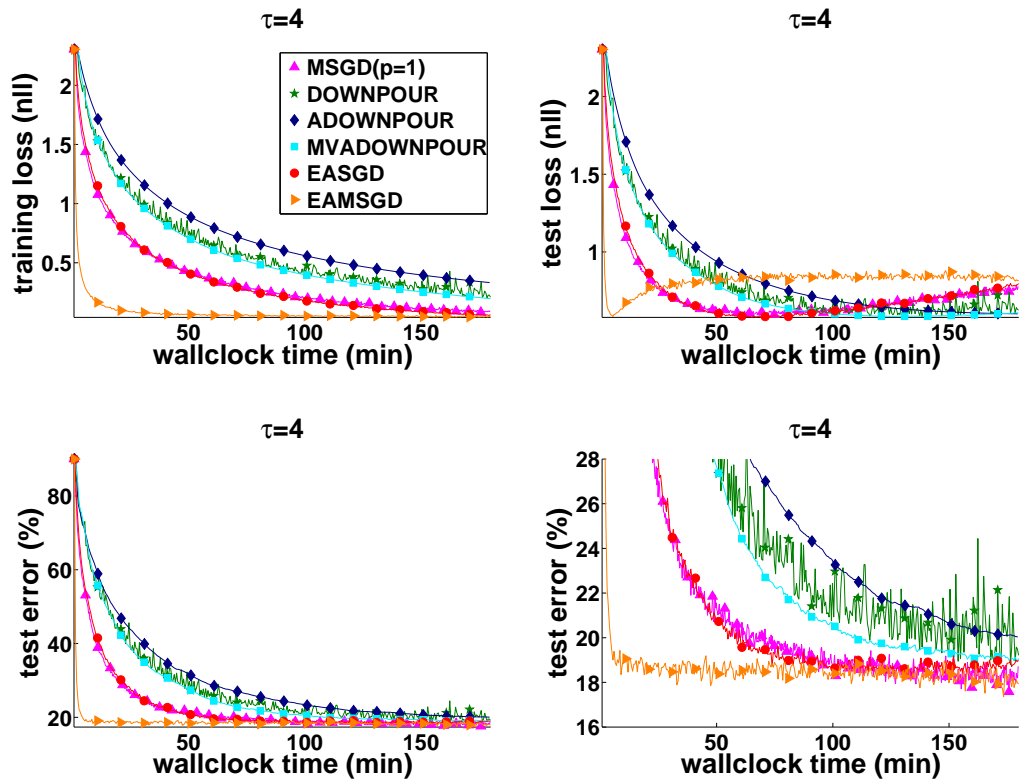


Figure 4.2: Training and test loss and the test error for the center variable versus a wallclock time for communication period  $\tau = 4$  on *CIFAR* dataset with the 7-layer convolutional neural network.  $p = 4$ .

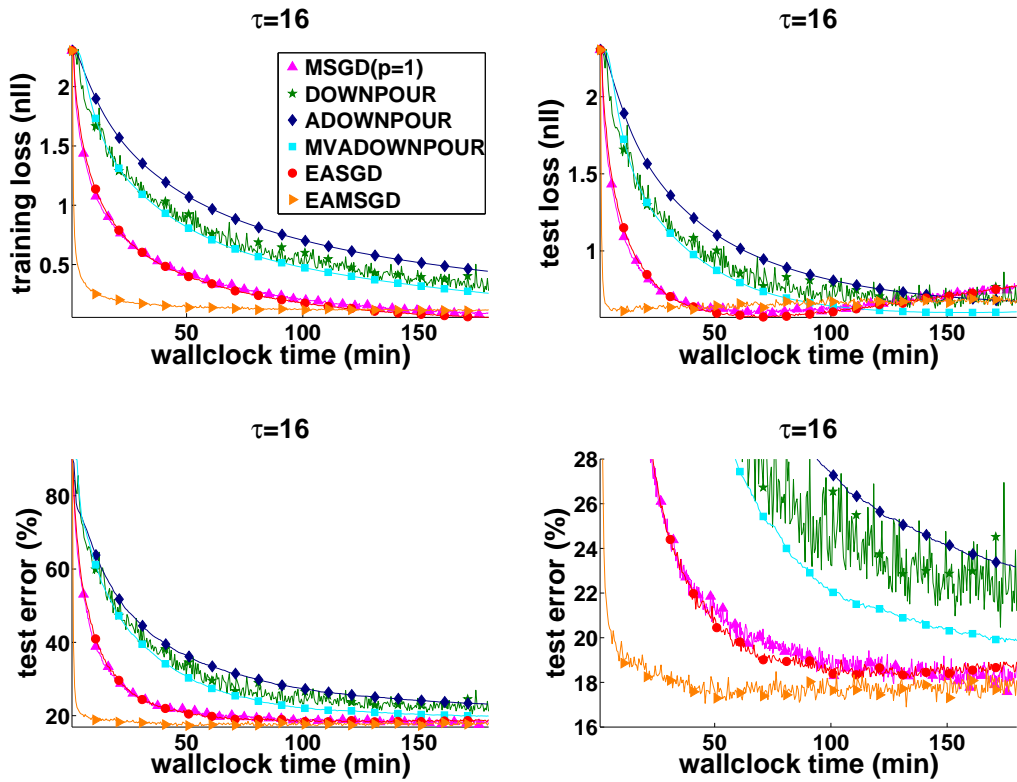


Figure 4.3: Training and test loss and the test error for the center variable versus a wallclock time for communication period  $\tau = 16$  on *CIFAR* dataset with the 7-layer convolutional neural network.  $p = 4$ .

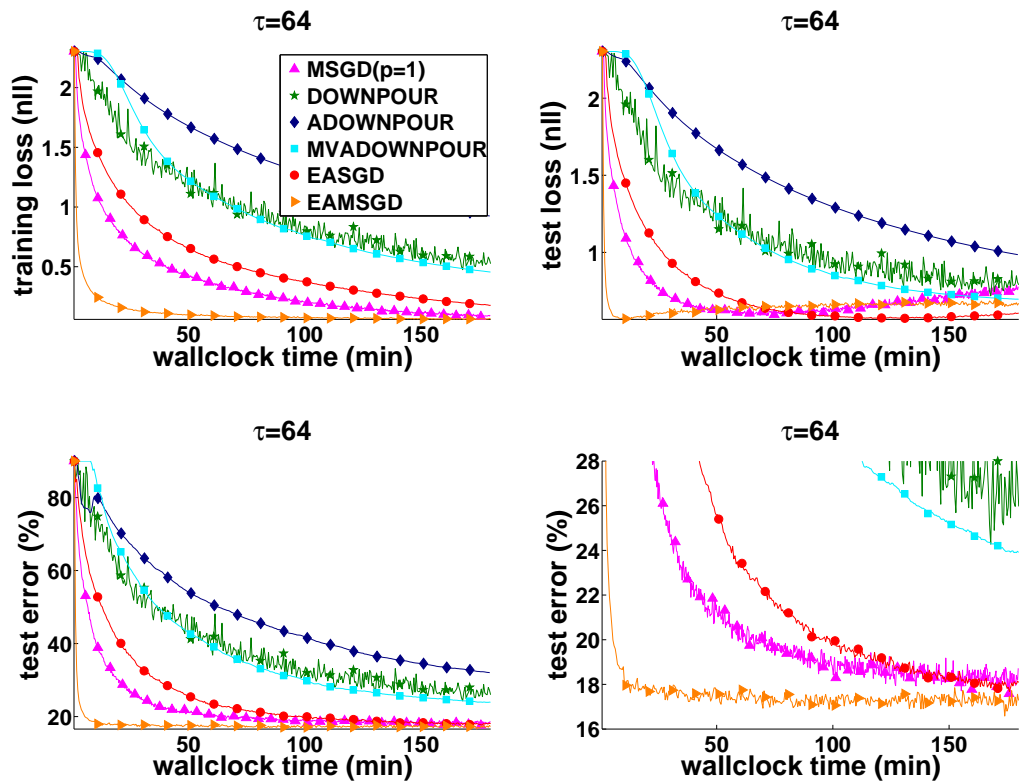


Figure 4.4: Training and test loss and the test error for the center variable versus a wallclock time for communication period  $\tau = 64$  on *CIFAR* dataset with the 7-layer convolutional neural network.  $p = 4$ .

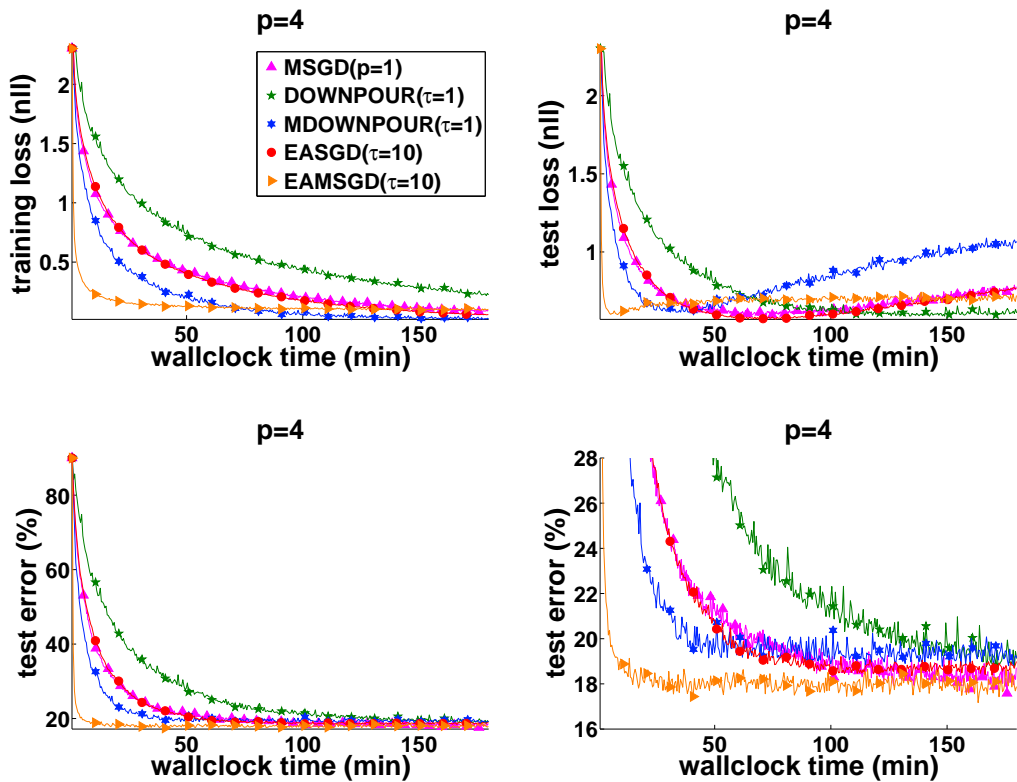


Figure 4.5: Training and test loss and the test error for the center variable versus a wallclock time with the number of local workers  $p = 4$  for parallel methods on *CIFAR* with the 7-layer convolutional neural network.

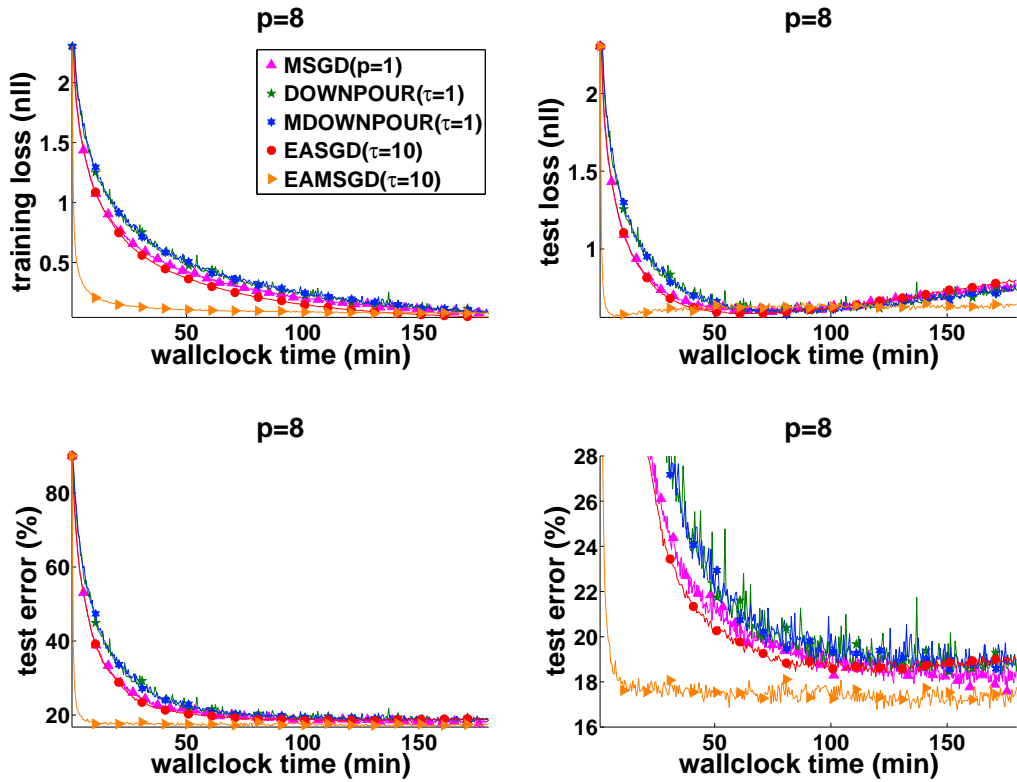


Figure 4.6: Training and test loss and the test error for the center variable versus a wallclock time with the number of local workers  $p = 8$  for parallel methods on *CIFAR* with the 7-layer convolutional neural network.

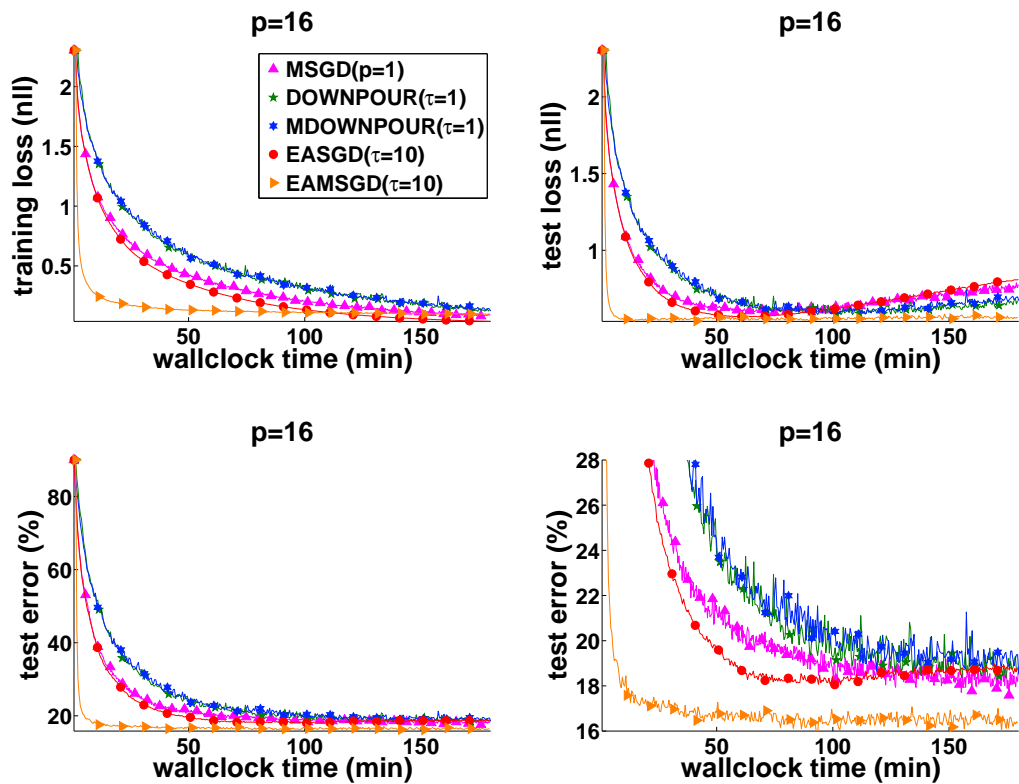


Figure 4.7: Training and test loss and the test error for the center variable versus a wallclock time with the number of local workers  $p = 16$  for parallel methods on *CIFAR* with the 7-layer convolutional neural network.

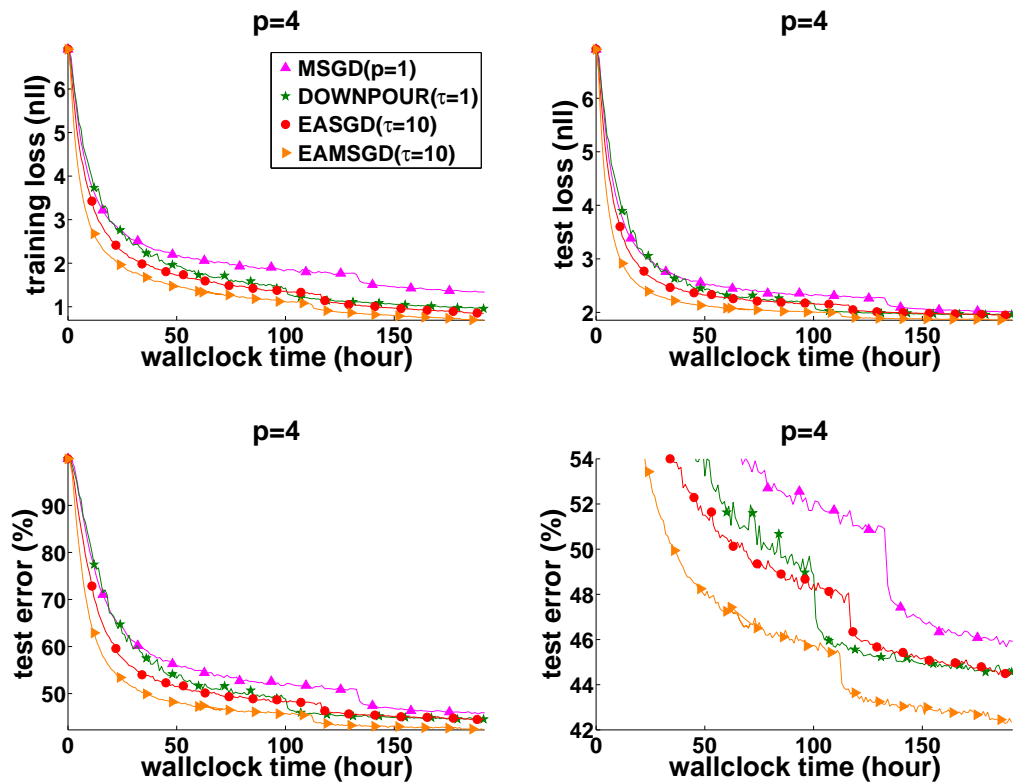


Figure 4.8: Training and test loss and the test error for the center variable versus a wallclock time with the number of local workers  $p = 4$  on *ImageNet* with the 11-layer convolutional neural network.



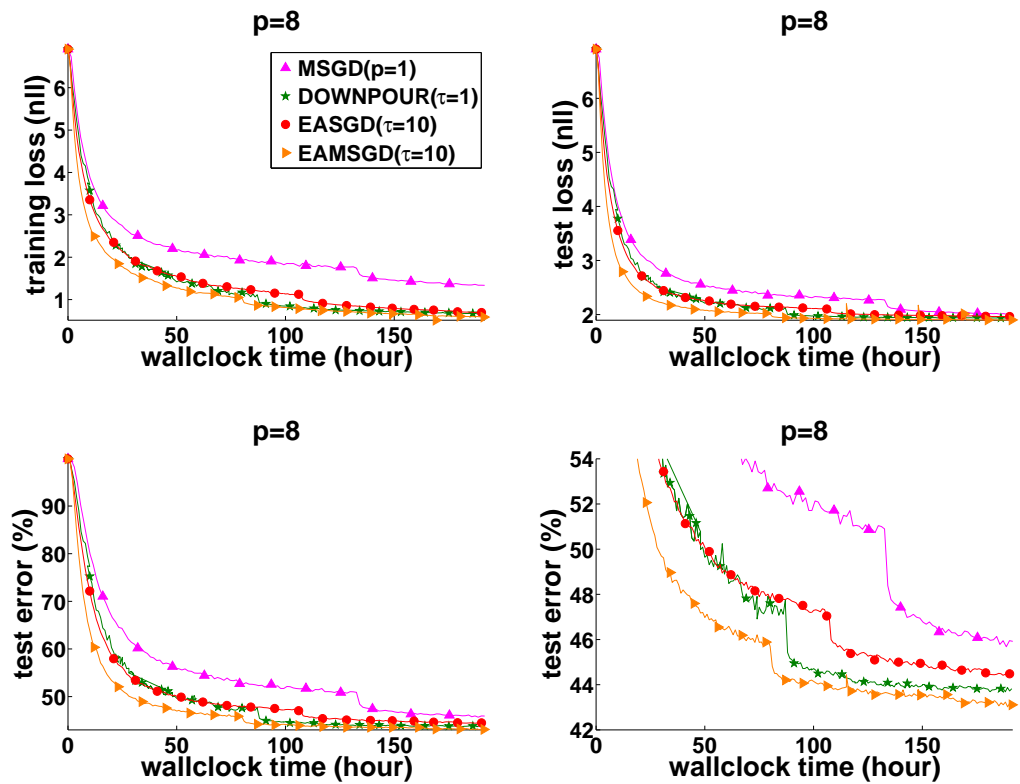


Figure 4.9: Training and test loss and the test error for the center variable versus a wallclock time with the number of local workers  $p = 8$  on *ImageNet* with the 11-layer convolutional neural network.

We next explore different number of local workers  $p$  from the set  $p = \{4, 8, 16\}$  for the *CIFAR* experiment, and  $p = \{4, 8\}$  for the *ImageNet* experiment<sup>6</sup>. For the *ImageNet* experiment we report the results of one run with the best setting we have found. *EASGD* and *EAMSGD* were run with  $\tau = 10$  whereas *DOWNPOUR* and *MDOWNPOUR* were run with  $\tau = 1$ .

For the *CIFAR* experiment, the results are in Figure 4.5, 4.6 and 4.7. *EAMSGD* achieves significant accelerations compared to other methods, e.g. the relative speedup for  $p = 16$  (the best comparator method is then *MSGD*) to achieve the test error 21% equals 11.1. It’s noticeable that the smallest achievable test error by either *EASGD* or *EAMSGD* decreases with larger  $p$ . This can potentially be explained by the fact that larger  $p$  allows for more exploration of the parameter space. In the next section, we discuss further the trade-off between exploration and exploitation as a function of the learning rate (section 4.3.2) and the communication period (section 4.3.3).

For the *ImageNet* experiment, the results are in Figure 4.8 and 4.9. The difficulty in this task is that we need to manually reduce the learning rate, otherwise the training loss will stagnate. Thus our initial learning rate is decreased twice over time, by a factor of 5 and then 2, when we observe that the online predictive loss [12] stagnates. *EAMSGD* again achieves significant accelerations compared to other methods, e.g. the relative speedup for  $p = 8$  (the best comparator method is then *DOWNPOUR*) to achieve the test error 49% equals 1.8, and simultaneously it reduces the communication overhead (*DOWNPOUR* uses communication period  $\tau = 1$  and *EAMSGD* uses  $\tau = 10$ ). However, there’s an annealing effect here in the sense that depending on the time the learning rate is reduced, the final test performance can be quite different. This makes the performance comparison difficult to define. In general, this is also a difficulty in comparing the NP-hard problem solvers.

---

<sup>6</sup>For the *ImageNet* experiment, the training loss is measured on a subset of the training data of size 50,000.

## 4.3 Further discussion and understanding

### 4.3.1 Comparison of *SGD*, *ASGD*, *MVASGD* and *MSGD*

For comparison we also report the performance of *MSGD* which outperformed *SGD*, *ASGD* and *MVASGD* on the test dataset. Recall that the way we compare the performance between different methods is based on the smallest achievable test error. Since the test dataset is fixed a priori, we may have the tendency to overfit this test dataset. Indeed, as we shall see in the Figure 4.10. One could use cross-validation to remedy this, we however emphasize that the point here is not to seek the best possible test accuracy, but to see all the possibilities that we can find, i.e. the richness of the dynamics arising from the neural network. We are aware that we could not exhaust all the possibilities.

Figure 4.10 shows the convergence of the training and test loss (negative log-likelihood) and the test error computed for the center variable as a function of wallclock time for *SGD*, *ASGD*, *MVASGD* and *MSGD* ( $p = 1$ ) on the *CIFAR* experiment. We observe that the final test performance of *ASGD* and *MSGD* are quite close to each other. But *ASGD* is much faster from the beginning. This explains why we do not see much speedup of the *EASGD* method (e.g. in Figure 4.1, 4.2, 4.3) and can sometimes be even slower (e.g. in Figure 4.4). It is caused by the sensitivity of the test performance to the choice of the learning rate. We shall discuss this phenomenon further in the next section 4.3.2.

Figure 4.11 shows the convergence of the training and test loss (negative log-likelihood) and the test error computed for the center variable as a function of wallclock time for *SGD*, *ASGD*, *MVASGD* and *MSGD* ( $p = 1$ ) on the *ImageNet* experiment. Note that for all *CIFAR* experiments we always start the averaging for the *ADOWNPOUR* and *ASGD* methods from the very beginning of each experiment. But for the *ImageNet* experiments we start the averaging for the *ASGD* and *MVASGD* at the first time when we reduce the learning rate. We have tried to start the averaging from the right beginning, both of the training and test performance are poor and they look very similar to the




Figure 4.10: Convergence of the training and test loss (negative log-likelihood) and the test error (original and zoomed) computed for the center variable as a function of wallclock time for *SGD*, *ASGD*, *MVASGD* and *MSGD* ( $p = 1$ ) on the *CIFAR* experiment.

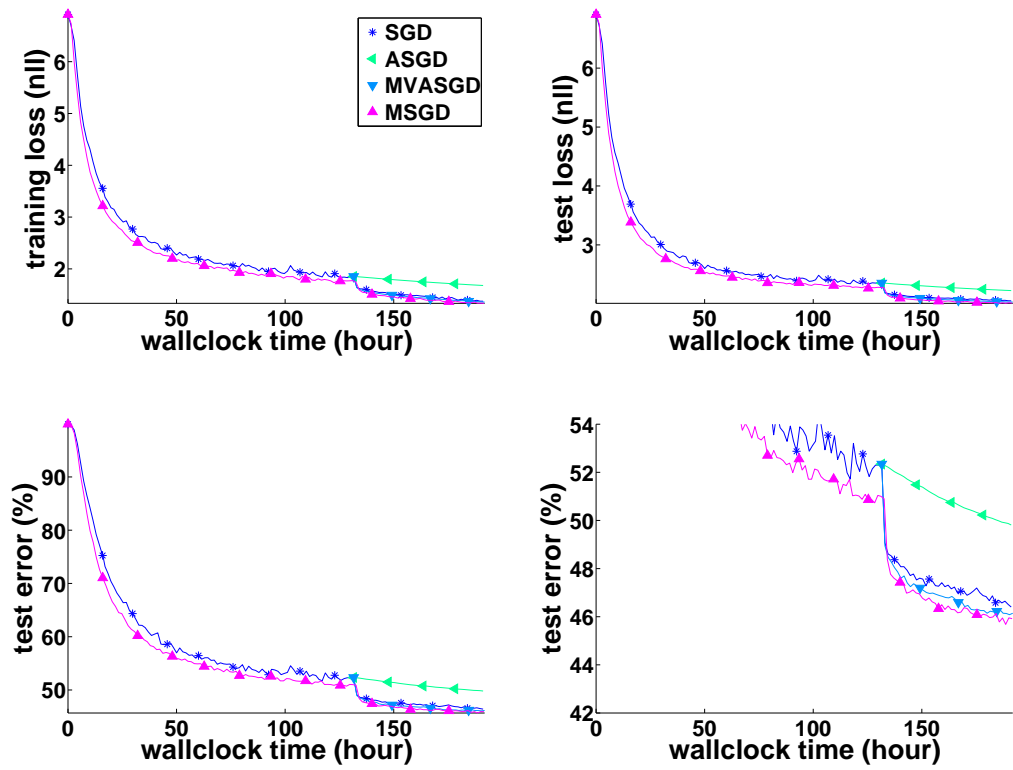


Figure 4.11: Convergence of the training and test loss (negative log-likelihood) and the test error (original and zoomed) computed for the center variable as a function of wallclock time for *SGD*, *ASGD*, *MVASGD* and *MSGD* ( $p = 1$ ) on the *ImageNet* experiment.

*ASGD* curve in Figure 4.11. The big difference between *ASGD* and *MVASGD* is quite striking and worth further study.

### 4.3.2 Dependence of the learning rate

This section discusses the dependence of the trade-off between exploration and exploitation on the learning rate. We compare the performance of respectively *EAMSGD* and *EASGD* for different learning rates  $\eta$  when  $p = 16$  and  $\tau = 10$  on the *CIFAR* experiment. We observe in Figure 4.12 that higher learning rates  $\eta$  lead to better test performance for the *EAMSGD* algorithm which potentially can be justified by the fact that they sustain higher fluctuations of the local workers. We conjecture that higher fluctuations lead to more exploration and simultaneously they also impose higher regularization. This picture however seems to be opposite for the *EASGD* algorithm for which larger learning rates hurt the performance of the method and lead to overfitting. Interestingly in this experiment for both *EASGD* and *EAMSGD* algorithm, the learning rate for which the best training performance was achieved simultaneously led to the worst test performance.

### 4.3.3 Dependence of the communication period

This section discusses the dependence of the trade-off between exploration and exploitation on the communication period. We observe in Figure 4.13 that *EASGD* algorithm exhibits very similar convergence behavior when  $\tau = 1$  up to even  $\tau = 1000$  for the *CIFAR* experiment, whereas *EAMSGD* can get trapped at a quite high energy level (of the objective) when  $\tau = 100$ . This trapping behavior is due to the non-convexity of the objective function. It can be avoided by gradually decreasing the learning rate, i.e. increasing the penalty term  $\rho$  (recall  $\alpha = \eta\rho$ ), as shown in Figure 4.13. In contrast, the *EASGD* algorithm does not seem to get trapped by any saddle point at all along its trajectory.

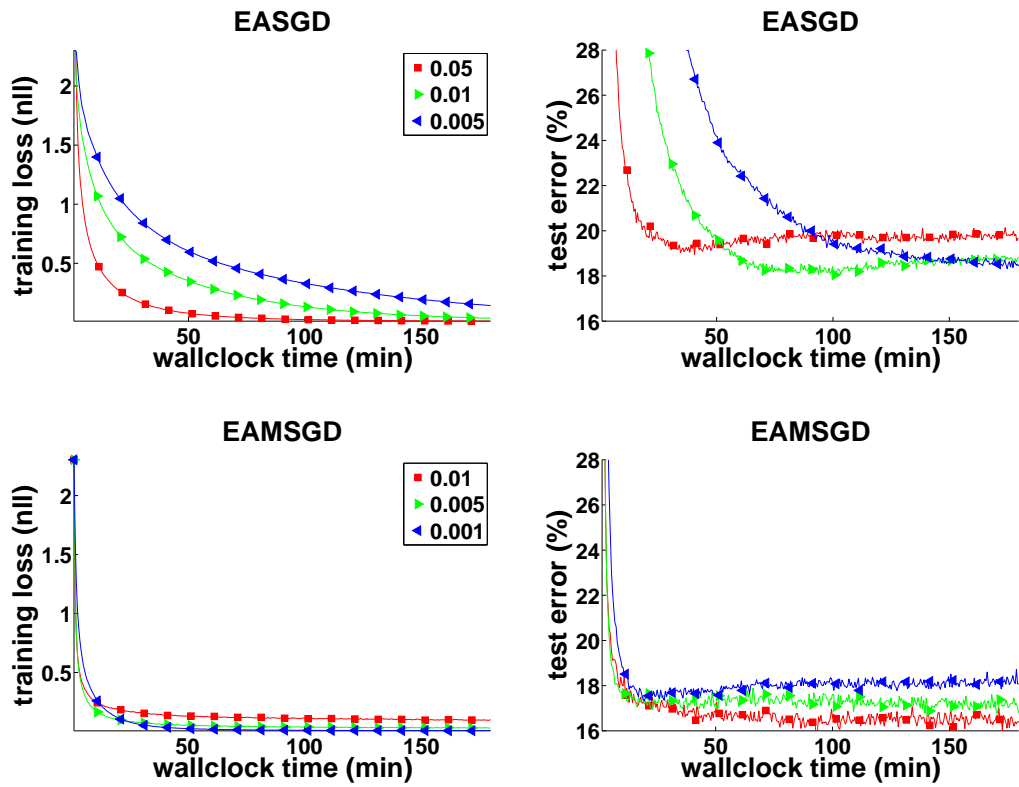


Figure 4.12: Convergence of the training loss (negative log-likelihood, original) and the test error (zoomed) computed for the center variable as a function of wallclock time for *EAMSGD* and *EASGD* run with different values of  $\eta$  on the *CIFAR* experiment.  $p = 16$ ,  $\tau = 10$ .

The performance<sup>7</sup> of *EASGD* being less sensitive to the communication period compared to *EAMSGD* is another striking observation.

It’s also very important to notice the tail behavior of the asynchronous *EASGD* method, i.e. what would happen if some local worker had finished the gradient updates and stopped the communication with the master. In the *EAMSGD* case in Figure 4.13, we see that the final training loss and test error can both become worse. This is due to the situation that some of the local workers have stopped earlier than the others, so that the averaging effect on the center variable is diminished.

#### 4.3.4 The tradeoff between data and parameter communication

In addition, we report in Table 4.4 the breakdown of the total running time for *EASGD* when  $\tau = 10$  (the time breakdown for *EAMSGD* is almost identical) and *DOWNPOUR* when  $\tau = 1$  into computation time, data loading time and parameter communication time. For the *CIFAR* experiment the reported time corresponds to processing  $400 \times 128$  data samples whereas for the *ImageNet* experiment it corresponds to processing  $1024 \times 128$  data samples. For  $\tau = 1$  and  $p \in \{8, 16\}$  we observe that the communication time accounts for significant portion of the total running time whereas for  $\tau = 10$  the communication time becomes negligible compared to the total running time (recall that based on previous results *EASGD* and *EAMSGD* achieve best performance with larger  $\tau$  which is ideal in the setting when communication is time-consuming).

We shall now examine the data communication cost in detail. Let’s focus on the *ImageNet* case. Based on the Table 4.4, for each single GPU ( $p = 1, \tau = 1$ ), it takes around 1248 seconds to process 1024 mini-batches of size 128. This is approximately processing one mini-batch per second. Each mini-batch consists of  $128 \times 3 \times 221 \times 221$  pixels. If each

---

<sup>7</sup>Compared to all earlier results, the experiment in this section is re-run three times with a new random seed and with faster cuDNN package on two Tesla K80 nodes ([developer.nvidia.com/cuDNN](https://developer.nvidia.com/cuDNN) and [github.com/soumith/cudnn.torch](https://github.com/soumith/cudnn.torch)). Also to clarify, the random initialization we use is by default in Torch’s implementation. All our methods are implemented in Torch ([torch.ch](https://torch.ch)). The Message Passing Interface implementation MVAPICH2 ([mvapich.cse.ohio-state.edu](https://mvapich.cse.ohio-state.edu)) is used for the GPU-CPU communication.



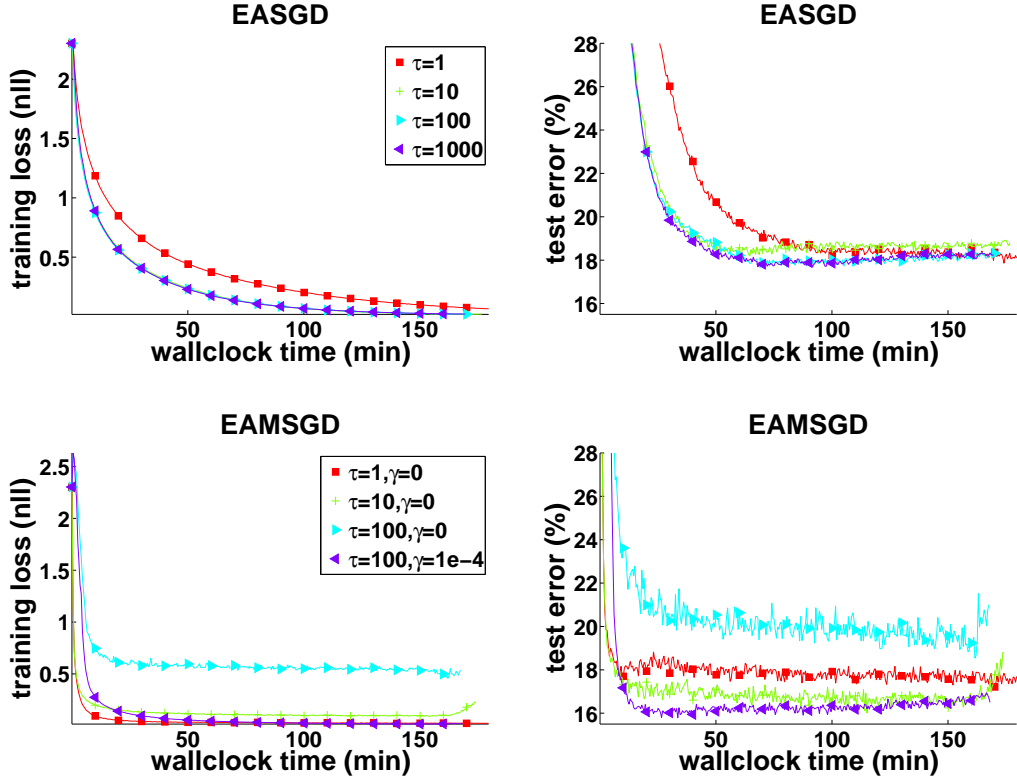


Figure 4.13: Convergence of the training loss (negative log-likelihood, original) and the test error (zoomed) computed for the center variable as a function of wallclock time for *EASGD* and *EAMSGD* ( $p = 16, \eta = 0.01, \beta = 0.9, \delta = 0.99$ ) on the *CIFAR* experiment with various communication period  $\tau$  and learning rate decay  $\gamma$ . The learning rate is decreased gradually over time based each local worker's own clock  $t$  with  $\eta_t = \eta / (1 + \gamma t)^{0.5}$ .

	$p = 1$	$p = 4$	$p = 8$	$p = 16$		$p = 1$	$p = 4$	$p = 8$
$\tau = 1$	12/1/0	11/2/3	11/2/5	11/2/9	$\tau = 1$	1248/20/0	1323/24/173	1239/61/284
$\tau = 10$	NA	11/2/1	11/2/1	12/2/1	$\tau = 10$	NA	1254/58/7	1266/84/11

Table 4.4: Approximate computation time, data loading time and parameter communication time [sec] for *DOWNPOUR* (top line for  $\tau = 1$ ) and *EASGD* (the time breakdown for *EAMSGD* is almost identical) (bottom line for  $\tau = 10$ ). Left time corresponds to *CIFAR* experiment and right table corresponds to *ImageNet* experiment. The computation and the communication time may have some overlap, due to the MPI implementation.

pixel value is represented by one byte, then each mini-batch is around 18 MB. As the whole dataset has around 1,300,000 images, it would take 174 GB. In fact, we have compressed the images as JPEG, so that the whole dataset is only around 36 GB. Thus we gain a compression ratio around  $1/5$ , i.e. we can assume each mini-batch size is  $18/5 = 3.6$  MB. The required data communication rate is thus 3.6 MB/sec. On the other hand, the *DOWNPOUR* method with  $\tau = 1$  requires communicating the whole model parameter per mini-batch. As the parameter size is around 233 MB, we need the network bandwidth at least 233 MB/sec per local worker (we have not even accounted for the gradient communication, which can double this cost). The parameter communication cost is thus at least 66 times of the data communication cost. In the *ImageNet* case, having access to the full dataset by the local workers is indeed a good tradeoff.

### 4.3.5 Time speed-up

In Figure 4.14 and 4.15, we summarize the wall clock time needed to achieve the same level of the test error for all the methods in the *CIFAR* and *ImageNet* experiment as a function of the number of local workers  $p$ . For the *CIFAR* (Figure 4.14) we examined the following levels: {21%, 20%, 19%, 18%} and for the *ImageNet* (Figure 4.15) we examined: {49%, 47%, 45%, 43%}. If some method does not appear on the figure for a given test error level, it indicates that this method never achieved this level. For the *CIFAR* experiment we observe that from among *EASGD*, *DOWNPOUR* and *MDOWNPOUR* methods, the *EASGD* method needs less time to achieve a particular level of test error.

We observe that with higher  $p$  each of these methods does not necessarily need less time to achieve the same level of test error. This seems counter intuitive though recall that the learning rate for the methods is selected based on the smallest achievable test error. For larger  $p$  smaller learning rates were selected than for smaller  $p$  which explains our results. Meanwhile, the *EAMSGD* method achieves significant speed-up over other methods for all the test error levels. For the *ImageNet* experiment we observe that all methods outperform *MSGD* and furthermore with  $p = 4$  or  $p = 8$  each of these methods requires less time to achieve the same level of test error.

## 4.4 Additional pseudo-codes of the algorithms

### DOWNPOUR pseudo-code

Algorithm 3 captures the pseudo-code of the implementation of *DOWNPOUR* used in this paper. Similar to the asynchronous behavior of *EASGD* that we have described in Chapter 2 (Section 2.2), *DOWNPOUR* method also performs several steps of local gradient updates by each worker  $i$  before pushing back the accumulated gradients  $v^i$  to the center variable. To be more precise, at the beginning of each period, the  $i$ -th worker reads a new center variable  $\tilde{x}$  from the parameter server. Then it performs  $\tau$  local *SGD* steps from the new  $\tilde{x}$ . All the gradients are accumulated (added) and at the end of that period, the total sum  $v^i$  is pushed (added) back to the parameter server. The center variable is updated by summing the accumulated gradients from any of the local workers. Notice that we do not use any adaptive learning scheme as having been done in [16].

### MDOWNPOUR pseudo-code

Algorithms 4 and 5 capture the pseudo-codes of the implementation of momentum *DOWNPOUR* (*MDOWNPOUR*) used in this paper. Algorithm 4 describes the behavior of each local worker and Algorithm 5 describes the behavior of the master. Note that

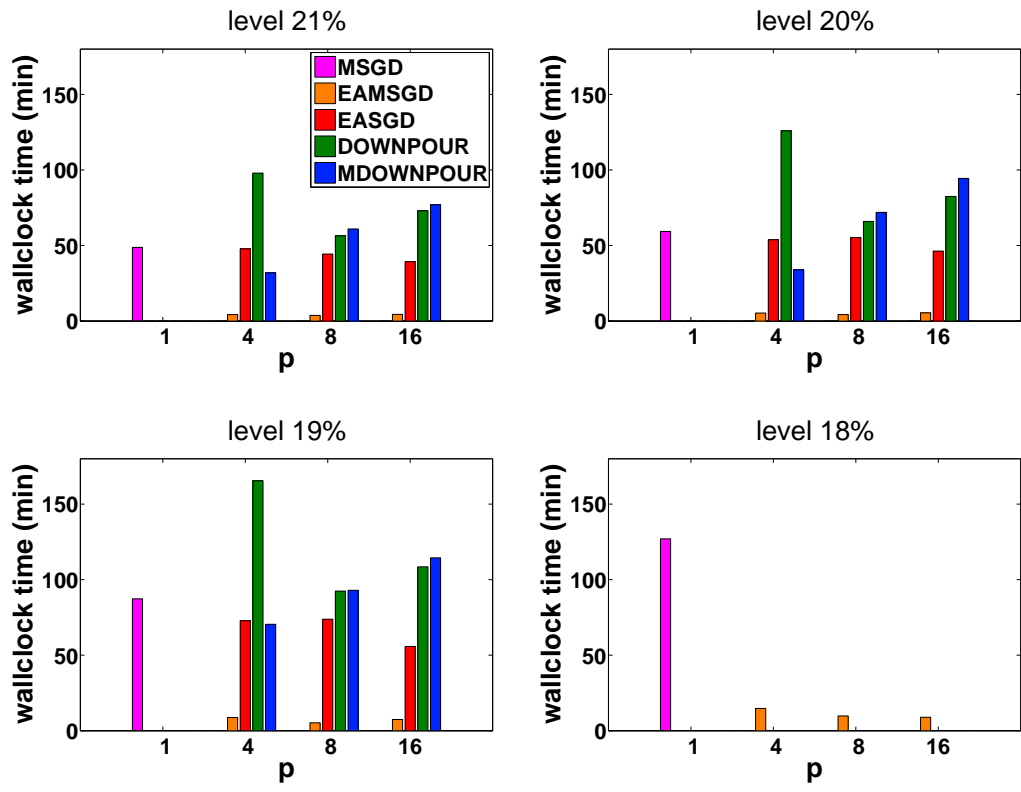


Figure 4.14: The wall clock time needed to achieve the same level of the test error  $thr$  as a function of the number of local workers  $p$  on the *CIFAR* dataset. From left to right:  $thr = \{21\%, 20\%, 19\%, 18\%\}$ . Missing bars denote that the method never achieved specified level of test error.

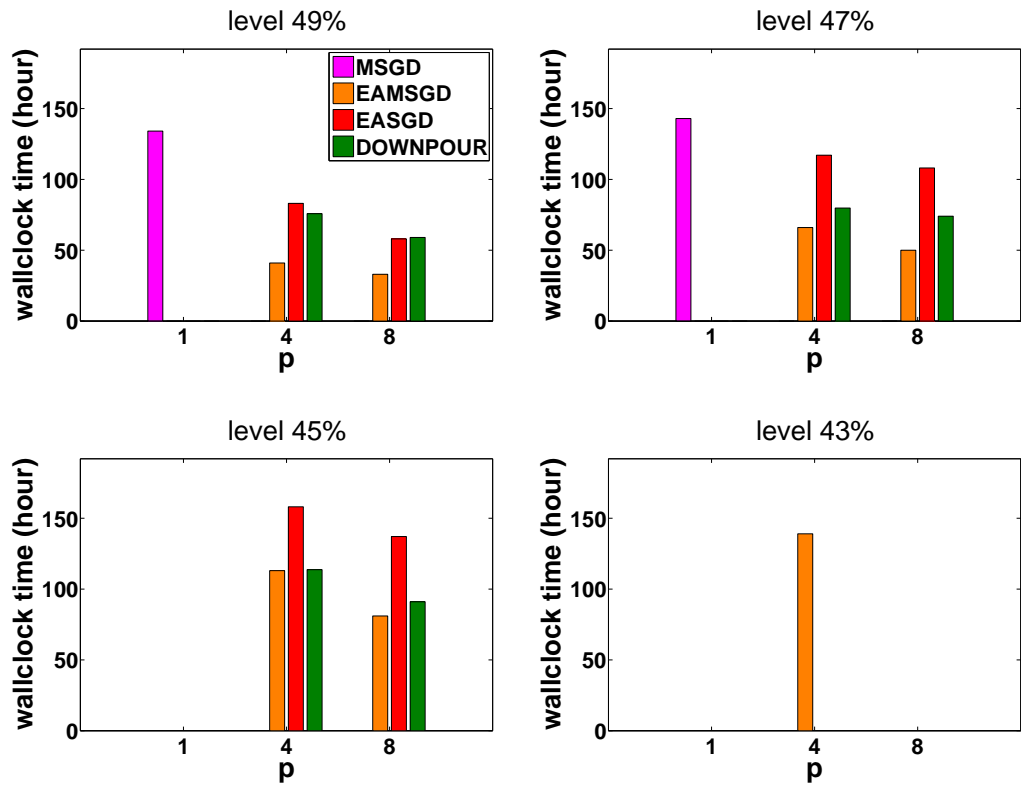


Figure 4.15: The wall clock time needed to achieve the same level of the test error  $thr$  as a function of the number of local workers  $p$  on the *ImageNet* dataset. From left to right:  $thr = \{49\%, 47\%, 45\%, 43\%\}$ . Missing bars denote that the method never achieved specified level of test error.

<b>Algorithm 3:</b> DOWNPOUR: Processing by worker $i$ and the master
<b>Input:</b> learning rate $\eta$ , communication period $\tau \in \mathbb{N}$
<b>Initialize:</b> $\tilde{x}$ is initialized randomly, $x^i = \tilde{x}$ , $v^i = 0$ , $t^i = 0$
<b>Repeat</b> <b>if</b> ( $\tau$ divides $t^i$ ) <b>then</b> $\tilde{x} \leftarrow \tilde{x} + v^i$ $x^i \leftarrow \tilde{x}$ $v^i \leftarrow 0$ <b>end</b> $x^i \leftarrow x^i - \eta g_{t^i}^i(x^i)$ $v^i \leftarrow v^i - \eta g_{t^i}^i(x^i)$ $t^i \leftarrow t^i + 1$ <b>Until forever</b>

unlike the *DOWNPOUR* method, we do not use the communication period  $\tau$ . This is because the Nesterov's momentum is applied to the center variable. Each worker reads an interpolated variable  $\tilde{x} + \delta v$  from the master, and then sends back the stochastic gradient evaluated at that point. In case  $p = 1$ , *MDOWNPOUR* is equivalent to the *MSGD* method.

<b>Algorithm 4:</b> MDOWNPOUR: Processing by worker $i$
<b>Initialize:</b> $x^i = \tilde{x}$
<b>Repeat</b> Receive $\tilde{x} + \delta v$ from the master: $x^i \leftarrow \tilde{x} + \delta v$ Compute gradient $g^i = g^i(x^i)$ Send $g^i$ to the master <b>Until forever</b>

<b>Algorithm 5:</b> MDOWNPOUR: Processing by the master
<b>Input:</b> learning rate $\eta$ , momentum term $\delta$
<b>Initialize:</b> $\tilde{x}$ is initialized randomly, $v^i = 0$ ,
<b>Repeat</b> Receive $g^i$ $v \leftarrow \delta v - \eta g^i$ $\tilde{x} \leftarrow \tilde{x} + v$ Send $\tilde{x} + \delta v$ <b>Until forever</b>

## Chapter 5

# The Limit in Speedup

This chapter studies the limitation in speedup of several stochastic optimization methods: the mini-batch *SGD*, the momentum *SGD* and the *EASGD* method. In Section 5.1, we study first the asymptotic phase of these methods using an additive noise model. The continuous-time SDE (stochastic differential equation) approximation of its *SGD* update is an Ornstein Uhlenbeck process. Then we study the initial phase of these methods using a multiplicative noise model in Section 5.2. The continuous-time SDE approximation of its *SGD* update is a Geometric Brownian motion. In Section 5.3, we study the stability of the critical points of a simple non-convex problem and discuss when the *EASGD* method can get trapped by a saddle point.

### 5.1 Additive noise

We (re-)study the simple additive noise model: one-dimensional quadratic objective with Gaussian noise (as in Section 3.1.1). The objective function evaluated at state  $x \in \mathbb{R}$  is defined to be the average loss of the quadratic form  $(hx - \xi)^2$ , i.e.

$$\min_{x \in \mathbb{R}} \mathbb{E}[(hx - \xi)^2]. \tag{5.1}$$

Here  $h > 0$  is a scalar, and the expectation is taken over the random variable  $\xi$ , which follows a Gaussian distribution. For simplicity, we assume further that  $\xi$  is zero mean and has a constant variance  $\sigma^2 > 0$ .

### 5.1.1 SGD with mini-batch

The update rule for the *SGD* method for solving the problem in Equation 5.1 is

$$x_{t+1} = x_t - \eta(hx_t - \xi_t), \quad (5.2)$$

where  $x_0$  is the initial starting point.

Notice that in the continuous-time limit, i.e. for small  $\eta$ , we can approximate the process in Equation 5.2 by an Ornstein Uhlenbeck process [32] as follows,

$$dX(t) = -hX(t)dt + \sqrt{\eta}\sigma dB(t).$$

In the discrete-time case, the bias term (the first-order moment  $\mathbb{E}x_t$ ) in Equation 5.2 will decrease at a linear rate  $1 - \eta h$ , i.e.

$$\mathbb{E}x_{t+1} = (1 - \eta h)\mathbb{E}x_t.$$

The second-order moment  $\mathbb{E}x_t^2$  changes as

$$\mathbb{E}x_{t+1}^2 = (1 - \eta h)^2\mathbb{E}x_t^2 + \eta^2\sigma^2. \quad (5.3)$$

Thus the variance will increase as

$$\begin{aligned} \mathbb{V}x_{t+1} &= \mathbb{E}x_{t+1}^2 - (\mathbb{E}x_{t+1})^2 = (1 - \eta h)^2\mathbb{E}x_t^2 + \eta^2\sigma^2 - (1 - \eta h)^2(\mathbb{E}x_t)^2 \\ &= (1 - \eta h)^2\mathbb{V}x_t + \eta^2\sigma^2, \end{aligned}$$



with  $\mathbb{V}x_0 = 0$ . As  $t \rightarrow \infty$ , we get  $\mathbb{V}x_\infty = \frac{\eta^2}{1-(1-\eta h)^2}\sigma^2$ . If we use mini-batch of size  $p$ , the variance of the noise  $\xi_t$  is then reduced by  $p$ , and this asymptotic variance becomes  $\frac{\eta^2}{1-(1-\eta h)^2}\frac{\sigma^2}{p}$ . The convergence rate of the bias term,  $1 - \eta h$ , is however not improved by increasing the mini-batch size  $p$ .

### 5.1.2 Momentum *SGD*

We now study the update rule for the *MSGD* (Nesterov's momentum) method for solving Equation 5.1,

$$\begin{aligned} v_{t+1} &= \delta v_t - \eta(h(x_t + \delta v_t) - \xi_t), \\ x_{t+1} &= x_t + v_{t+1}, \end{aligned} \tag{5.4}$$

where  $x_0$  is the initial starting point and  $v_0 = 0$  is the initial velocity set to zero.

Let  $\delta_h = \delta(1 - \eta h)$ ,  $\eta_h = \eta h$ , then Equation 5.4 is equivalent to

$$\begin{aligned} v_{t+1} &= \delta_h v_t - \eta_h x_t + \eta \xi_t, \\ x_{t+1} &= x_t + v_{t+1}. \end{aligned}$$

In case that  $\delta_h$  is chosen independently of  $h$ , the update is no different to the heavy ball method [44]. But the momentum term  $\delta_h$  above equals  $\delta(1 - \eta h)$ , thus it implicitly depends on  $h$ . As we usually choose  $\delta$  to be smaller than one,  $\delta_h$  is upper bounded by  $1 - \eta h$ . This fact saves us from the variance explosion as  $\delta$  tends to one.

More precisely, the second-order moment equation can be computed as follows,

$$\begin{aligned} v_{t+1}^2 &= (\delta_h v_t - \eta_h x_t + \eta \xi_t)^2 \\ &= \delta_h^2 v_t^2 + \eta_h^2 x_t^2 + \eta^2 \xi_t^2 - 2\delta_h \eta_h v_t x_t + 2\delta_h \eta v_t \xi_t - 2\eta_h \eta x_t \xi_t \\ x_{t+1}^2 &= (x_t + v_{t+1})^2 = x_t^2 + 2v_{t+1}x_t + v_{t+1}^2 \\ v_{t+1}x_t &= (\delta_h v_t - \eta_h x_t + \eta \xi_t)x_t = \delta_h v_t x_t - \eta_h x_t^2 + \eta x_t \xi_t \\ v_{t+1}x_{t+1} &= v_{t+1}(x_t + v_{t+1}) = v_{t+1}x_t + v_{t+1}^2. \end{aligned} \tag{5.5}$$

Now taking expectation on both sides of the Equation 5.5, we obtain the following recursive relation

$$\begin{pmatrix} \mathbb{E}v_{t+1}^2 \\ \mathbb{E}v_{t+1}x_{t+1} \\ \mathbb{E}x_{t+1}^2 \end{pmatrix} = \underbrace{\begin{pmatrix} \delta_h^2 & -2\delta_h\eta_h & \eta_h^2 \\ \delta_h^2 & \delta_h(1-2\eta_h) & -\eta_h(1-\eta_h) \\ \delta_h^2 & 2\delta_h(1-\eta_h) & (1-\eta_h)^2 \end{pmatrix}}_M \begin{pmatrix} \mathbb{E}v_t^2 \\ \mathbb{E}v_t x_t \\ \mathbb{E}x_t^2 \end{pmatrix} + \begin{pmatrix} \eta^2\sigma^2 \\ \eta^2\sigma^2 \\ \eta^2\sigma^2 \end{pmatrix}. \quad (5.6)$$

To see the asymptotic behavior, assume  $\underline{v}_\infty^2 = \lim_{t \rightarrow \infty} \mathbb{E}v_t^2$ ,  $\underline{vx}_\infty = \lim_{t \rightarrow \infty} \mathbb{E}v_t x_t$ , and  $\underline{x}_\infty^2 = \lim_{t \rightarrow \infty} \mathbb{E}x_t^2$ . Then by solving

$$\begin{pmatrix} \underline{v}_\infty^2 \\ \underline{vx}_\infty \\ \underline{x}_\infty^2 \end{pmatrix} = \begin{pmatrix} \delta_h^2 & -2\delta_h\eta_h & \eta_h^2 \\ \delta_h^2 & \delta_h(1-2\eta_h) & -\eta_h(1-\eta_h) \\ \delta_h^2 & 2\delta_h(1-\eta_h) & (1-\eta_h)^2 \end{pmatrix} \begin{pmatrix} \underline{v}_\infty^2 \\ \underline{vx}_\infty \\ \underline{x}_\infty^2 \end{pmatrix} + \begin{pmatrix} \eta^2\sigma^2 \\ \eta^2\sigma^2 \\ \eta^2\sigma^2 \end{pmatrix},$$

we obtain

$$\begin{aligned} \underline{v}_\infty^2 &= \frac{2}{(1-\delta_h)(2(1+\delta_h)-\eta_h)}\eta^2\sigma^2, \\ \underline{vx}_\infty &= \frac{1}{(1-\delta_h)(2(1+\delta_h)-\eta_h)}\eta^2\sigma^2, \\ \underline{x}_\infty^2 &= \frac{1+\delta_h}{\eta_h(1-\delta_h)(2(1+\delta_h)-\eta_h)}\eta^2\sigma^2. \end{aligned} \quad (5.7)$$

From Equation 5.7, we see that for the asymptotic variance  $\underline{x}_\infty^2$  to be strictly positive, we should assume  $-1 < \delta_h < 1$  and  $0 < \eta_h < 2(1+\delta_h)$ . Moreover, compared to the asymptotic variance of *SGD* (the case that  $\delta = 0$ ), we can check, for example, that in the region  $\eta_h \in (0, 1)$  and  $\delta_h \in (0, 1)$ , the asymptotic variance of *MSGD* is always larger.

On the other hand, the above condition  $-1 < \delta_h < 1$  and  $0 < \eta_h < 2(1+\delta_h)$  is also the condition for the matrix  $M$  in Equation 5.6 to remain (strictly) stable, i.e. the largest absolute eigenvalue is (strictly) smaller than one. In fact, we have three eigenvalues for

the matrix  $M$  as follows,

$$\begin{aligned} z_1 &= \delta_h, \\ z_2 &= \frac{(1-\eta_h)^2 - 2\eta_h\delta_h + \delta_h^2 - \sqrt{((1-\eta_h)^2 - 2\eta_h\delta_h + \delta_h^2)^2 - 4\delta_h^2}}{2}, \\ z_3 &= \frac{(1-\eta_h)^2 - 2\eta_h\delta_h + \delta_h^2 + \sqrt{((1-\eta_h)^2 - 2\eta_h\delta_h + \delta_h^2)^2 - 4\delta_h^2}}{2}. \end{aligned} \tag{5.8}$$

Let  $2b = z_2 + z_3 = (1 - \eta_h)^2 - 2\eta_h\delta_h + \delta_h^2$  and  $c = z_2z_3 = \delta_h^2$ . Then  $z_2$  and  $z_3$  are the two roots (in  $z$ ) of  $z^2 - 2bz + c = 0$ . In general, if  $z_2, z_3$  are real-valued, i.e.  $b^2 - c > 0$ , then we have two cases:

- $b > \sqrt{c}$  implies  $z_3 = b + \sqrt{b^2 - c} > \sqrt{c} > z_2 = b - \sqrt{b^2 - c} > 0$ ,
- $b < -\sqrt{c}$  implies  $z_2 = b - \sqrt{b^2 - c} < -\sqrt{c} < z_3 = b + \sqrt{b^2 - c} < 0$ .

On the other hand, if  $z_2, z_3$  are not real-valued, i.e.  $b^2 - c \leq 0$ , then  $|z_3| = |z_1| = |z_2| = \sqrt{c}$ . Recall that  $|z_1| = |\delta_h| = \sqrt{c}$ .

We see thus for the (strict) stability of the matrix  $M$  in Equation 5.6, we need  $\sqrt{c} = |\delta_h| < 1$  and  $|z_3| < 1$  because the condition  $b < -\sqrt{c}$  is never satisfied for any real  $\delta_h$ . The condition for  $|z_3| = b + \sqrt{b^2 - c} < 1$  is that  $2b - c < 1$ , i.e.  $0 < \eta_h < 2(1 + \delta_h)$ .

Moreover, as in [32], we can try to minimize  $|z_3|$  with respect to  $\delta$  such that the rate of convergence of the second order moment is maximized for a given  $\eta$ . We shall prove that the minimal  $|z_3|$  is achieved at  $\delta_h = (\sqrt{\eta_h} - 1)^2$ . In fact, it happens when  $b = \sqrt{c}$ , i.e. the optimal rate is at the edge where the eigenvalues transit from real-valued to the complex-valued. Notice that  $b = \sqrt{c}$  gives us two positive solutions:  $\delta_h = (\sqrt{\eta_h} - 1)^2$  and  $\delta_h = (\sqrt{\eta_h} + 1)^2$ . Since  $b = \frac{1}{2}[(\delta_h - \eta_h)^2 + 1 - 2\eta_h]$  is a quadratic function in  $\delta_h$ , the fact that  $b = \sqrt{c}$  has two positive solutions means that the quadratic function intersects twice with the line  $\delta_h$  in the first orthant. If  $\eta_h \geq 1/4$ , the first intersection point is to the left of the minimum of  $b$ , and the second intersection point is to the right of the minimum, i.e.  $(\sqrt{\eta_h} - 1)^2 \leq \eta_h < (\sqrt{\eta_h} + 1)^2$ . But if  $0 < \eta_h < 1/4$ , both intersection points will be both to the right of the minimum of  $b$ , i.e.  $\eta_h < (\sqrt{\eta_h} - 1)^2 < (\sqrt{\eta_h} + 1)^2$ .

Nevertheless, in either case, we can show  $b > \sqrt{c}$  whenever  $\delta_h < (\sqrt{\eta_h} - 1)^2$ . Thus, in the range  $\delta_h < (\sqrt{\eta_h} - 1)^2$ , we have  $z_3 > \sqrt{c} > z_2 > 0$ . We thus only need to find the minimum of  $z_3$  in the range  $\delta_h \in (-1, (\sqrt{\eta_h} - 1)^2]$ , because  $|z_3| = \delta_h$  is monotonically increasing in the range  $1 > \delta_h > (\sqrt{\eta_h} - 1)^2$ . We now show that the minimal value of  $z_3$  is  $(\sqrt{\eta_h} - 1)^2$  in this range. In fact, if  $\eta_h \geq 1/4$ ,  $b$  is monotonically decreasing in this range, thus  $z_3 \geq b$  reaches its minimum at  $\delta_h = (\sqrt{\eta_h} - 1)^2$  with  $z_3 = b$ . If  $0 < \eta_h < 1/4$ ,  $b$  firstly decreases for  $\delta_h < \eta_h$ , then increases for  $\eta_h \leq \delta_h \leq (1 - \sqrt{\eta_h})^2$ . We show that  $z_3$  is still monotonically decreasing in these two ranges of  $\delta_h$ . We check whether  $\frac{\partial z_3}{\partial \delta_h} = (\delta_h - \eta_h)(1 + \frac{b}{\sqrt{b^2 - c}}) - \delta_h \frac{1}{\sqrt{b^2 - c}} < 0$  holds. For  $\delta_h < \eta_h$ , this is equivalent to  $z_3 > \frac{\delta_h}{\delta_h - \eta_h}$ ; for  $\eta_h < \delta_h \leq (1 - \sqrt{\eta_h})^2$ , this is equivalent to  $z_3 < \frac{\delta_h}{\delta_h - \eta_h}$ . The former one is true in the range  $0 < \delta_h < \eta_h$ , because  $z_3$  is always positive. One can check that this is still true in the range  $-1 < \delta_h < 0$ . The latter one is equivalent to  $b + \sqrt{b^2 - c} < \frac{\delta_h}{\delta_h - \eta_h}$ . Taking square on both sides, we can check that  $b < \frac{\delta_h}{\delta_h - \eta_h}$  and  $b^2 - c < (\frac{\delta_h}{\delta_h - \eta_h} - b)^2$ , based on  $0 < \delta_h - \eta_h < 1 - 2\sqrt{\eta_h} < 1$ .

Thus we conclude that for a fixed  $\eta_h$  such that  $0 < \eta_h < 2(1 + \delta_h)$ , the minimal  $|z_3|$  over  $-1 < \delta_h < 1$  is obtained at

$$\delta_h = (\sqrt{\eta_h} - 1)^2,$$

with the minimal value  $z_3 = b + \sqrt{b^2 - c} = b = \sqrt{c} = \delta_h = (\sqrt{\eta_h} - 1)^2$ . Compared to the rate  $(1 - \eta_h)^2$  in the *SGD* case (Equation 5.3), *MSGD* can indeed help if  $\eta_h$  is small enough (usually  $\eta_h = \eta h$  is close to the inverse of the condition number of the Hessian in higher dimensional case). To check our above reasoning, we have computed numerically the spectral norm of the matrix  $M$ ,  $\text{sp}(M)$ , which is given in Figure 5.1. Note that when  $\eta_h > 1$ , we have the optimal momentum rate being negative, i.e.  $\delta = \frac{(\sqrt{\eta_h} - 1)^2}{(1 - \eta_h)} < 0$ .

Perhaps the most special behavior for *MSGD* is that when the momentum rate  $\delta$  gets close to 1, the asymptotic variance of  $\underline{x}_\infty^2$  in Equation 5.7 still remains bounded, i.e.  $\underline{x}_\infty^2 = \frac{1 + 1 - \eta_h}{\eta_h^2(2(1 + 1 - \eta_h) - \eta_h)} \eta^2 \sigma^2 = \frac{2 - \eta_h}{4 - 3\eta_h} \frac{\sigma^2}{h^2}$  for  $\delta = 1$  (or  $\delta_h = 1 - \eta_h$ ). This contrasts to the behavior of the heavy ball method, whose asymptotic variance tends to infinity as

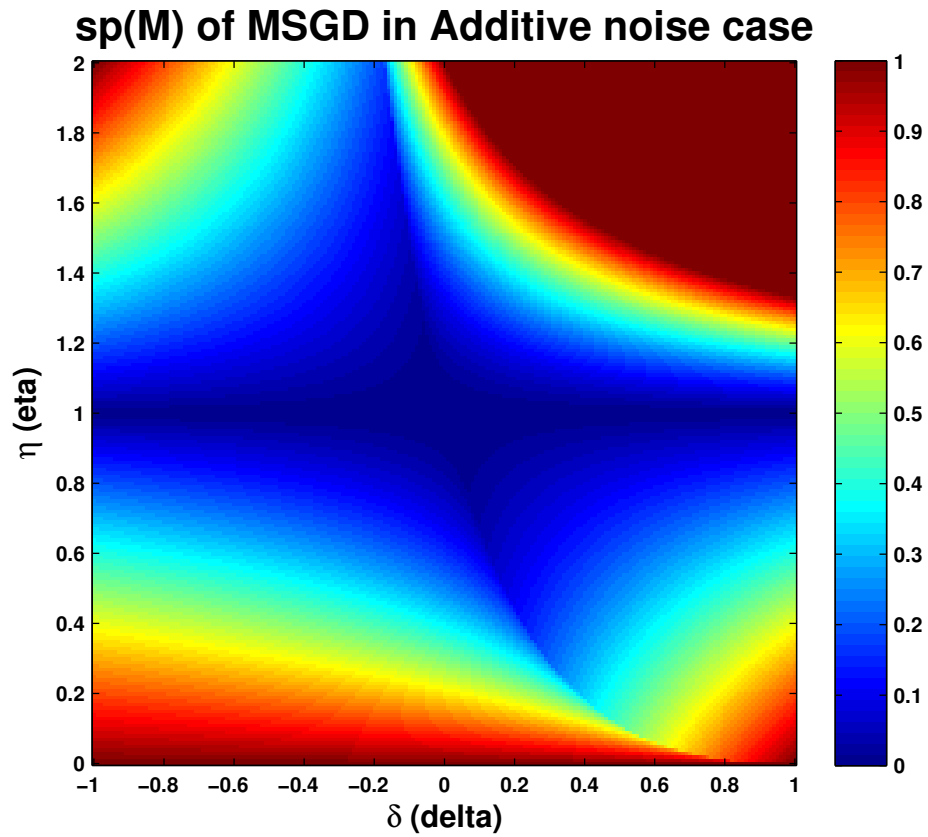


Figure 5.1: The largest absolute eigenvalue of the matrix  $M$  ( $\text{sp}(M)$ ) in Equation 5.6 as a function of the learning rate  $\eta \in (0, 2)$  and the momentum rate  $\delta \in (-1, 1)$ .  $h = 1$ .

$\delta \rightarrow 1$  [32].

### 5.1.3 *EASGD* and *EAMSGD*

We can now study the moment equation for the *EASGD* and *EAMSGD* method. For *EASGD*, we have the following update rules

$$\begin{aligned} x_{t+1}^i &= x_t^i - \eta(hx_t^i - \xi_t^i) + \alpha(\tilde{x}_t - x_t^i), \\ \tilde{x}_{t+1} &= \tilde{x}_t + \beta\left(\frac{1}{p}\sum_{i=1}^p x_t^i - \tilde{x}_t\right). \end{aligned} \tag{5.9}$$

Denote  $y_t = \frac{1}{p}\sum_{i=1}^p x_t^i$  and  $\xi_t = \frac{1}{p}\sum_{i=1}^p \xi_t^i$ , then Equation 5.9 can be reduced to

$$\begin{aligned} y_{t+1} &= y_t - \eta(hy_t - \xi_t) + \alpha(\tilde{x}_t - y_t), \\ \tilde{x}_{t+1} &= \tilde{x}_t + \beta(y_t - \tilde{x}_t). \end{aligned} \tag{5.10}$$

Similar to *MSGD*, the second-order moment equation for Equation 5.10 is as follows,

$$\begin{aligned} y_{t+1}^2 &= ((1 - \eta h - \alpha)y_t + \alpha\tilde{x}_t + \eta\xi_t)^2 \\ &= (1 - \eta h - \alpha)^2 y_t^2 + \alpha^2 \tilde{x}_t^2 + \eta^2 \xi_t^2 \\ &\quad + 2\alpha(1 - \eta h - \alpha)y_t \tilde{x}_t + 2\eta(1 - \eta h - \alpha)y_t \xi_t + 2\alpha\eta \tilde{x}_t \xi_t, \\ y_{t+1}\tilde{x}_{t+1} &= ((1 - \eta h - \alpha)y_t + \alpha\tilde{x}_t + \eta\xi_t)((1 - \beta)\tilde{x}_t + \beta y_t) \\ &= (1 - \eta h - \alpha)\beta y_t^2 + (1 - \eta h - \alpha)(1 - \beta)y_t \tilde{x}_t + \alpha\beta y_t \tilde{x}_t + \alpha(1 - \beta)\tilde{x}_t^2 \\ &\quad + \eta\xi_t((1 - \beta)\tilde{x}_t + \beta y_t), \\ \tilde{x}_{t+1}^2 &= (1 - \beta)^2 \tilde{x}_t^2 + \beta^2 y_t^2 + 2\beta(1 - \beta)y_t \tilde{x}_t. \end{aligned} \tag{5.11}$$

Taking expectation on both sides of the Equation 5.11, we get

$$\begin{pmatrix} \mathbb{E}y_{t+1}^2 \\ \mathbb{E}y_{t+1}\tilde{x}_{t+1} \\ \mathbb{E}\tilde{x}_{t+1}^2 \end{pmatrix} = \underbrace{\begin{pmatrix} (1-\eta h-\alpha)^2 & 2\alpha(1-\eta h-\alpha) & \alpha^2 \\ (1-\eta h-\alpha)\beta & (1-\eta h-\alpha)(1-\beta)+\alpha\beta & \alpha(1-\beta) \\ \beta^2 & 2\beta(1-\beta) & (1-\beta)^2 \end{pmatrix}}_M \begin{pmatrix} \mathbb{E}y_t^2 \\ \mathbb{E}y_t\tilde{x}_t \\ \mathbb{E}\tilde{x}_t^2 \end{pmatrix} + \begin{pmatrix} \eta^2\frac{\sigma^2}{p} \\ 0 \\ 0 \end{pmatrix}. \quad (5.12)$$

Similar to the analysis of *MSGD*, we get the following asymptotic variance,

$$\underline{y}_\infty^2 = \lim_{t \rightarrow \infty} \mathbb{E}y_t^2 = \frac{(2-\beta)(1-\beta)\eta_h + \beta(2-\alpha-\beta)}{\eta_h[(2-\beta)(2-\eta_h) - 2\alpha][\alpha + \beta + \eta_h(1-\beta)]} \frac{\eta^2\sigma^2}{p}, \quad (5.13)$$

$$\underline{y\tilde{x}}_\infty = \lim_{t \rightarrow \infty} \mathbb{E}y_t\tilde{x}_t = \frac{\beta((2-\beta)(1-\eta_h) - \alpha)}{\eta_h[(2-\beta)(2-\eta_h) - 2\alpha][\alpha + \beta + \eta_h(1-\beta)]} \frac{\eta^2\sigma^2}{p},$$

$$\underline{\tilde{x}}_\infty^2 = \lim_{t \rightarrow \infty} \mathbb{E}\tilde{x}_t^2 = \frac{-\beta(1-\beta)\eta_h + \beta(2-\alpha-\beta)}{\eta_h[(2-\beta)(2-\eta_h) - 2\alpha][\alpha + \beta + \eta_h(1-\beta)]} \frac{\eta^2\sigma^2}{p}. \quad (5.14)$$

For the above asymptotic variance to be positive, in particular  $\underline{\tilde{x}}_\infty^2$ , we can assume

$$\begin{aligned} \eta_h &> 0, \\ \beta &> 0, \\ (2-\beta)(2-\eta_h) - 2\alpha &> 0, \\ \frac{2-\alpha-\beta-\eta_h+\beta\eta_h}{\alpha+\beta+\eta_h(1-\beta)} &> 0. \end{aligned} \quad (5.15)$$

Moreover, if  $0 < \beta < 1$ , the asymptotic variance of the center variable  $\underline{\tilde{x}}_\infty^2$  is strictly smaller than that of the spatial average  $\underline{y}_\infty^2$  (by comparing Equation 5.13 and 5.14). Interestingly, if  $\beta > 1$ , it becomes strictly bigger.

Note that it's still not very clear whether Equation 5.15 is the necessary and sufficient condition for the matrix  $M$  (in Equation 5.12) to be (strictly) stable. However, if we look back to the earlier result in Section 3.1.1 about the condition in Equation 3.4, they are nearly the same, except maybe for the last formula. The last formula in Equation 5.15 reads  $0 < \alpha + \beta + \eta_h - \beta\eta_h < 2$ . The left inequality implies  $\beta\eta_h < \alpha + \beta + \eta_h$ . Based on

the third formula in Equation 5.15, we have thus  $\alpha + \beta + \eta_h < 2 + \frac{\beta\eta_h}{2} < 2 + \frac{\alpha + \beta + \eta_h}{2}$ . This gives us the last formula for the condition in Equation 3.4, which is  $\alpha + \beta + \eta_h < 4$ . Conversely, assuming  $\alpha > 0$  (as assumed in Section 3.1.1 for  $\beta = p\alpha$ ), then one can check that  $\alpha + \beta + \eta_h < 4$  implies the last formula in Equation 5.15, which also reads  $\alpha - 1 < (\eta_h - 1)(\beta - 1) < \alpha + 1$ .

Perhaps the most striking observation is that the optimal  $\alpha$  such that the convergence rate (of the moment Equation 5.12) is maximized turns out to be negative, given  $\eta_h > 0$  and  $\beta > 0$  fixed.

The three eigenvalues of the matrix  $M$  in Equation 5.12 are

$$\begin{aligned} z_1 &= -\alpha + (1 - \eta_h)(1 - \beta), \\ z_2 &= b - \sqrt{b^2 - c}, \\ z_3 &= b + \sqrt{b^2 - c}, \end{aligned} \tag{5.16}$$

where  $b = \frac{1}{2}[(\alpha - (1 - \eta_h - \beta))^2 + 1 - 2\beta\eta_h]$ ,  $c = [\alpha - (1 - \eta_h)(1 - \beta)]^2 = z_1^2$ . Denote  $\alpha' = z_1 = -\alpha + (1 - \eta_h)(1 - \beta)$ , then  $b = \frac{1}{2}[(\alpha' - \eta_h\beta)^2 + 1 - 2\beta\eta_h]$  and  $c = (\alpha')^2$ . Using exactly the same analysis and the result from the *MSGD* case, we get that the minimal  $|z_3|$  over  $-1 < \alpha' < 1$  is obtained at

$$\alpha' = (\sqrt{\beta\eta_h} - 1)^2,$$

which is equivalent to

$$\alpha = -(\sqrt{\beta} - \sqrt{\eta_h})^2. \tag{5.17}$$

The situation that the coupling constant  $\alpha$  being negative, while  $\beta$  being positive suggests a very different perspective to understand *EASGD*. It seems that our earlier condition  $\beta = p\alpha$  is unnecessary and is sub-optimal in this case. To check our above results, we have computed numerically the spectral norm of the matrix  $M$  for a fixed  $\beta = 0.9$ , which is given in Figure 5.2. However, we are in danger this time if we were to simulate



*EASGD* using the optimal  $\alpha$  given in Equation 5.17. In Figure 5.3, we illustrate an unstable behavior in such optimal case. The reason is that our above analysis is based on the reduced Equation 5.10, rather than the original Equation 5.9.

In the original Equation 5.9, we have the following form of the drift matrix

$$M_p = \begin{bmatrix} 1 - \alpha - \eta_h & 0 & \dots & 0 & \alpha \\ 0 & 1 - \alpha - \eta_h & 0 & \dots & \alpha \\ \dots & 0 & \dots & 0 & \dots \\ 0 & \dots & 0 & 1 - \alpha - \eta_h & \alpha \\ \beta' & \beta' & \dots & \beta' & 1 - \beta \end{bmatrix}, \quad (5.18)$$

whose first  $p$  rows correspond to the local workers' updates, and the last row correspond to the master's update. The  $p + 1$  eigenvalues  $M_p$  can be computed recursively as follows: let  $I_{p+1}(z) = \det(M_p - z)$ , then we have  $I_{p+1}(z) = (1 - \alpha - \eta_h - z)I_p(z) - \alpha\beta'(1 - \alpha - \eta_h - z)^{p-1} = \dots = (1 - \alpha - \eta_h - z)^{p-1}[(1 - \beta - z)(1 - \alpha - \eta_h - z) - p\alpha\beta']$ . Notice that  $\beta' = \beta/p$ , thus we have two eigenvalues which do not depend on  $p$ , i.e.  $(1 - \beta - z)(1 - \alpha - \eta_h - z) - \alpha\beta = 0$ , and an extra eigenvalue which only shows up for  $p > 1$ , i.e.  $z = 1 - \alpha - \eta_h$ . This extra eigenvalue  $z = 1 - \alpha - \eta_h$  is completely ignored in our reduced Equation 5.10. Then what is the optimal  $\alpha$  for the matrix  $M_p$  instead? The three eigenvalues of the matrix  $M_p$  in Equation 5.18 are

$$\begin{aligned} z_1 &= 1 - \alpha - \eta_h, \\ z_2 &= b - \sqrt{b^2 - c}, \\ z_3 &= b + \sqrt{b^2 - c}, \end{aligned} \quad (5.19)$$

where  $b = \frac{1}{2}(2 - \beta - \eta_h - \alpha)$ ,  $c = (1 - \eta_h)(1 - \beta) - \alpha$ .

Given  $\eta_h > 0$  and  $\beta > 0$  fixed, we shall prove that

- if  $\beta > \eta_h$ : the optimal  $\alpha = 0$ .
- if  $\beta < \eta_h$ : the optimal  $\alpha = -(\sqrt{\beta} - \sqrt{\eta_h})^2$ .

### sp(M) of EASGD ( $\beta=0.9$ ) in Additive noise case




Figure 5.2: The largest absolute eigenvalue of the matrix  $M$  in Equation 5.12 as a function of the learning rate  $\eta \in (0, 2)$  and the moving rate  $\alpha \in (-1, 1)$ .  $h = 1$  and  $\beta = 0.9$ .

EASGD ( $\eta=0.1, \beta=0.9, p=4$ ) in Additive noise case




Figure 5.3: Three independent simulations of *EASGD* using the elastic averaging  $\alpha = \beta/p$  and the optimal  $\alpha$  given in Equation 5.17. The x-axis is the time step  $t$  in the *EASGD* updates of Equation 5.9. The y-axis is the squared distance of the center variable to the optimum zero, i.e.  $\tilde{x}_t^2$ . We have chosen  $h = 1$ ,  $\sigma = 10^{-2}$ ,  $p = 4$ ,  $\eta = 0.1$  and  $\beta = 0.9$ .

The proof idea is similar to the *MSGD* case, so we shall be more brief. Let's focus on the variable  $c = (1 - \eta_h)(1 - \beta) - \alpha$  instead of the variable  $\alpha$  itself. Rewrite Equation 5.19 with  $z_1 = c + \beta(1 - \eta_h)$ ,  $b = \frac{1}{2}(c - \beta\eta_h + 1)$ . We find that for  $b^2 > c$ , we only need  $c > c_2$  or  $c < c_1$ , where  $c_1 = (\sqrt{\beta\eta_h} - 1)^2$ , and  $c_2 = (\sqrt{\beta\eta_h} + 1)^2$ . We also observe that the line  $z_1$  as a function of  $c$  intersects with the quadratic curve  $z_2$  or  $z_3$  at  $c_0 = (1 - \eta_h)(1 - \beta)$  (i.e. at  $\alpha = 0$ ). At  $c = c_0$ ,  $z_1 = 1 - \eta_h$ ,  $z_2 = 1 - \frac{1}{2}(\beta + \eta_h) - \frac{1}{2}|\beta - \eta_h|$ , and  $z_3 = 1 - \frac{1}{2}(\beta + \eta_h) + \frac{1}{2}|\beta - \eta_h|$ . If  $\beta > \eta_h$ , then  $z_3 = z_1$  at  $c = c_0$ . Otherwise  $z_2 = z_1$  at  $c = c_0$ . We can check that the negative optimal  $\alpha = -(\sqrt{\beta} - \sqrt{\eta_h})^2$  is given at  $c = c_1$ , where  $z_2$  and  $z_3$  meet and transit from real-valued to complex-valued. Note that  $z_1$  is an increasing function of  $c$ , and if it intersects with  $z_3$ , the optimal  $\alpha$  becomes 0 (rather than being negative). They are illustrated in Figure 5.4 and 5.5 as a function of  $\alpha$  under the two conditions, i.e.  $\beta > \eta_h$  and  $\beta < \eta_h$ . We also computed numerically the spectral norm of the matrix  $M_p$  for a fixed  $\beta = 0.9$ , which is given in Figure 5.6. Finally, we show in Figure 5.7 the optimal case of *EASGD* under the condition  $\beta < \eta_h$ .

We have studied the the second-order moment equation 5.12 of the reduced system (Equation 5.10), and the first-order moment equation 5.18 of the original system (Equation 5.9) for the *EASGD* method. We found that the reduced system can lose critical information about the stability of the original system. However, the eigenvalues are still closely related in the first-order moment matrix  $M_p$  and the second-order moment  $M$ , i.e. the  $z_2$  and  $z_3$  in Equation 5.16 and 5.19. Thus for the *EAMSGD* method, we shall only focus on the first-order moment equation. Its asymptotic variance, which can obtained from the second-order moment equation, is rather complicated and will not be discussed.

For *EAMSGD*, we have the following update rules

$$\begin{aligned} v_{t+1}^i &= \delta v_t^i - \eta(h(x_t^i + \delta v_t^i) - \xi_t^i), \\ x_{t+1}^i &= x_t^i + v_{t+1}^i + \alpha(\tilde{x}_t - x_t^i), \\ \tilde{x}_{t+1} &= \tilde{x}_t + \beta\left(\frac{1}{p} \sum_{i=1}^p x_t^i - \tilde{x}_t\right). \end{aligned}$$




Figure 5.4: The absolute value of the eigenvalues of  $z_1$ ,  $z_2$  and  $z_3$  in Equation 5.19.  $\eta_h = 0.1$  and  $\beta = 0.9$ .




Figure 5.5: The absolute value of the eigenvalues of  $z_1$ ,  $z_2$  and  $z_3$  in Equation 5.19.  $\eta_h = 1.5$  and  $\beta = 0.9$ .

### sp(M) of EASGD ( $\beta=0.9, p=2$ ) in Additive noise case




Figure 5.6: The largest absolute eigenvalue of the matrix  $M_p$  in Equation 5.18 as a function of the learning rate  $\eta \in (0, 2)$  and the moving rate  $\alpha \in (-1, 1)$ .  $h = 1$  and  $\beta = 0.9$ . Note that we computed this spectrum using  $p = 2$ , as we have discussed in the text, it is independent of the choice of  $p$  for  $p > 1$ .

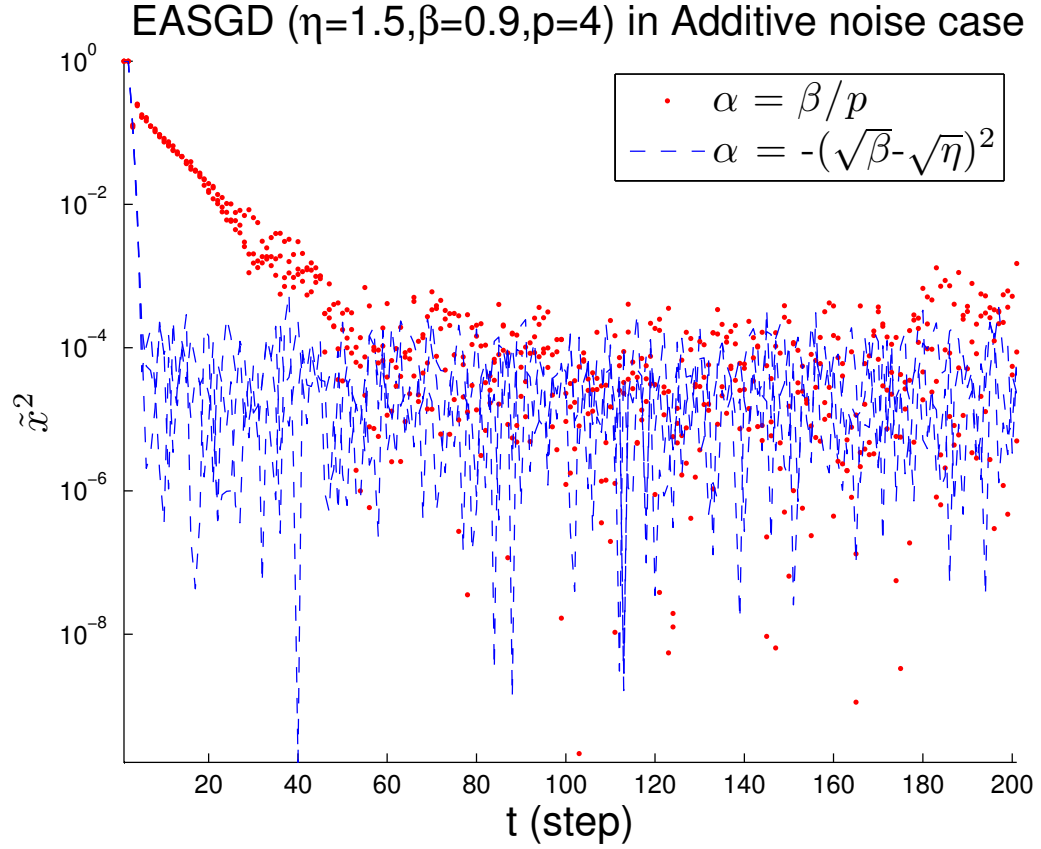


Figure 5.7: Three independent simulations of *EASGD* using the elastic averaging  $\alpha = \beta/p$  and the optimal  $\alpha$  given in Equation 5.17. The x-axis is the time step  $t$  in the *EASGD* updates of Equation 5.9. The y-axis is the squared distance of the center variable to the optimum zero, i.e.  $\tilde{x}_t^2$ . We have chosen  $h = 1$ ,  $\sigma = 10^{-2}$ ,  $p = 4$ ,  $\eta = 1.5$  and  $\beta = 0.9$ .

We have the following form of the drift matrix for the first-order moment equation,

$$M_p = \begin{bmatrix} \delta_h & -\eta_h & 0 & 0 & \dots & \dots & 0 \\ \delta_h & 1 - \eta_h - \alpha & 0 & 0 & \dots & \dots & \alpha \\ 0 & 0 & \delta_h & -\eta_h & \dots & \dots & 0 \\ 0 & 0 & \delta_h & 1 - \eta_h - \alpha & \dots & \dots & \alpha \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & \beta' & 0 & \beta' & \dots & \dots & 1 - \beta \end{bmatrix}, \quad (5.20)$$

such that  $[v_{t+1}^1, x_{t+1}^1, v_{t+1}^2, x_{t+1}^2, \dots, \tilde{x}_{t+1}]^T = M_p[v_t^1, x_t^1, v_t^2, x_t^2, \dots, \tilde{x}_t]^T$ . Recall that  $\delta_h = \delta(1 - \eta h)$ ,  $\eta_h = \eta h$ , and  $\beta' = \beta/p$ . We can compute the eigenvalues of the above drift matrix  $M_p$  recursively, and one can check that they are again independent of the choice of  $p$  for  $p > 1$ , as in the *EASGD* case. More precisely, we have

$$\det|M_p - z| = (u(z))^{p-1}v(z) = 0, \quad (5.21)$$

where  $u(z) = z^2 - (1 - \eta - \alpha - \delta)z + \delta(1 - \alpha)$  and  $v(z) = (\delta - z)(1 - \eta - \alpha - z)(1 - \beta - z) + \eta\delta(1 - \beta - z) - \alpha\beta(\delta - z)$ .

The solution of the Equation 5.21 is quite complicated as it involves a third-order polynomial which is irreducible. It would be difficult to obtain the optimal  $\delta$  or  $\alpha$  as before. So we computed numerically the spectral norm of the matrix  $M_p$  in Equation 5.20 for a fixed  $\beta = 0.9$  and  $\delta = 0.99$  (which we used in the experiment in Chapter 4). It is given in Figure 5.8. This result suggests that the optimal  $\alpha$  increases as the  $\eta$  (or  $\eta_h$  as  $h = 1$ ) decreases. Recall our result of *MSGD* in Figure 5.1: the optimal  $\delta$  increases as  $\eta$  decreases. This is consistent with the choice of the learning rate and the momentum rate scheduling in the Nesterov's optimal methods in literature [28]. We have seen that there's an intimate connection between *MSGD* and *EASGD* when we were studying the optimal momentum rate  $\delta$  and the optimal moving rate  $\alpha$ . It suggests that we may also



find interesting rate scheduling for *EASGD*, as well as *EAMSGD* based on the convex analysis. However, we should be careful as the optimal  $\alpha$  for *EASGD* is either zero or negative, while for *EAMSGD* it can be positive as well.

## 5.2 Multiplicative noise

In this section, we study a multiplicative noise model, which attempts to capture the initial behavior of the stochastic optimization method. It is complementary to the additive noise model, which captures the asymptotic behavior.

Our starting point is the following linear regression problem

$$\min_{a \in \mathbb{R}} \mathbb{E}[(v - au)^2],$$

where the expectation  $\mathbb{E}$  is taken over the joint distribution of  $(u, v)$ . If the input data  $(u, v)$  satisfies  $v = a^*u$ , we may reduce the above problem to the following one-dimensional case,

$$\min_{x \in \mathbb{R}} \mathbb{E}[(xu)^2]. \tag{5.22}$$

An interesting perspective of this problem is that we can assume that  $u^2$  of the input data follows a Gamma distribution  $\Gamma(\lambda, \omega)$ , with mean  $\lambda/\omega$  and variance  $\lambda/\omega^2$ . Note that if  $u$  follows a Gaussian distribution with mean zero and variance  $\sigma^2$ , then  $\lambda = 1/2$  and  $\omega = 1/(2\sigma^2)$ .

### 5.2.1 *SGD* with mini-batch

Given  $x_0 \in \mathbb{R}$ , the mini-batch *SGD* method for solving the multiplicative noise problem in Equation 5.22 is as follows,

$$x_{t+1} = x_t - \eta u_t^2 x_t, \tag{5.23}$$

where  $\eta > 0$ , and  $u_t$  is an i.i.d input process.

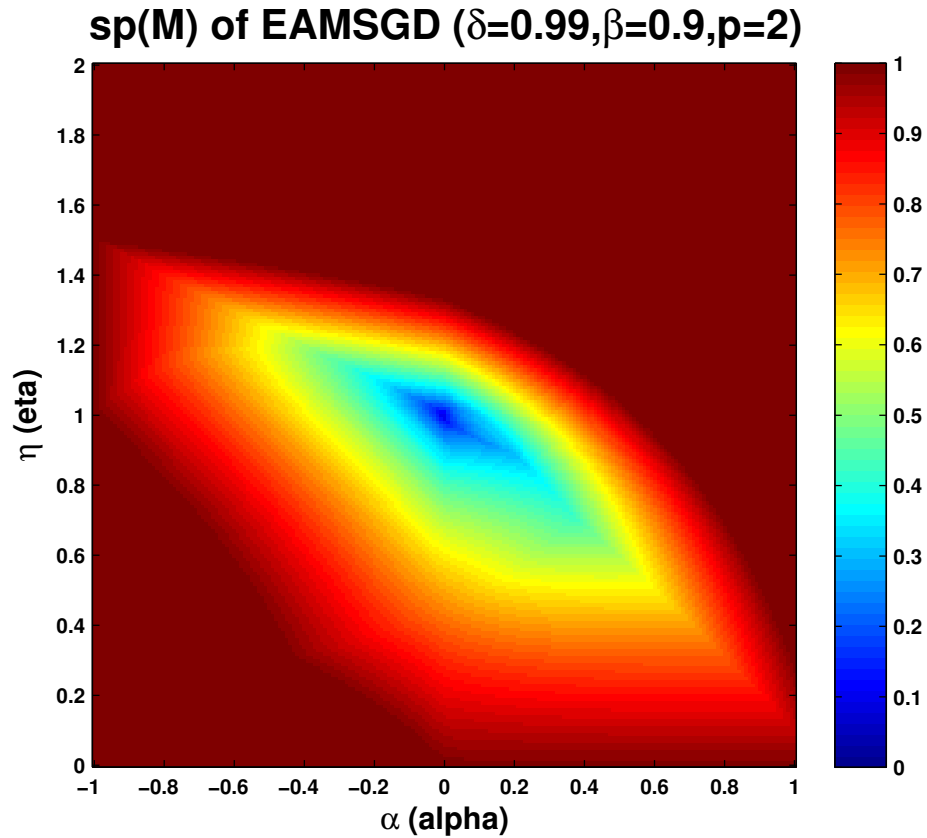


Figure 5.8: The largest absolute eigenvalue of the matrix  $M_p$  in Equation 5.20 as a function of the learning rate  $\eta \in (0, 2)$  and the moving rate  $\alpha \in (-1, 1)$ .  $h = 1$ ,  $\beta = 0.9$  and  $\delta = 0.99$ . Note that we computed this spectrum again using  $p = 2$ , as we have discussed in the text, it is independent of the choice of  $p$  for  $p > 1$ .

As we assumed that  $u_t^2$  follow a Gamma distribution  $\Gamma(\lambda, \omega)$ , it's easy to verify that if we use mini-batch of size  $p$ , Equation 5.23 becomes

$$x_{t+1} = x_t - \eta \frac{1}{p} \sum_{i=1}^p (u_t^i)^2 x_t, \quad (5.24)$$

and this mini-batch  $\frac{1}{p} \sum_{i=1}^p (u_t^i)^2$  also follows a Gamma distribution  $\Gamma(p\lambda, p\omega)$ , since the mean of the mini-batch is not changed and the variance is divided by  $p$ .

We also notice that in the continuous-time limit, i.e. for small  $\eta$ , we can approximate the process in Equation 5.24 by a Geometric Brownian motion [32] as follows,

$$dX(t) = -\frac{\lambda}{\omega} X(t) dt + \sqrt{\eta} \sqrt{\frac{\lambda}{p\omega^2}} X(t) dB(t).$$

An interesting behavior of the geometric Brownian Motion is that its variance can explode (go to infinity), yet the path of the stochastic process still converges almost surely towards zero. In such case, we can observe along the path a few extreme large values. These extreme values may however cause trouble to the stability of the nonlinear dynamics in the training of deep learning models. So our goal here is to control this variance by studying again the second-order moment equation.

Going back to the discrete process defined by Equation 5.24. Denote  $\xi_t = \frac{1}{p} \sum_{i=1}^p (u_t^i)^2$ , we have

$$x_{t+1}^2 = (1 - \eta \xi_t)^2 x_t^2. \quad (5.25)$$

Taking expectation on both sides of the Equation 5.25, we obtain

$$\mathbb{E}[x_{t+1}^2] = \left(1 - 2\eta \frac{p\lambda}{p\omega} + \eta^2 \frac{p\lambda(p\lambda + 1)}{(p\omega)^2}\right) \mathbb{E}[x_t^2] = \left(1 - 2\eta \frac{\lambda}{\omega} + \eta^2 \frac{\lambda(p\lambda + 1)}{p\omega^2}\right) \mathbb{E}[x_t^2]. \quad (5.26)$$

The rate of convergence defined by the above second-moment equation 5.26 is  $1 - 2\eta \frac{\lambda}{\omega} + \eta^2 \frac{\lambda(p\lambda + 1)}{p\omega^2}$ . It is monotone decreasing as  $p$  increases and will saturate at a limit  $1 - 2\eta \frac{\lambda}{\omega} +$

$\eta^2 \frac{\lambda^2}{\omega^2} = (1 - \eta \frac{\lambda}{\omega})^2$ . Given  $p$ , we can optimize this convergence rate with respect to the choice of the learning rate  $\eta$ . This optimal learning rate is given by

$$\eta_p = \frac{p\omega}{p\lambda + 1} = \frac{\omega}{\lambda + 1/p}. \quad (5.27)$$

From this formula, we can see that as  $p$  evolves,  $\eta_p$  will change a lot if  $\lambda$  is much smaller than  $1/p$ . In other words, if  $\lambda$  is very big compared to  $1/p$ , then the change in  $\eta_p$  is not that much as we increases the mini-batch size  $p$ . Thus we can expect that the speedup we would gain depends heavily on the distribution of the input data, in particular the shape parameter  $\lambda$  of the Gamma distribution. Suppose for  $p = 1$ , the  $\xi_t$  follows the Gamma distribution  $\Gamma(\lambda, \omega)$ . Its probability density function is  $\frac{\omega^\lambda}{\Gamma(\omega)} \xi^{\lambda-1} e^{-\omega\xi}$ . We illustrate this probability density function in Figure 5.9 for the standard Gaussian case ( $\lambda = 1/2$ ,  $\omega = 1/2$ ) with  $p = 1$ ,  $p = 2$  and  $p = 4$ . We see that when  $\lambda$  is smaller than one, the probability density function has a pole at zero and has also a slower decay (heavier tail) toward infinity. By increasing the mini-batch size  $p$ , the probability density function becomes more and more concentrated near its mean. Thus the mini-batch is more effective for such input distribution with very large spread (i.e. small  $\lambda$ ).

### 5.2.2 Momentum *SGD*

Given  $x_0 \in \mathbb{R}$  and  $v_0 = 0$ , the *MSGD* method for solving the multiplicative noise problem in Equation 5.22 is as follows,

$$\begin{aligned} v_{t+1} &= \delta v_t - \eta \xi_t (x_t + \delta v_t), \\ x_{t+1} &= x_t + v_{t+1}, \end{aligned} \quad (5.28)$$

where  $\xi_t$  follows a Gamma distribution  $\Gamma(\lambda, \omega)$ .

Denote  $\delta_t = \delta(1 - \eta\xi_t)$  and  $\eta_t = \eta\xi_t$ , the second-order moment equation can be computed

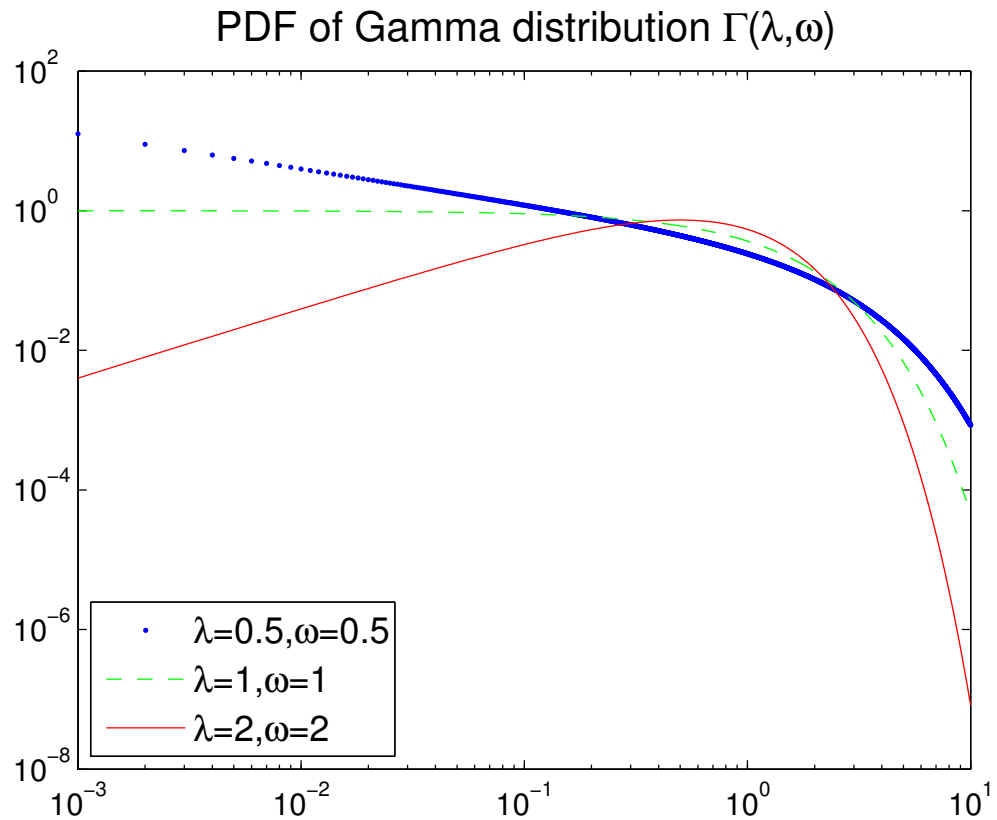


Figure 5.9: The probability density function of the Gamma distribution  $\Gamma(\lambda, \omega)$ . The x-axis and y-axis are both in log-scale to enlarge the singularity at zero and the decay of the tail toward infinity.

as follows,

$$\begin{aligned}
v_{t+1}^2 &= (\delta(1 - \eta\xi_t)v_t - \eta\xi_t x_t)^2 = (\delta_t v_t - \eta_t x_t)^2 \\
&= \delta_t^2 v_t^2 - 2\eta_t \delta_t v_t x_t + \eta_t^2 x_t^2, \\
v_{t+1} x_t &= (\delta_t v_t - \eta_t x_t) x_t = \delta_t v_t x_t - \eta_t x_t^2, \\
x_{t+1}^2 &= (x_t + v_{t+1})^2 = x_t^2 + 2v_{t+1} x_t + v_{t+1}^2 \\
&= x_t^2 + 2(\delta_t v_t x_t - \eta_t x_t^2) + \delta_t^2 v_t^2 - 2\eta_t \delta_t v_t x_t + \eta_t^2 x_t^2, \\
&= \delta_t^2 v_t^2 + (1 - 2\eta_t + \eta_t^2) x_t^2 + (2\delta_t - 2\eta_t \delta_t) v_t x_t, \\
v_{t+1} x_{t+1} &= v_{t+1} (x_t + v_{t+1}) = v_{t+1} x_t + v_{t+1}^2 \\
&= \delta_t v_t x_t - \eta_t x_t^2 + \delta_t^2 v_t^2 - 2\eta_t \delta_t v_t x_t + \eta_t^2 x_t^2 \\
&= \delta_t^2 v_t^2 + (-\eta_t + \eta_t^2) x_t^2 + \delta_t (1 - 2\eta_t) v_t x_t.
\end{aligned} \tag{5.29}$$

Taking expectation on both sides of the Equation 5.29, we obtain the following recursive relation  $(\mathbb{E}v_{t+1}^2, \mathbb{E}x_{t+1}^2, \mathbb{E}v_{t+1}x_{t+1})^T = M(\mathbb{E}v_t^2, \mathbb{E}x_t^2, \mathbb{E}v_t x_t)^T$ , where

$$M = \begin{pmatrix} \delta^2(1 - 2\eta u_1 + \eta^2 u_2) & \eta^2 u_2 & -2\delta\eta(u_1 - \eta u_2) \\ \delta^2(1 - 2\eta u_1 + \eta^2 u_2) & 1 - 2\eta u_1 + \eta^2 u_2 & 2\delta(1 - \eta u_1) - 2\delta\eta(u_1 - \eta u_2) \\ \delta^2(1 - 2\eta u_1 + \eta^2 u_2) & -\eta u_1 + \eta^2 u_2 & \delta(1 - \eta u_1) - 2\delta\eta(u_1 - \eta u_2) \end{pmatrix}, \tag{5.30}$$

$$u_1 = \frac{\lambda}{\omega} \text{ and } u_2 = \frac{\lambda(\lambda+1)}{\omega^2}.$$

Now we can compare the numerical spectral norm of the matrix  $M$  for various choices of the input data distribution parameterized by  $\lambda$  and  $\omega$ . In Figure 5.10, we illustrate the standard Gaussian case, i.e.  $u_t^i$  follows i.i.d.  $\mathcal{N}(0, 1)$  in Equation 5.24. In Figure 5.11 and 5.12 we illustrate the resulting effect if we use mini-batch of size  $p = 2$  and  $p = 4$ . Recall that for  $\delta = 0$ , we have the optimal learning rate in Equation 5.27. It achieves nearly the optimal convergence rate (smallest  $\text{sp}(M)$ ) in these Figures. We also observe that the optimal momentum rate is actually 0 at these optimal learning rates. One can check such observation in Figure 5.13. Thus contrary to our earlier *MSGD* results in the additive noise case (Figure 5.1), the momentum can slow down the optimal convergence

rate. However, we can yet ask, given  $\eta$  and  $\delta$  fixed, what is the region of input data distribution  $\lambda$  and  $\omega$  such that the momentum helps. We find that the answer is for relatively small slope  $\frac{\lambda}{\omega}$  (c.f. Figure 5.14). This phenomenon of acceleration is consistent with our earlier analysis of *MSGD* in the additive noise case (c.f. Figure 5.1).

### 5.2.3 *EASGD* and *EAMSGD*

We now focus on the *EASGD* updates for solving the problem in Equation 5.22,

$$\begin{aligned} x_{t+1}^i &= x_t^i - \eta \xi_t^i x_t^i + \alpha(\tilde{x}_t - x_t^i), \\ \tilde{x}_{t+1} &= \tilde{x}_t - \frac{\beta}{p} \sum_{i=1}^p (\tilde{x}_t - x_t^i), \end{aligned} \tag{5.31}$$

where  $\xi_t^i$  follows a Gamma distribution  $\Gamma(\lambda, \omega)$ .

The second-order moment equation can be computed as follows,

$$\begin{aligned} (x_{t+1}^i)^2 &= (1 - \alpha - \eta \xi_t^i)^2 (x_t^i)^2 + 2\alpha(1 - \alpha - \eta \xi_t^i) \tilde{x}_t x_t^i + \alpha^2 (\tilde{x}_t)^2, \\ (\tilde{x}_{t+1})^2 &= (1 - \beta)^2 (\tilde{x}_t)^2 + 2\beta(1 - \beta) \left( \frac{1}{p} \sum_{i=1}^p \tilde{x}_t x_t^i \right) + \frac{\beta^2}{p^2} \sum_{i=1, j=1}^p x_t^i x_t^j, \\ \tilde{x}_{t+1} x_{t+1}^i &= (1 - \beta)(1 - \alpha - \eta \xi_t^i) \tilde{x}_t x_t^i + \alpha(1 - \beta) (\tilde{x}_t)^2 \\ &\quad + (1 - \alpha - \eta \xi_t^i) \frac{\beta}{p} \sum_{j=1}^p x_t^i x_t^j, \\ x_{t+1}^i x_{t+1}^j &= (1 - \alpha - \eta \xi_t^i)(1 - \alpha - \eta \xi_t^j) x_t^i x_t^j + \alpha^2 (\tilde{x}_t)^2 \\ &\quad + \alpha(1 - \alpha - \eta \xi_t^i) \tilde{x}_t x_t^i + \alpha(1 - \alpha - \eta \xi_t^j) \tilde{x}_t x_t^j, \quad i \neq j. \end{aligned} \tag{5.32}$$

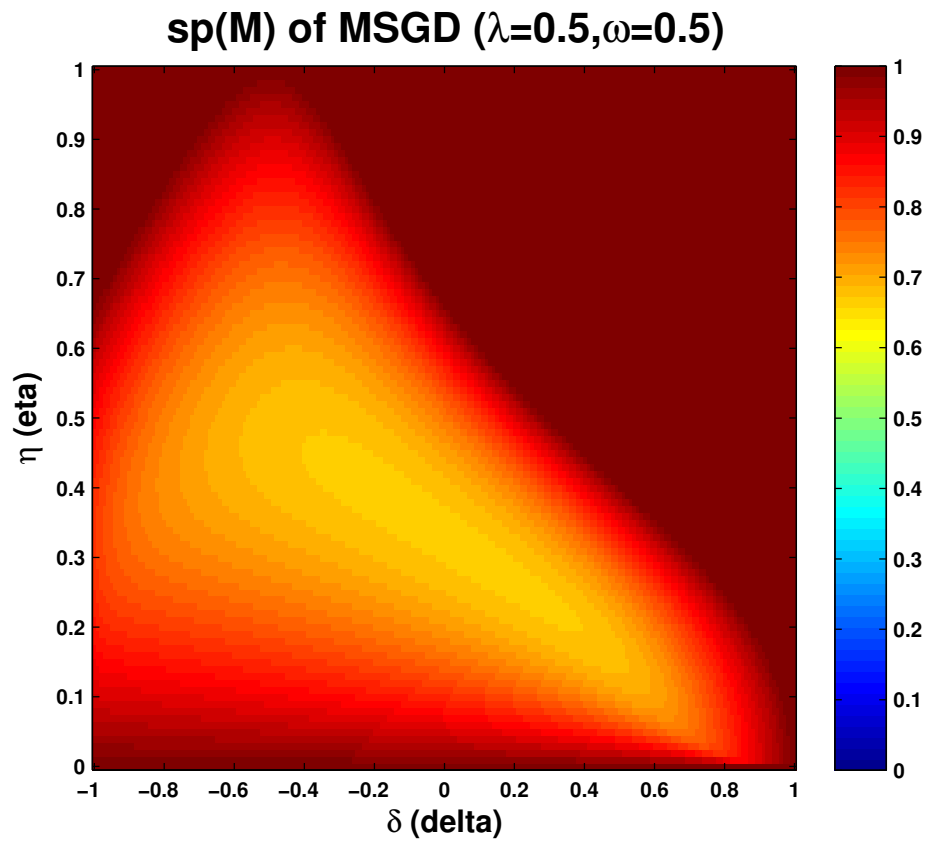


Figure 5.10: The largest absolute eigenvalue of the matrix  $M$  in Equation 5.30 as a function of the learning rate  $\eta \in (0, 1)$  and the momentum rate  $\delta \in (-1, 1)$ .  $\lambda = 0.5$ ,  $\omega = 0.5$ .






Figure 5.11: The largest absolute eigenvalue of the matrix  $M$  in Equation 5.30 as a function of the learning rate  $\eta \in (0, 1)$  and the momentum rate  $\delta \in (-1, 1)$ .  $\lambda = 1$ ,  $\omega = 1$ .




Figure 5.12: The largest absolute eigenvalue of the matrix  $M$  in Equation 5.30 as a function of the learning rate  $\eta \in (0, 1)$  and the momentum rate  $\delta \in (-1, 1)$ .  $\lambda = 2$ ,  $\omega = 2$ .

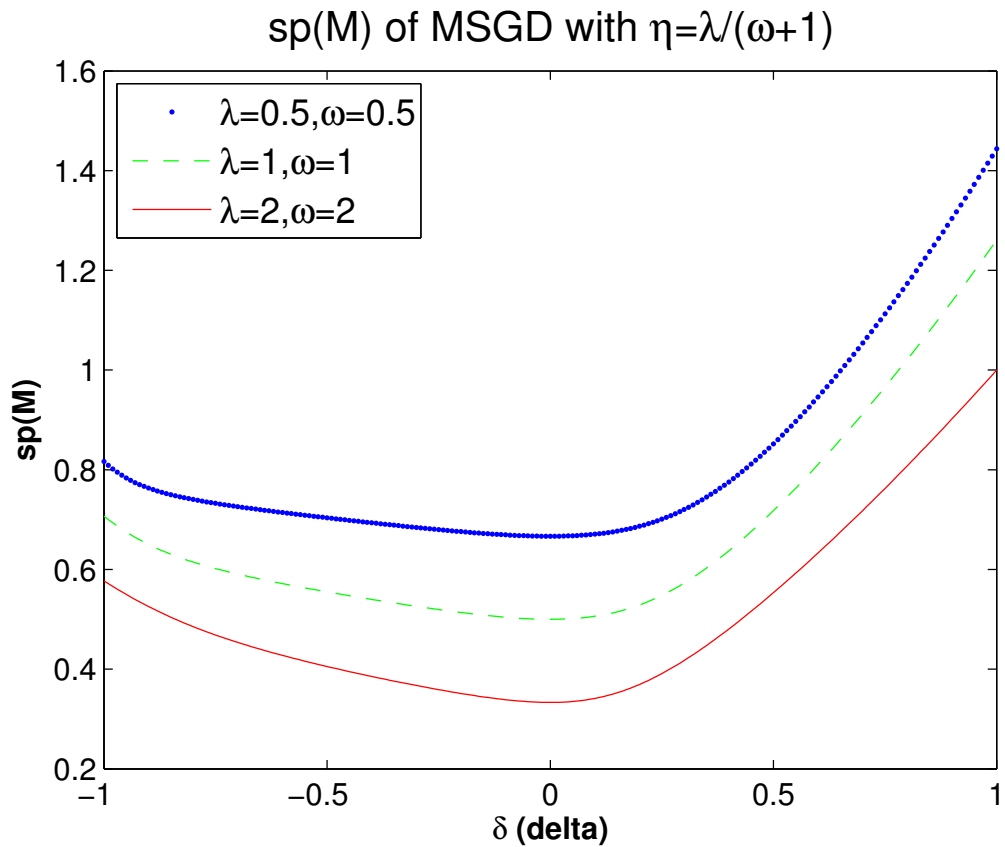


Figure 5.13: The largest absolute eigenvalue of the matrix  $M$  in Equation 5.30 as a function of the momentum rate  $\delta \in (-1, 1)$  at  $\eta = \frac{\lambda}{\omega+1}$ .  $(\lambda, \omega) \in \{(0.5, 0.5), (1, 1), (2, 2)\}$ .




Figure 5.14: The largest absolute eigenvalue of the matrix  $M$  in Equation 5.30 as a function of the input Gamma distribution  $\Gamma(\lambda, \omega)$  in the range  $\omega \in (0, 100)$  and the  $\lambda \in (0, 100)$ .  $(\eta, \delta) \in \{(1, 0), (0.1, 0), (0.1, 0.9)\}$ .

Taking expectation on both sides of the Equation 5.32, we get

$$\begin{aligned}
\mathbb{E}(x_{t+1}^i)^2 &= [(1 - \alpha - \eta \frac{\lambda}{\omega})^2 + \eta^2 \frac{\lambda}{\omega^2}] \mathbb{E}(x_t^i)^2 + 2\alpha(1 - \alpha - \eta \frac{\lambda}{\omega}) \mathbb{E}(\tilde{x}_t x_t^i) + \alpha^2 \mathbb{E}(\tilde{x}_t)^2, \\
\mathbb{E}(\tilde{x}_{t+1})^2 &= (1 - \beta)^2 \mathbb{E}(\tilde{x}_t)^2 + 2\beta(1 - \beta) \frac{1}{p} \sum_{i=1}^p \mathbb{E}(\tilde{x}_t x_t^i) + \frac{\beta^2}{p^2} \sum_{i=1, j=1}^p \mathbb{E}(x_t^i x_t^j), \\
\mathbb{E}(\tilde{x}_{t+1} x_{t+1}^i) &= (1 - \beta)(1 - \alpha - \eta \frac{\lambda}{\omega}) \mathbb{E}(\tilde{x}_t x_t^i) + \alpha(1 - \beta) \mathbb{E}(\tilde{x}_t)^2 \\
&\quad + (1 - \alpha - \eta \frac{\lambda}{\omega}) \frac{\beta}{p} \sum_{j=1}^p \mathbb{E}(x_t^i x_t^j) + \alpha\beta(\frac{1}{p} \sum_{j=1}^p \mathbb{E}(\tilde{x}_t x_t^j)), \\
\mathbb{E}(x_{t+1}^i x_{t+1}^j) &= (1 - \alpha - \eta \frac{\lambda}{\omega})(1 - \alpha - \eta \frac{\lambda}{\omega}) \mathbb{E}(x_t^i x_t^j) + \alpha^2 \mathbb{E}(\tilde{x}_t)^2 \\
&\quad + \alpha(1 - \alpha - \eta \frac{\lambda}{\omega}) \mathbb{E}(\tilde{x}_t x_t^i) + \alpha(1 - \alpha - \eta \frac{\lambda}{\omega}) \mathbb{E}(\tilde{x}_t x_t^j), \quad i \neq j.
\end{aligned} \tag{5.33}$$

We can reduce the above system (Equation 5.33) into a closed form by introducing

$$a_t = \mathbb{E}[(\tilde{x}_t)^2], \quad b_t = \frac{1}{p} \sum_{i=1}^p \mathbb{E}[(x_t^i)^2], \quad c_t = \frac{1}{p} \sum_{i=1}^p \mathbb{E}[\tilde{x}_t x_t^i], \quad d_t = \frac{1}{p^2} \sum_{i=1}^p \sum_{j=1}^p \mathbb{E}[x_t^i x_t^j].$$

We get  $(a_{t+1}, b_{t+1}, c_{t+1}, d_{t+1})^T = M(a_t, b_t, c_t, d_t)^T$ , where

$$M = \begin{pmatrix} (1 - \beta)^2 & 0 & 2\beta(1 - \beta) & \beta^2 \\ \alpha^2 & (1 - \alpha - \eta \frac{\lambda}{\omega})^2 + \eta^2 \frac{\lambda}{\omega^2} & 2\alpha(1 - \alpha - \eta \frac{\lambda}{\omega}) & 0 \\ \alpha(1 - \beta) & 0 & (1 - \beta)(1 - \alpha - \eta \frac{\lambda}{\omega}) + \alpha\beta & (1 - \alpha - \eta \frac{\lambda}{\omega})\beta \\ \alpha^2 & \eta^2 \frac{\lambda}{p\omega^2} & 2\alpha(1 - \alpha - \eta \frac{\lambda}{\omega}) & (1 - \alpha - \eta \frac{\lambda}{\omega})^2 \end{pmatrix}. \tag{5.34}$$

We will analyze the spectral norm of the matrix  $M$  in the limit  $p$  goes to infinity. We focus on the question that whether the stability region (i.e. the spectral norm be smaller than one) with respect to the learning rate  $\eta$  will be enlarged for large  $p$ .

**Case I:**  $\alpha = \beta/p$

The characteristic polynomial of  $M$  in Equation 5.34 can be written as

$$P(z) = [z - (1 - \beta)^2][z - (1 - \beta)(1 - \eta \frac{\lambda}{\omega})][z - (1 - 2\eta \frac{\lambda}{\omega} + \eta^2 \frac{\lambda(\lambda + 1)}{\omega^2})][z - (1 - \eta \frac{\lambda}{\omega})^2] + O_z(\frac{1}{p}),$$

where  $O_z(\frac{1}{p})$  is the remaining polynomial in  $z$  of order  $\frac{1}{p}$  and higher. The four leading order eigenvalues of  $P(z)$  are  $z_1 = (1-\beta)^2$ ,  $z_2 = (1-\beta)(1-\eta\frac{\lambda}{\omega})$ ,  $z_3 = (1-2\eta\frac{\lambda}{\omega} + \eta^2\frac{\lambda(\lambda+1)}{\omega^2})$ ,  $z_4 = (1-\eta\frac{\lambda}{\omega})^2$ . We need necessarily  $\beta \in (0, 2)$ , and in such case the stability condition for  $\eta$  is the same as the *SGD* case (c.f. Equation 5.26), i.e.  $|1 - 2\eta\frac{\lambda}{\omega} + \eta^2\frac{\lambda(\lambda+1)}{\omega^2}| < 1$ . Thus the stability region of *EASGD* remains the same as *SGD* in the limit  $p$  goes to infinity, i.e.

$$0 < \eta < \frac{2\omega}{\lambda + 1}.$$

In the Figure 5.15, 5.16, 5.17 and 5.18, we illustrate the stability region for different  $\lambda$ ,  $\omega$  and  $p$ . Compared with the *MSGD* method in the Figure 5.13, *EASGD* improves its optimal convergence rate:

- $\lambda = 0.5, \omega = 0.5$ :  $\text{sp}(M) = 0.5742$  at  $p = 6$  and  $\eta = 0.3814$  for *EASGD* vs.  $\text{sp}(M) = 0.6667$  for *MSGD* with  $\eta = \frac{\lambda}{\omega+1} = \frac{1}{3}$  and  $\delta = 0$ .
- $\lambda = 1, \omega = 1$ :  $\text{sp}(M) = 0.4317$  at  $p = 7$  and  $\eta = 0.5225$  vs.  $\text{sp}(M) = 0.5$  for *MSGD* with  $\eta = \frac{\lambda}{\omega+1} = \frac{1}{2}$  and  $\delta = 0$ .
- $\lambda = 2, \omega = 2$ :  $\text{sp}(M) = 0.2945$  at  $p = 9$  and  $\eta = 0.6647$  vs.  $\text{sp}(M) = 0.3333$  for *MSGD* with  $\eta = \frac{\lambda}{\omega+1} = \frac{2}{3}$  and  $\delta = 0$ .

Note that *EASGD*'s optimal convergence rate is achieved for some finite  $p$ . This is in contrast to the mini-batch *SGD* method (c.f. Equation 5.27).

### Case II: $\alpha$ is independent of $\beta$ and $p$

The characteristic polynomial of  $M$  of Equation 5.34 can be written as

$$P(z) = (z - z_1)(z - z_2)(z - z_3)(z - z_4) + O_z(\frac{1}{p}),$$

where  $z_1 = (1 - \beta)(1 - \eta\frac{\lambda}{\omega}) - \alpha$ ,  $z_2 = (1 - \alpha)^2 - 2(1 - \alpha)\eta\frac{\lambda}{\omega} + \eta^2\frac{\lambda(\lambda+1)}{\omega^2}$ ,  $z_3 + z_4 = 2 + \alpha^2 - (2 - \beta)\beta - (2 - \eta\frac{\lambda}{\omega})\eta\frac{\lambda}{\omega} - 2\alpha(1 - \beta - \eta\frac{\lambda}{\omega})$ ,  $z_3z_4 = z_1^2$ .




Figure 5.15: The largest absolute eigenvalue of the matrix  $M$  in Equation 5.34 as a function of the learning rate  $\eta \in (0, 1)$  and the number of workers  $p \in [1, 64]$ .  $\lambda = 0.5$ ,  $\omega = 0.5$ ,  $\beta = 0.9$ ,  $\alpha = \beta/p$ .




Figure 5.16: The largest absolute eigenvalue of the matrix  $M$  in Equation 5.34 as a function of the learning rate  $\eta \in (0, 1)$  and the number of workers  $p \in [1, 64]$ .  $\lambda = 1$ ,  $\omega = 1$ ,  $\beta = 0.9$ ,  $\alpha = \beta/p$ .






Figure 5.17: The largest absolute eigenvalue of the matrix  $M$  in Equation 5.34 as a function of the learning rate  $\eta \in (0, 1)$  and the number of workers  $p \in [1, 64]$ .  $\lambda = 2$ ,  $\omega = 2$ ,  $\beta = 0.9$ ,  $\alpha = \beta/p$ .




Figure 5.18: The largest absolute eigenvalue of the matrix  $M$  in Equation 5.34 as a function of the learning rate  $\eta \in (0, 2)$  and the number of workers  $p \in [1, 64]$ .  $\lambda = 10$ ,  $\omega = 10$ ,  $\beta = 0.9$ ,  $\alpha = \beta/p$ . The minimal  $\text{sp}(M) = 0.0868$  is achieved at  $p = 29$  and  $\eta = 0.8929$ .

These four eigenvalues depend on  $\lambda$ ,  $\omega$ ,  $\beta$ ,  $\alpha$ ,  $p$  and  $\eta$ . It's not as clear here on how to choose an optimal  $\alpha$ . We can however still find the stability region with respect of  $\eta$  in the limit  $p \rightarrow \infty$ . We ask given  $\lambda$ ,  $\omega$ ,  $\beta$  and  $\alpha$ , what is the range of  $\eta$  such that the four eigenvalues lie on the unit circle of the complex plane. We also observe that the optimal  $\alpha$  can be positive for large  $p$ , in contrast to the additive noise case (c.f. Figure 5.19 and 5.6). Thus to simplify our analysis, we only consider  $\eta > 0$ ,  $\alpha \in (0, 1)$  and  $\beta \in (0, 1)$ . It turns out that in this case,

- $|z_1| < 1$ :  $0 < \eta < \frac{2-\alpha-\beta}{1-\beta} \frac{\omega}{\lambda}$ .
- $|z_2| < 1$ :  $0 < \eta < \frac{\omega(1-\alpha)}{\lambda+1} + \frac{\omega^2}{\lambda(\lambda+1)} \sqrt{\frac{\lambda^2}{\omega^2} + \frac{\lambda}{\omega^2}(2\alpha - \alpha^2)}$ .
- $|z_3| < 1$  and  $|z_4| < 1$ :  $0 < \eta < \frac{4-2\alpha-2\beta}{2-\beta} \frac{\omega}{\lambda}$ .

We can also check that  $\frac{2-\alpha-\beta}{1-\beta} \frac{\omega}{\lambda} > \frac{4-2\alpha-2\beta}{2-\beta} \frac{\omega}{\lambda}$ , thus we only need to maximize the  $\min_{0 < \alpha < 1} \left\{ \frac{4-2\alpha-2\beta}{2-\beta} \frac{\omega}{\lambda}, \frac{\omega(1-\alpha)}{\lambda+1} + \frac{\omega}{\lambda(\lambda+1)} \sqrt{\lambda^2 + \lambda(2\alpha - \alpha^2)} \right\}$ . We now prove that the maximum is achieved at  $\alpha = 1 - \sqrt{\lambda}$ . In fact, this is the maximum of the second term, and at that  $\alpha$ , the first term equals  $\frac{4-2\alpha-2\beta}{2-\beta} \frac{\omega}{\lambda} = \frac{2+2\sqrt{\lambda}-2\beta}{2-\beta} \frac{\omega}{\lambda}$ , and the second term equals  $\frac{\sqrt{\lambda}}{\lambda+1} (\omega + \frac{\omega}{\lambda}) = \frac{\omega}{\sqrt{\lambda}}$ . The first term being larger than the second term, thus we conclude that the maximal value  $\frac{\omega}{\sqrt{\lambda}}$  is indeed achieved at  $\alpha = 1 - \sqrt{\lambda}$ .

Based on the Equation 5.27, the stability region for *SGD* is  $0 < \eta < \frac{2\omega}{\lambda+1}$ ; for mini-batch *SGD* with  $p \rightarrow \infty$ , it is  $0 < \eta < \frac{2\omega}{\lambda}$ . For *EASGD*, it is  $0 < \eta < \frac{\omega}{\sqrt{\lambda}}$ , with  $p \rightarrow \infty$  and  $\alpha = 1 - \sqrt{\lambda}$ . In case  $\lambda = 0.5$ ,  $\omega = 0.5$ , we can check from Figure 5.19 that  $\alpha = 1 - \sqrt{0.5} = 0.2929$  achieves the widest range of  $\eta \in (0, \sqrt{0.5})$  for  $M$  to be stable.

For the *EAMSGD* method, it's much more difficult to analyze the second-order moment equation. It involves a linear system of nine variables:  $a_t = \mathbb{E}(\tilde{x}_t^2)$ ,  $b_t = \frac{1}{p} \sum_i \mathbb{E}(x_t^i)^2$ ,  $c_t = \frac{1}{p} \sum_i \mathbb{E}(\tilde{x}_t x_t^i)$ ,  $d_t = \frac{1}{p^2} \sum_{i,j} \mathbb{E}(x_t^i x_t^j)$ ,  $e_t = \frac{1}{p} \mathbb{E}(\tilde{x}_t v_t^i)$ ,  $f_t = \frac{1}{p^2} \sum_{i,j} \mathbb{E}(v_t^i v_t^j)$ ,  $g_t = \frac{1}{p} \mathbb{E}(v_t^i)^2$ ,  $h_t = \frac{1}{p} \sum_i \mathbb{E}(x_t^i v_t^i)$ ,  $k_t = \frac{1}{p^2} \sum_{i,j} \mathbb{E}(x_t^i v_t^j)$ . The linear matrix is quite complicated, so we will not present it here.




Figure 5.19: The largest absolute eigenvalue of the matrix  $M$  in Equation 5.34 as a function of the learning rate  $\eta \in (0, 1)$  and the moving rate  $\alpha \in (-1, 1)$ .  $\lambda = 0.5$ ,  $\omega = 0.5$ ,  $\beta = 0.9$ ,  $p = 100$ . The minimal  $\text{sp}(M) = 0.5024$  is achieved at  $\eta = 0.4343$  and  $\alpha = 0.2525$ .

### 5.3 A non-convex case

Here's an amusing non-convex case which sheds some light on when *EASGD* will work and when it will not work. Recall our formalization of the problem in Chapter 1, Equation 1.2:

$$\min_{x^1, \dots, x^p, \tilde{x}} \sum_{i=1}^p \mathbb{E}[f(x^i, \xi^i)] + \frac{\rho}{2} \|x^i - \tilde{x}\|^2,$$

where we refer to  $x^i$ 's as local variables and we refer to  $\tilde{x}$  as a center variable.

If  $f(x^i, \xi^i) = \frac{1}{4}(1 - (x^i)^2)^2$ , which is deterministic (independent of  $\xi^i$ ) and one-dimensional. We see that  $f$  has two minimum  $x^i = 1$  and  $x^i = -1$ . It also has a saddle point at  $x^i = 0$ . We ask how large  $\rho$  should be such that the  $x^i$ 's will have a common minimum. Let's assume that  $p = 2$ , we see that if  $\rho = 0$ , then  $x^1 = -1$  and  $x^2 = 1$  can be a stable solution. This is the scenario that the 'elasticity' of *EASGD* is broken. We have observed a related phenomenon in *EAMSGD* when the communication period is too large ( $\tau$  big) in the Figure 4.13 of Chapter 4. More formally, we have the following objective for  $p = 2$  workers,

$$\min_{x, y, z \in \mathbb{R}} \frac{1}{4}(1 - x^2)^2 + \frac{1}{4}(1 - y^2)^2 + \frac{\rho}{2} \|x - z\|^2 + \frac{\rho}{2} \|y - z\|^2. \quad (5.35)$$

The partial derivative of our objective in Equation 5.35 with respect to  $x$ ,  $y$  and  $z$  is:

$$\begin{aligned} \frac{\partial}{\partial x} &= (x^2 - 1)x + \rho(x - z) \\ \frac{\partial}{\partial y} &= (y^2 - 1)y + \rho(y - z) \\ \frac{\partial}{\partial z} &= \rho(z - x) + \rho(z - y). \end{aligned} \quad (5.36)$$

The question is that what are all the critical points of the objective in Equation 5.35.

By setting the derivative to zero in Equation 5.36, they should satisfy  $z = \frac{x+y}{2}$  and

$$\begin{aligned} (x^2 - 1)x + \rho(x - \frac{x+y}{2}) &= 0 \\ (y^2 - 1)y + \rho(y - \frac{x+y}{2}) &= 0. \end{aligned} \quad (5.37)$$

Observing three special cases:  $x = y = z = 1$ ,  $x = y = z = -1$  and  $x = y = z = 0$ . We can guess that  $x = -y$  is also a solution. In fact, they should satisfy  $z = 0$ ,  $(x^2 - 1)x + \rho(x - 0) = 0$  and  $(y^2 - 1)y + \rho(y - 0) = 0$ . We have thus either  $x = 0$  or  $x^2 = 1 - \rho$  (resp.  $y = 0$  or  $y^2 = 1 - \rho$ ). If  $\rho < 1$ , then we indeed find a real critical point  $x = \sqrt{1 - \rho}$ ,  $y = -\sqrt{1 - \rho}$ ,  $z = 0$ . Is this critical point stable? To answer this, we compute the Hessian of our objective in Equation 5.35 with respect to  $x$ ,  $y$  and  $z$ :

$$H = \begin{bmatrix} 3x^2 - 1 + \rho & 0 & -\rho \\ 0 & 3y^2 - 1 + \rho & -\rho \\ -\rho & -\rho & 2\rho. \end{bmatrix}. \quad (5.38)$$

Evaluating this Hessian at the critical point  $x = \sqrt{1 - \rho}$ ,  $y = -\sqrt{1 - \rho}$ ,  $z = 0$ , we can compute its smallest eigenvalue and see when it is positive definite. Figure 5.20 shows that the smallest eigenvalue is always positive in the range  $\rho \in (0, 2/3)$ . This suggests that this critical point can indeed be stable, i.e. it is a local optimum introduced by *EASGD* when the penalty term  $\rho$  is small enough.

For completeness, we now prove that all the critical points are either of the form  $x = y$  or of the form  $x = -y$ . In fact, adding the two Equations in 5.37 gives us  $(x^2 - 1)x + (y^2 - 1)y = 0$ , i.e.  $x^3 + y^3 = x + y = (x^2 - xy + y^2)(x + y)$ . Subtracting these two Equations gives us  $x^3 - x - y^3 + y + \rho(x - y) = 0$ , i.e.  $x^3 - y^3 = (1 - \rho)(x - y) = (x^2 + xy + y^2)(x - y)$ . If  $x \neq y$  and  $x \neq -y$ , then

$$\begin{aligned} x^2 - xy + y^2 &= 1 \\ x^2 + xy + y^2 &= 1 - \rho. \end{aligned} \quad (5.39)$$

Adding and subtracting again the two Equations in 5.39, we obtain  $x^2 + y^2 = 1 - \frac{\rho}{2}$ , and  $xy = -\frac{\rho}{2}$ . There's no real solution satisfying these two conditions for any  $\rho > 0$ .




Figure 5.20: The smallest eigenvalue of the Hessian matrix  $H$  in Equation 5.38 as a function of the penalty term  $\rho$ , evaluated at the critical point  $x = \sqrt{1-\rho}$ ,  $y = -\sqrt{1-\rho}$ ,  $z = 0$ .

## Chapter 6

# Scaling up Elastic Averaging SGD

This chapter discusses how to scale up the *EASGD* method to hundreds and thousands of processors. In Section 6.1, we first propose a tree-structured extension of the *EASGD* method called *EASGD Tree*. The basic idea and the design principle are discussed in Section 6.1.1, and the numerical results are presented in Section 6.1.2. We present two different communication schemes for the *EASGD Tree* method. As we had seen the advantage of *EAMSGD*, we also accelerate *EASGD Tree* with *Nesterov*'s momentum method. In Section 6.2, we unify the *DONWPOUR* method and the *EASGD* method by considering a *Gauss-Seidel* reformulation of the *EASGD* update rules in the synchronous scenario. This unification suggests the possibility of using both *DONWPOUR* and *EASGD* under the *EASGD Tree*. It also suggests that in-between the *DONWPOUR* and the *EASGD* method there may be some even better method.

### 6.1 *EASGD Tree*

The original motivation of *EASGD Tree* is to run *SGD* at multiple time scales, where each scale corresponds to the use of a different learning rate. It naturally gives rise to a hierarchical tree structured organization of the processors. In literature, the tree



idea has shown up in various contexts. For example, the tree was used to scale up the asynchronous *SGD* method by aggregating the delayed gradients computed by the intermediate nodes and the leaf nodes [2]. It can also be used to efficiently implement the Broadcast/AllReduce operation as in MPI [1, 62]. The benefit of using a tree is that the number of links to connect a very large number of nodes is minimal. The trade-off is that the connectivity of the whole tree is not robust to the link failure. However, the tree structure has its own charm for its simplicity. The main theoretical challenge is to understand its convergence property in terms of the root of the tree. This still remains open.

### 6.1.1 The algorithm

We start from introducing the *EASGD Tree* algorithm. There are two important aspects that can guide us to design such an algorithm, one is the computation, the other is the communication.

- Do we need intermediate node to perform the gradient computation? We had hoped that each tree node runs *SGD* in parallel with a smaller and smaller learning rates as the depth of the tree node decreases (from bottom to top). Each intermediate node (i.e. except the leaf node) is expected to average out the noise of its children. The root can thus achieve the smallest variance. One benefit that each node performs its own gradient computation is that it can reduce the bias of the spatial averaging of its children away from the local optimum. The disadvantage is that each node will have the overhead for both computation and communication. In case each tree node uses only one CPU core, it needs to tradeoff the computation speed for the communication throughput. In addition, in the HPC environment (where we run our experiments), we have not observed any visible performance improvement for the intermediate node to perform the gradient computation. It's also complicated to choose a proper learning rate decay for different levels of the

tree nodes. On the other hand, if we were to remove the intermediate node’s gradient computation, we can still preserve the multi-scale variance reduction property using the elastic averaging. Moreover, the root node of the tree behaves as the center variable of the *EASGD* method. Thus in the following we only discuss the *EASGD Tree* method without intermediate node’s gradient computation.

- There are at least two different ways to design the communication protocol. One is locally synchronous, the other is fully asynchronous. We observe that the local synchronization protocol (as in our asynchronous *EASGD* in Section 2.2) may cause extra waiting time for the parent and the children nodes, but it can give better convergence property (due to a smaller parameter staleness). The fully asynchronous protocol can in theory hide completely this waiting time by allowing each node asynchronously broadcasts its parameter to its neighbours without blocking any computation. In the scale that we simulate the *EASGD Tree*, we have to use the fully asynchronous protocol as it saves significantly the time which is spent on the chain of waiting. In our HPC environment, the fully asynchronous protocol works well if the network link is stable over time.
- The communication cost is different within a machine and across a machine. We would thus structure the tree nodes from bottom-up, i.e. the leaf nodes are allocated on one machine so that their variance can be averaged out the quickest possible. The intermediate nodes including the root of the tree will then have to communicate across the machine boundary (in our HPC environment, there are 20 physical CPU cores per machine). These communication should happen less frequently. We will thus introduce two levels of the communication periods. The first communication period is between the leaf nodes and their parents. The second (smaller) one is between the intermediate nodes. We will also distinguish the push up (toward parent) period and the push down (toward children) period. With this, we can trade-off more push up communication for less push down communication.

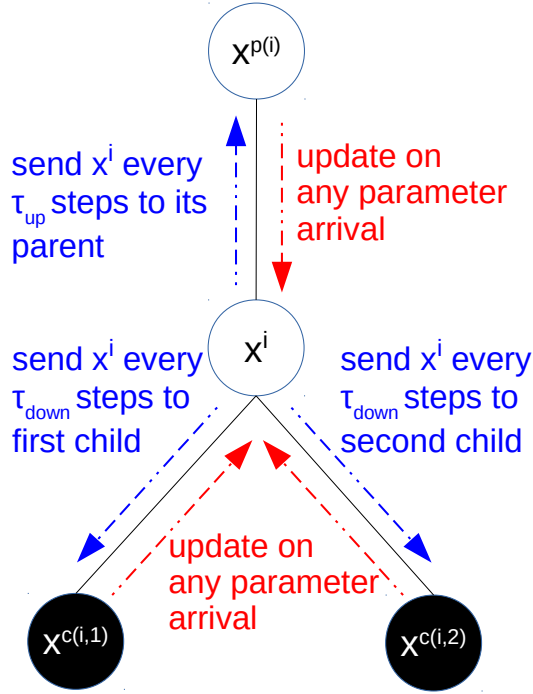


Figure 6.1: The behavior of *EASGD Tree* in Algorithm 6.

Based on the above discussion, we now present the *EASGD Tree* in Algorithm 6. Figure 6.1 illustrates this. This generic pseudocode does not distinguish the leaf node, the intermediate node and the root node, so let's describe their difference first. As in *EASGD*, each node starts with the same initial parameter  $x_0$  and sets its local clock  $t^i$  to 0. It performs non-blocking read (Irecv) and non-blocking send (Isend) with its parent and children (if any). Every  $\tau_{up}$  local steps, it sends its parameter to its parent. Every  $\tau_{down}$  local steps (number of local gradient updates), it sends its parameter to the  $d$  children. We denote  $d$  to be the degree of the  $d$ -ary tree we consider here. Each node also needs to check if there's any new arrival of the parameter from its parent or its  $d$  children. This needs to be interleaved with the gradient computation to achieve high I/O throughput. On receiving a new parameter, a moving average with the moving rate  $\alpha$  is performed. If it is the leaf node, then a stochastic gradient descent step with learning rate  $\eta$  is performed per local step; otherwise it replaces the gradient descent step with a while loop, looping for roughly the same amount of time as a gradient step, just to apply

the moving average on any new parameter arrival. Note that for each node, they can use different communication period  $\tau_{up}$  and  $\tau_{down}$ . We discuss two such possible schemes below. They are illustrated in Figure 6.2.

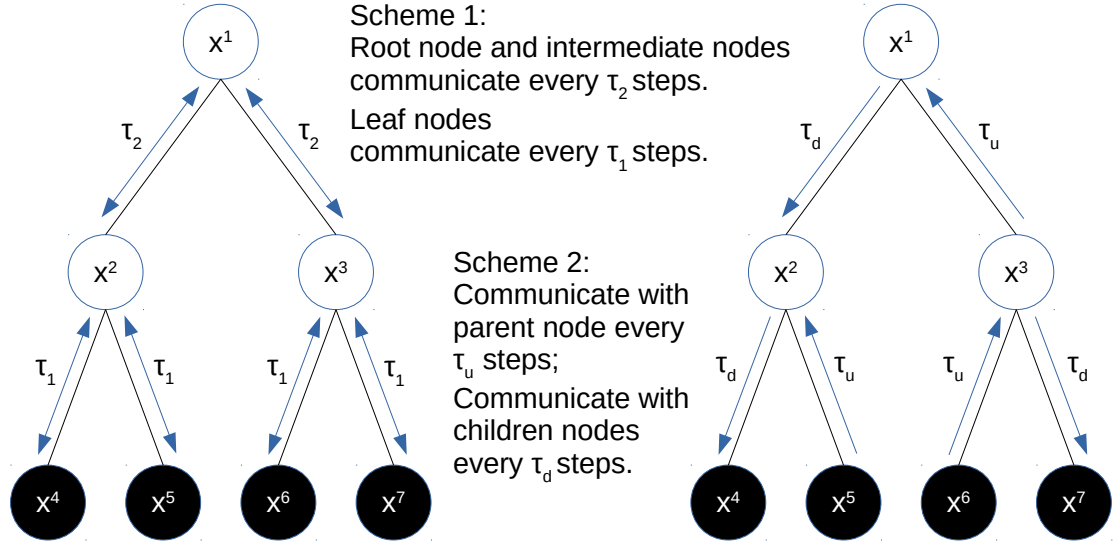


Figure 6.2: The two communication schemes of the *EASGD Tree*.

The first scheme is based on the idea of multi-scale averaging. We would like to have a fast averaging process at the bottom level, and a slower averaging process at the top and intermediate level. Let  $\tau_1$  be the fast communication period between the leaf nodes and their parents. Let  $\tau_2$  be the slow communication period between the intermediate nodes. For the leaf nodes, i.e. without child, we set  $\tau_{up} = \tau_1$ ,  $\tau_{down} = NaN$ . For their parents (which should sit on the same machine as their children), we set  $\tau_{up} = \tau_2$ ,  $\tau_{down} = \tau_1$ . For the rest of the intermediate nodes, except for the root node, we set  $\tau_{up} = \tau_{down} = \tau_2$ . Lastly, for the root node, we set  $\tau_{up} = NaN$ ,  $\tau_{down} = \tau_2$ .

The second scheme is to mimic the behavior of *EASGD* for  $\beta$  being large and  $\alpha$  being small (c.f. Equation 2.3 and 2.4). We use a large  $\tau_{up}$  to represent large  $\beta$ , and a small  $\tau_{down}$  to represent small  $\alpha$ . For simplicity, we set each tree node with the same  $\tau_{up} = \tau_u$  and  $\tau_{down} = \tau_d$ , except for the leaf node with  $\tau_{down} = NaN$  and the root node with  $\tau_{up} = NaN$ .

There’s a subtle difference between *EASGD* and *EASGD Tree*. As we discussed in Section 2.2, the design of *EASGD* follows the *Jacobi* form. In *EASGD Tree*, we use instead the *Gauss-Seidel* form to perform the moving average. This is because the time when the parameter will arrive is not predictable in general, so we’d rather perform the update just-in-time. Nevertheless, we should be careful to avoid performing the moving average during the gradient update. Moreover, there’s a deeper connection between the *Jacobi/Gauss-Seidel* method and *EASGD/DOWNPOUR* method that we shall discuss in the next Section 6.2. Thus *EASGD Tree* is somewhere in-between *EASGD* and *DOWNPOUR*.

<p><b>Algorithm 6:</b> <i>EASGD Tree</i>          Processing by worker <math>i</math>, with its parent <math>p(i)</math> and children <math>\{c(i, j)   j = 1, \dots, d\}</math>.</p> <hr/> <p><b>Input:</b> learning rate <math>\eta</math>, moving rate <math>\alpha</math>, communication period <math>\tau_{up}</math> and <math>\tau_{down} \in \mathbb{N}</math>  <b>Initialize:</b> <math>x^i = x_0, t^i = 0</math></p> <hr/> <p><b>Irecv</b> <math>x^{p(i)}</math> from parent <math>p(i)</math>  <b>Irecv</b> <math>x^{c(i, j)}</math> from children <math>\{c(i, j)   j = 1, \dots, d\}</math>  <b>Repeat</b>            <b>if</b> (<math>\tau_{up}</math> divides <math>t^i</math>) <b>then</b>              <b>Isend</b> <math>x^i</math> to parent <math>p(i)</math>            <b>end</b>            <b>if</b> (<math>\tau_{down}</math> divides <math>t^i</math>) <b>then</b>              <b>Isend</b> <math>x^i</math> to children <math>\{c(i, j)   j = 1, \dots, d\}</math>            <b>end</b>            <b>if</b> there is new arrival of <math>x^{p(i)}</math> <b>then</b>              <math>x^i \leftarrow x^i + \alpha(x^{p(i)} - x^i)</math>            <b>end</b>            <b>if</b> there is new arrival of <math>x^{c(i, j)}</math> <b>then</b>              <math>x^i \leftarrow x^i + \alpha(x^{c(i, j)} - x^i)</math>            <b>end</b>            <math>x^i \leftarrow x^i - \eta \nabla f(x^i, \xi_{t^i}^i)</math>            <math>t^i \leftarrow t^i + 1</math>  <b>Until forever</b></p>
--

### 6.1.2 The result

In this section, we show empirically the advantage and the challenge to scale up the *EASGD Tree* to a few hundreds of CPU cores. Since we are limited by the number

GPUs available, all the results are run on CPUs. The experiment setup is thus different to our earlier results in Section 4.1. We report the results based on CIFAR-10 dataset using a low-rank convolutional neural-network model [56]. The low-rank convolution approximation saves us a lot of CPU computation time, so we can simulate the long term behavior of the algorithm within 12 hours. We do not use any data augmentation as we did in Chapter 4, because it requires extra CPUs to pre-process the input data. Each data point is sampled uniformly random with replacement (different to what we did in Section 4.1). As in [56], we center the input pixel values by the mean and the standard deviation for each of the three channels using the training data. We also add  $l_2$ -regularization  $\frac{\lambda}{2} \|x\|^2$  to the loss function with  $\lambda = 0.0001$ .

For consistency, we describe the detailed model<sup>1</sup> using the notation in Section 4.1:

$$\begin{aligned}
(3, 32, 32) &\xrightarrow[(1,5,1,1,0,2)]{C} (10, 32, 32) \xrightarrow[(5,1,1,1,2,0)]{C,R} (96, 32, 32) \xrightarrow[(2,2,2,2)]{P} (96, 16, 16) \xrightarrow[(1,5,1,1,0,2)]{C} \\
(51, 16, 16) &\xrightarrow[(5,1,1,1,2,0)]{C,R} (128, 16, 16) \xrightarrow[(2,2,2,2)]{P} (128, 8, 8) \xrightarrow[(1,5,1,1,0,2)]{C} (51, 8, 8) \xrightarrow[(5,1,1,1,2,0)]{C,R} \\
(256, 8, 8) &\xrightarrow[(2,2,2,2)]{P} (256, 4, 4) \xrightarrow[(1,1,1,1,0,0)]{C,R} (64, 4, 4) \xrightarrow[0.5]{D,L,R} (256, 1, 1) \xrightarrow{D,L,S} (10, 1, 1).
\end{aligned}$$

In the following experiments, we call *CIFAR-lowrank* the experiment that is just described. We use a  $d$ -ary tree with degree  $d = 16$  and  $p = 256$  leaf nodes running on 16 machines. For simplicity, we set the moving rate  $\alpha$  at each node to be a constant  $0.9/(16 + 1)$ . First we examine the two communication schemes that we have discussed above. Figure 6.3 shows the performance of the first communication scheme using period  $\tau_1 = 10$  and  $\tau_2 = 100$ . Figure 6.4 shows the performance of the second communication scheme using  $\tau_u = 8$  and  $\tau_d = 80$ . We see that the first communication scheme achieves a faster convergence in terms of the training loss. All the six runs are nearly the same at the beginning. However, five runs out of the six have diverged in the middle of the training. We have chosen a quite large learning rate which can trigger this instability. The second communication scheme is more stable. Only one out of the six runs has diverged. It has a slower convergence in terms of the training loss, but it consistently

<sup>1</sup>see also in [https://github.com/chengtaipu/lowrankcnn/blob/master/cifar/models/baseline\\_lowrank\\_cudnn.lua](https://github.com/chengtaipu/lowrankcnn/blob/master/cifar/models/baseline_lowrank_cudnn.lua)

achieves a lower test error (the smallest one is 15.32%, compared to the smallest test error of the first scheme which is 16.86%). Also notice that there are two runs with a slower initial convergence rate. This may be related to the network traffic.

Next we examine the effect of the momentum when using a mini-batch size of 16. We use mini-batch as we have observed that the momentum can help more when the stochastic gradient variance is smaller (c.f. also Figure 5.14). Moreover, it would be nice to see whether using a relatively small mini-batch size can still hurt the test performance. As in *EAMSGD*, we can apply Nesterov’s momentum to each of the leaf node during the gradient computation. Compared to the earlier results (without using mini-batch) in Figure 6.3 and 6.4, we increase the total training time from 3 hours to 12 hours. This is because we do not have more CPU resources to parallelize the mini-batch computation. It takes 0.13sec/step in the minibatch case vs. 0.01sec/step in the case without using mini-batch.

In Figure 6.5, 6.6 and 6.7, we report the results of the momentum in the first communication scheme. We see that without using momentum (i.e.  $\delta = 0$ ), the training process is still not very stable as before. Three out of the six runs have diverged. Using momentum  $\delta = 0.9$  allows us to reduce the learning rate by a factor of ten (see also the discussion of Figure 5.14), and gives a much more stable training process. Also they look almost identical, except for the two curves marked in e and f. This difference is due to a longer initialization of all the nodes. Using momentum  $\delta = 0.99$  is also very stable, though the results are more varied.

In Figure 6.8, 6.9 and 6.10, we report the results of the momentum in the second communication scheme. Without using momentum is still not as stable. The curve a and c have stopped in the middle without any indication of divergence. But they did have diverged and ended with *NaN*. Nevertheless, the curve b gives a smallest test error 15.4%. The second communication scheme again leads to better test performance. Using momentum  $\delta = 0.9$  gives a smallest test error 14.91% in curve b. Using momentum  $\delta = 0.99$  also

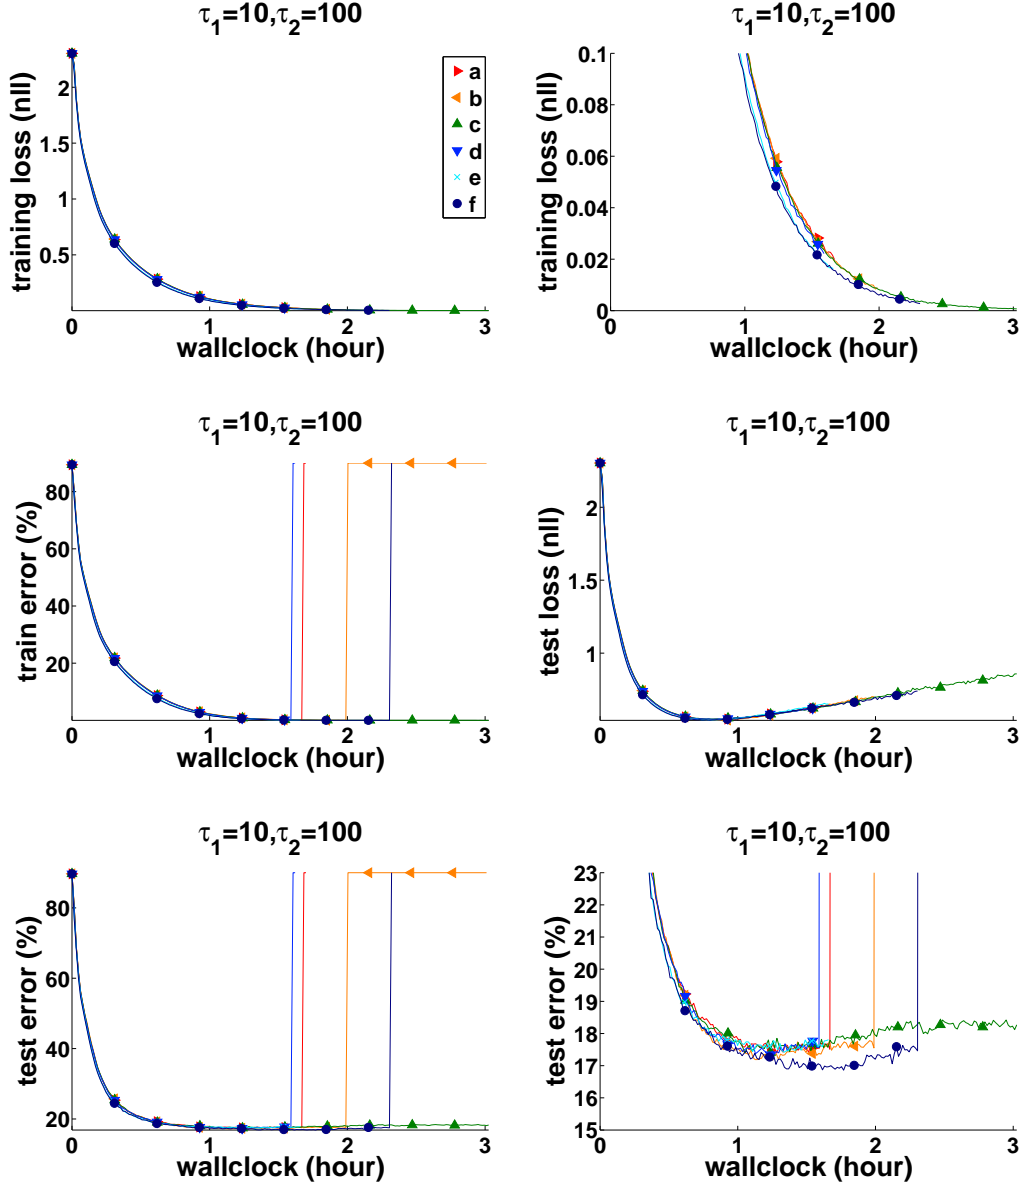


Figure 6.3: *EASGD Tree* on *CIFAR-lowrank* using the first communication scheme with  $\tau_1 = 10$  and  $\tau_2 = 100$ . Training loss and error, test loss and error of the root node versus a wallclock time. We run this experiment six times independently (labeled with a,b,c,d,e,f) with a same random initialization.  $p = 256, d = 16, \eta = 5e - 3, \alpha = 0.9/(d + 1)$ .



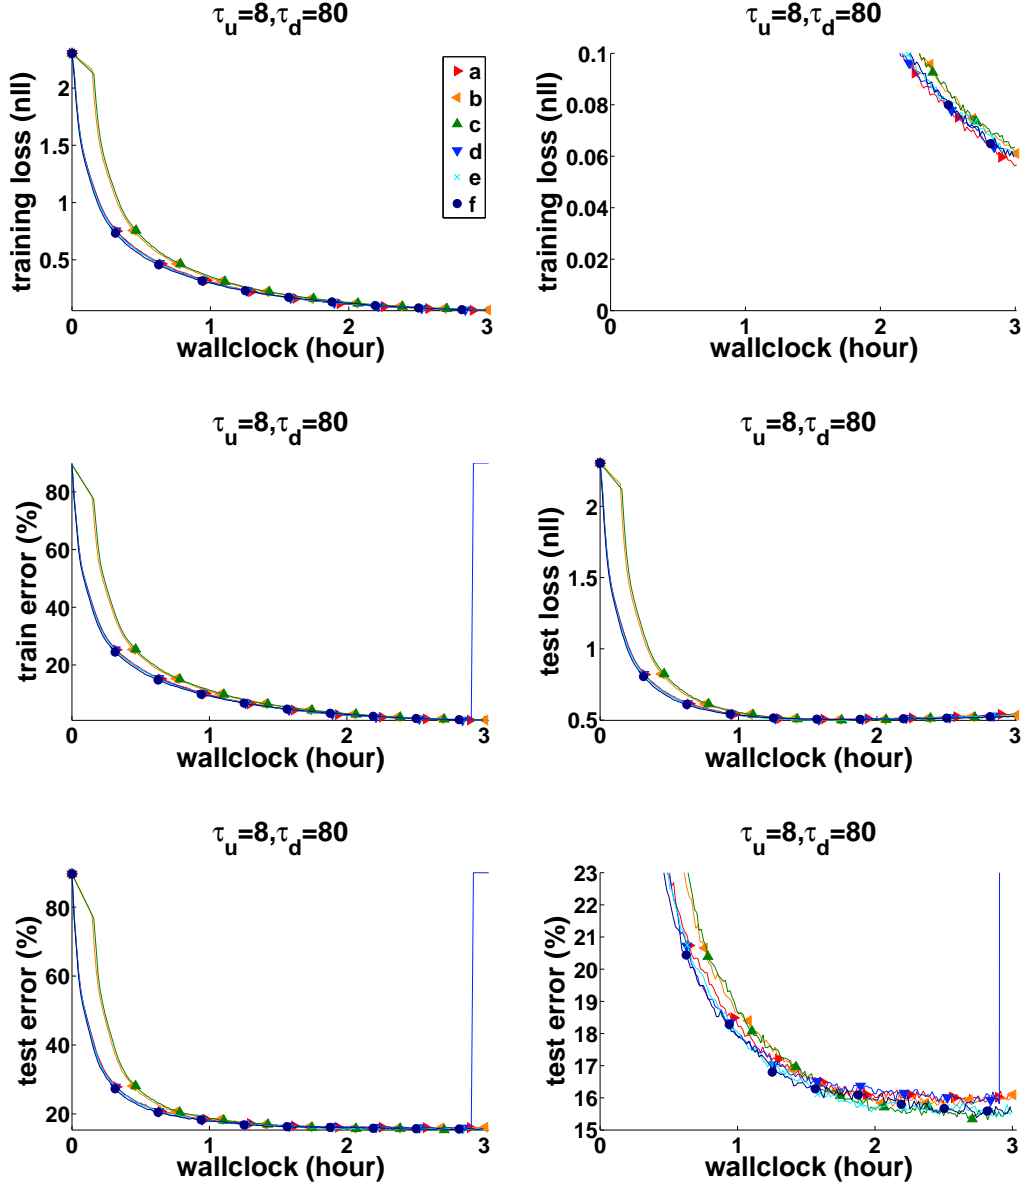


Figure 6.4: *EASGD Tree* on *CIFAR-lowrank* using the second communication scheme with  $\tau_u = 8$  and  $\tau_d = 80$ . Training loss and error, test loss and error of the root node versus a wallclock time. We run this experiment six times independently (labeled with a,b,c,d,e,f) with a same random initialization.  $p = 256, d = 16, \eta = 5e-3, \alpha = 0.9/(d+1)$ .

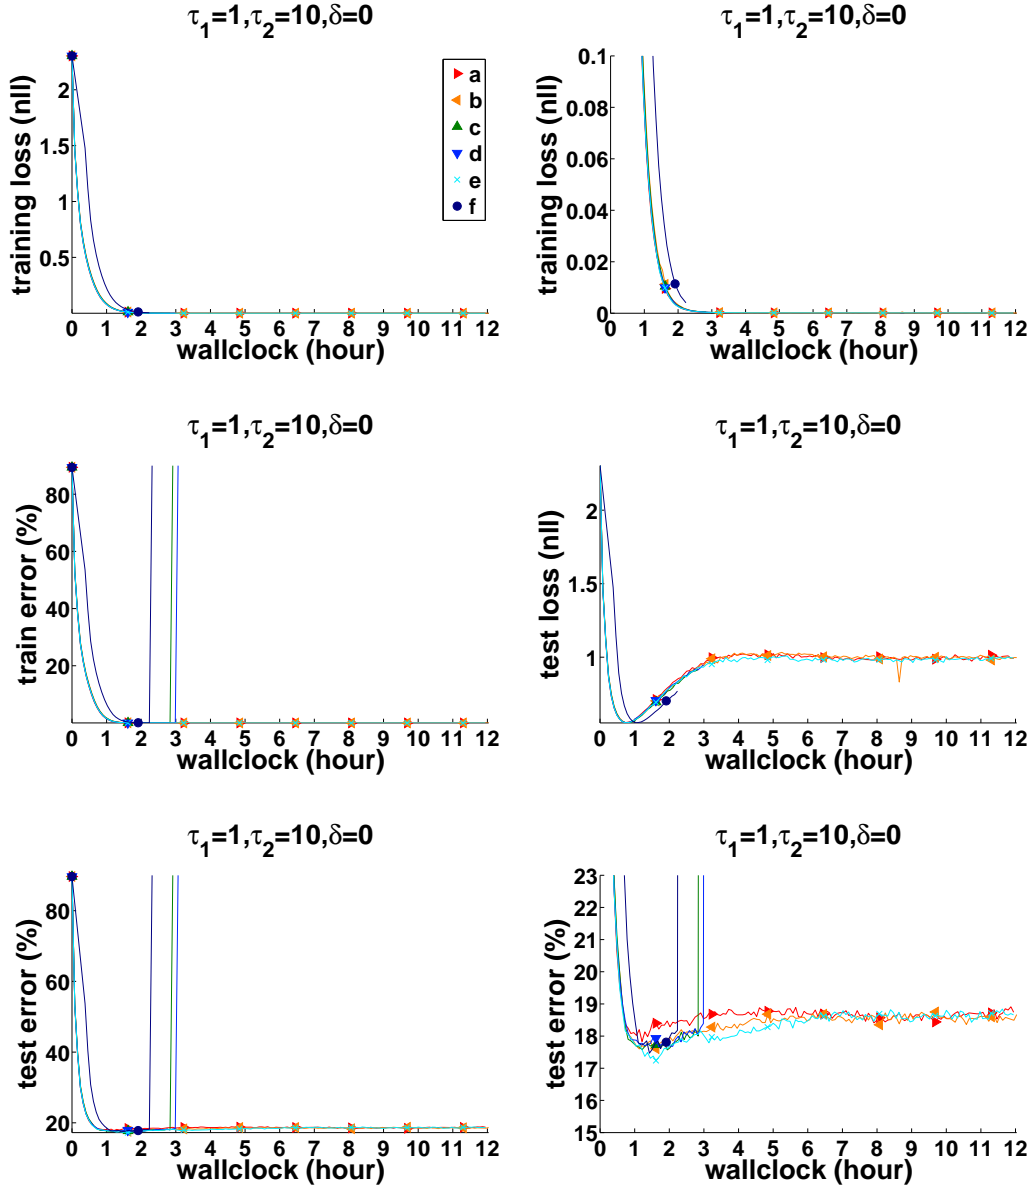


Figure 6.5: *EASGD Tree* on *CIFAR-lowrank* using the first communication scheme with  $\tau_1 = 1$  and  $\tau_2 = 10$ . Training loss and error, test loss and error of the root node versus a wallclock time. We run this experiment six times independently (labeled with a,b,c,d,e,f) with a same random initialization.  $p = 256, d = 16, \eta = 5e - 2, \alpha = 0.9/(d + 1)$ . The momentum rate  $\delta = 0$ .

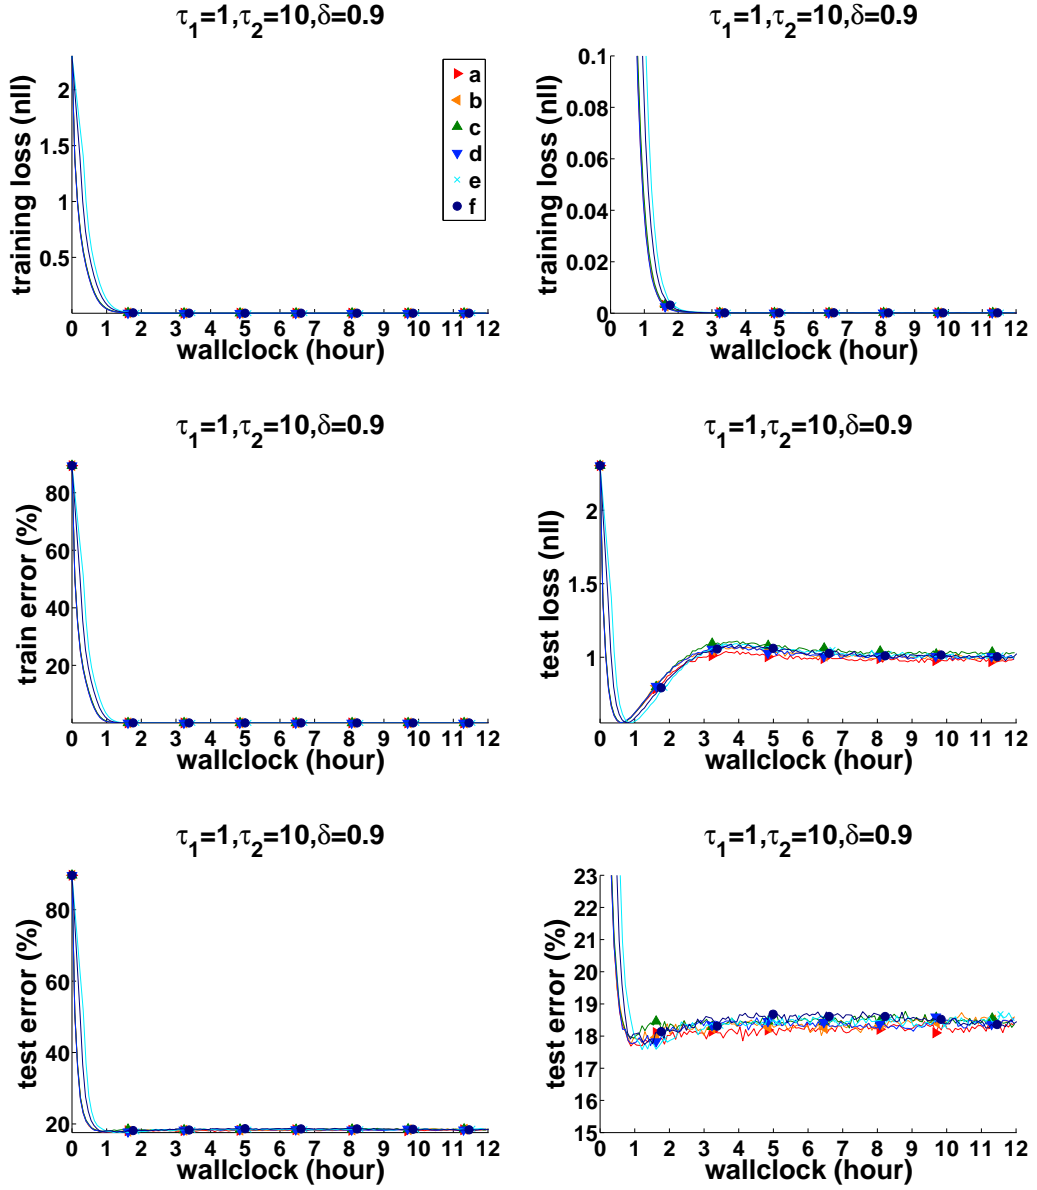


Figure 6.6: *EASGD Tree* on *CIFAR-lowrank* using the first communication scheme with  $\tau_1 = 1$  and  $\tau_2 = 10$ . Training loss and error, test loss and error of the root node versus a wallclock time. We run this experiment six times independently (labeled with a,b,c,d,e,f) with a same random initialization.  $p = 256, d = 16, \eta = 5e - 3, \alpha = 0.9/(d + 1)$ . The momentum rate  $\delta = 0.9$ .

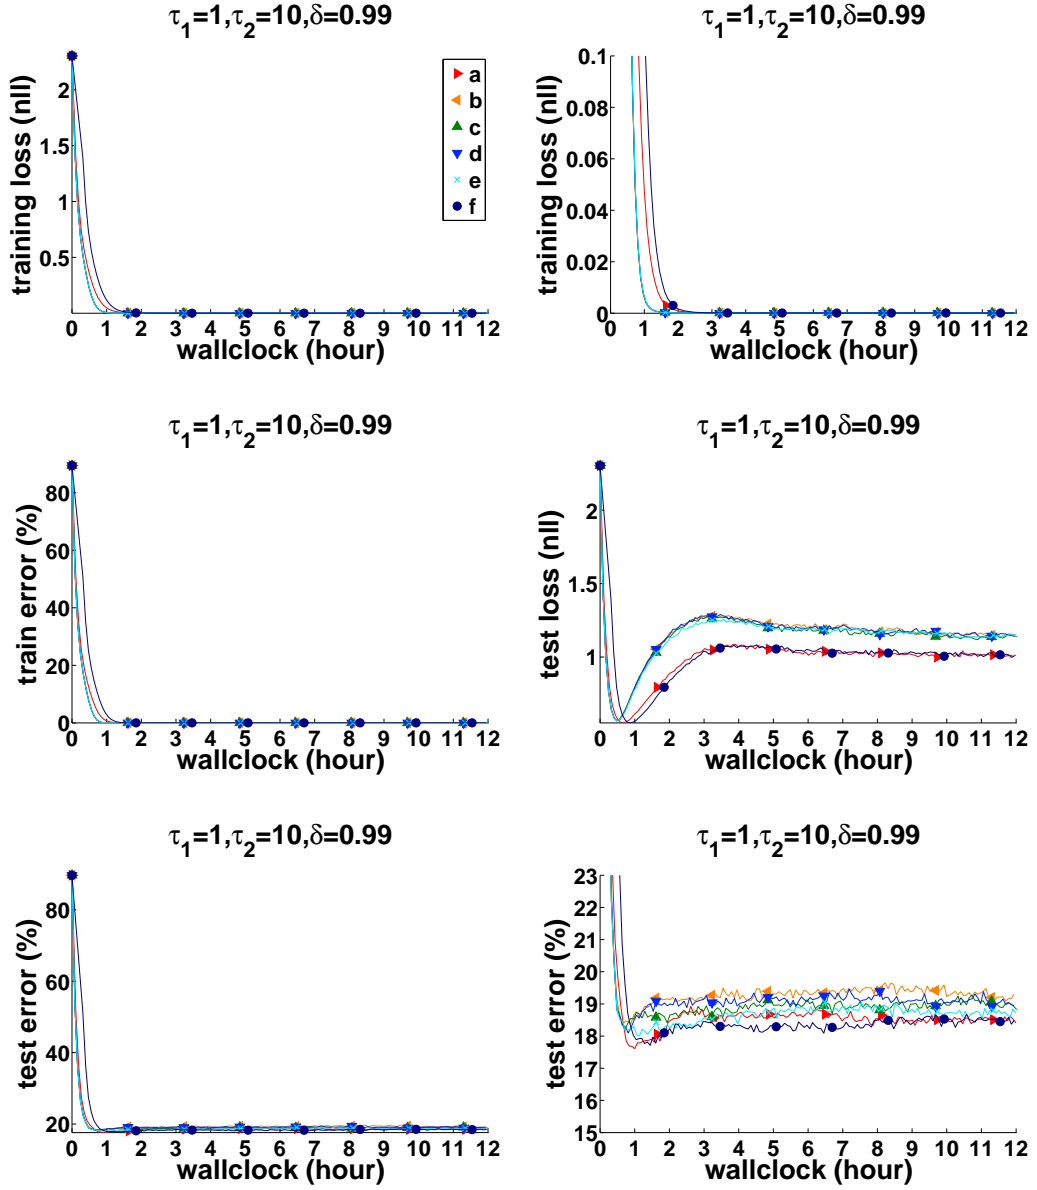


Figure 6.7: *EASGD Tree* on *CIFAR-lowrank* using the first communication scheme with  $\tau_1 = 1$  and  $\tau_2 = 10$ . Training loss and error, test loss and error of the root node versus a wallclock time. We run this experiment six times independently (labeled with a,b,c,d,e,f) with a same random initialization.  $p = 256, d = 16, \eta = 5e - 4, \alpha = 0.9/(d + 1)$ . The momentum rate  $\delta = 0.99$ .

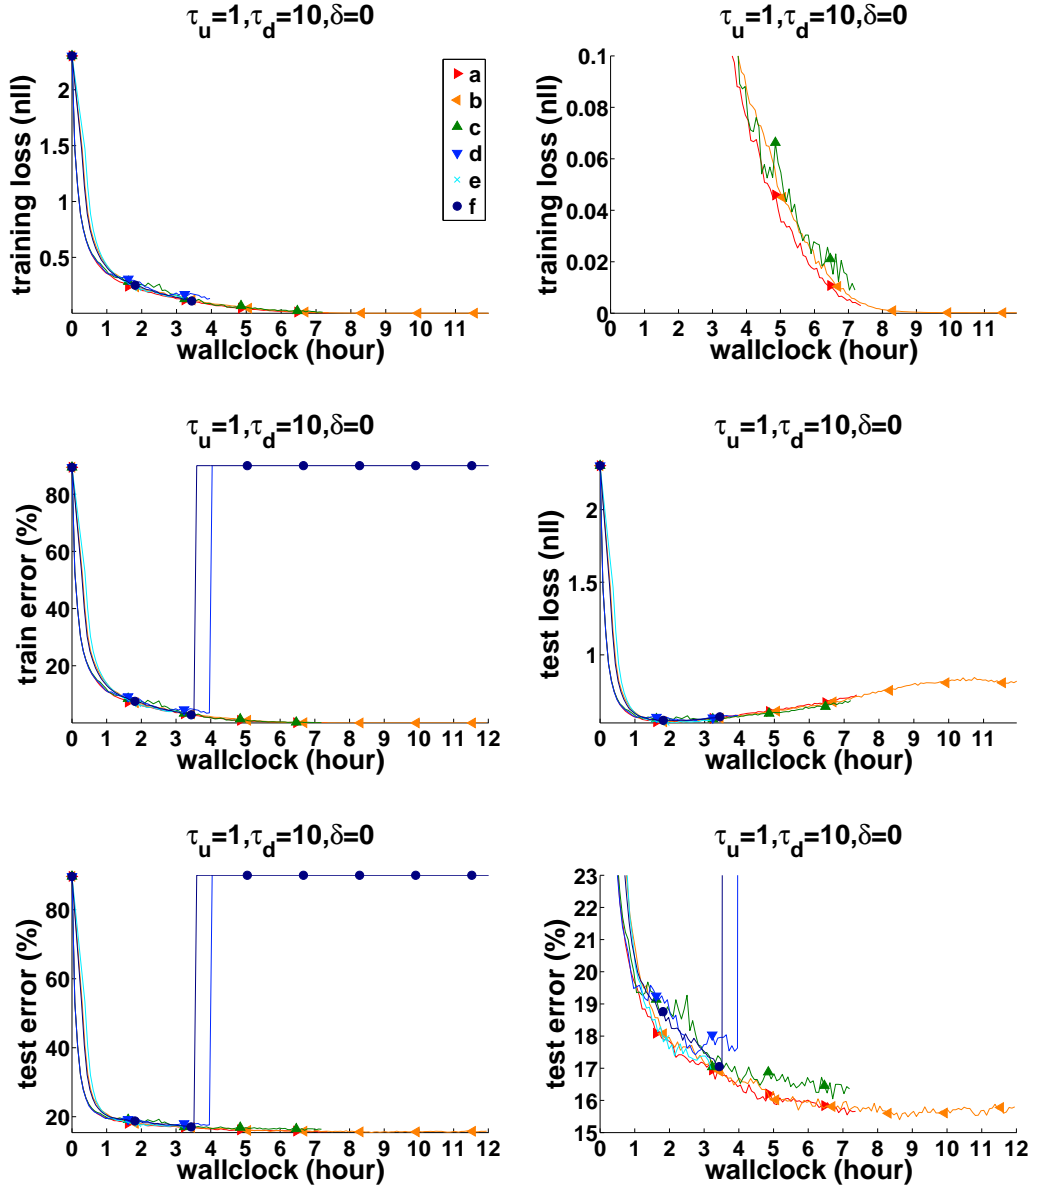


Figure 6.8: *EASGD Tree* on *CIFAR-lowrank* using the second communication scheme with  $\tau_u = 1$  and  $\tau_d = 10$ . Training loss and error, test loss and error of the root node versus a wallclock time. We run this experiment six times independently (labeled with a,b,c,d,e,f) with a same random initialization.  $p = 256, d = 16, \eta = 5e-2, \alpha = 0.9/(d+1)$ . The momentum rate  $\delta = 0$ .




Figure 6.9: *EASGD Tree* on *CIFAR-lowrank* using the second communication scheme with  $\tau_u = 1$  and  $\tau_d = 10$ . Training loss and error, test loss and error of the root node versus a wallclock time. We run this experiment six times independently (labeled with a,b,c,d,e,f) with a same random initialization.  $p = 256, d = 16, \eta = 5e-3, \alpha = 0.9/(d+1)$ . The momentum rate  $\delta = 0.9$ .

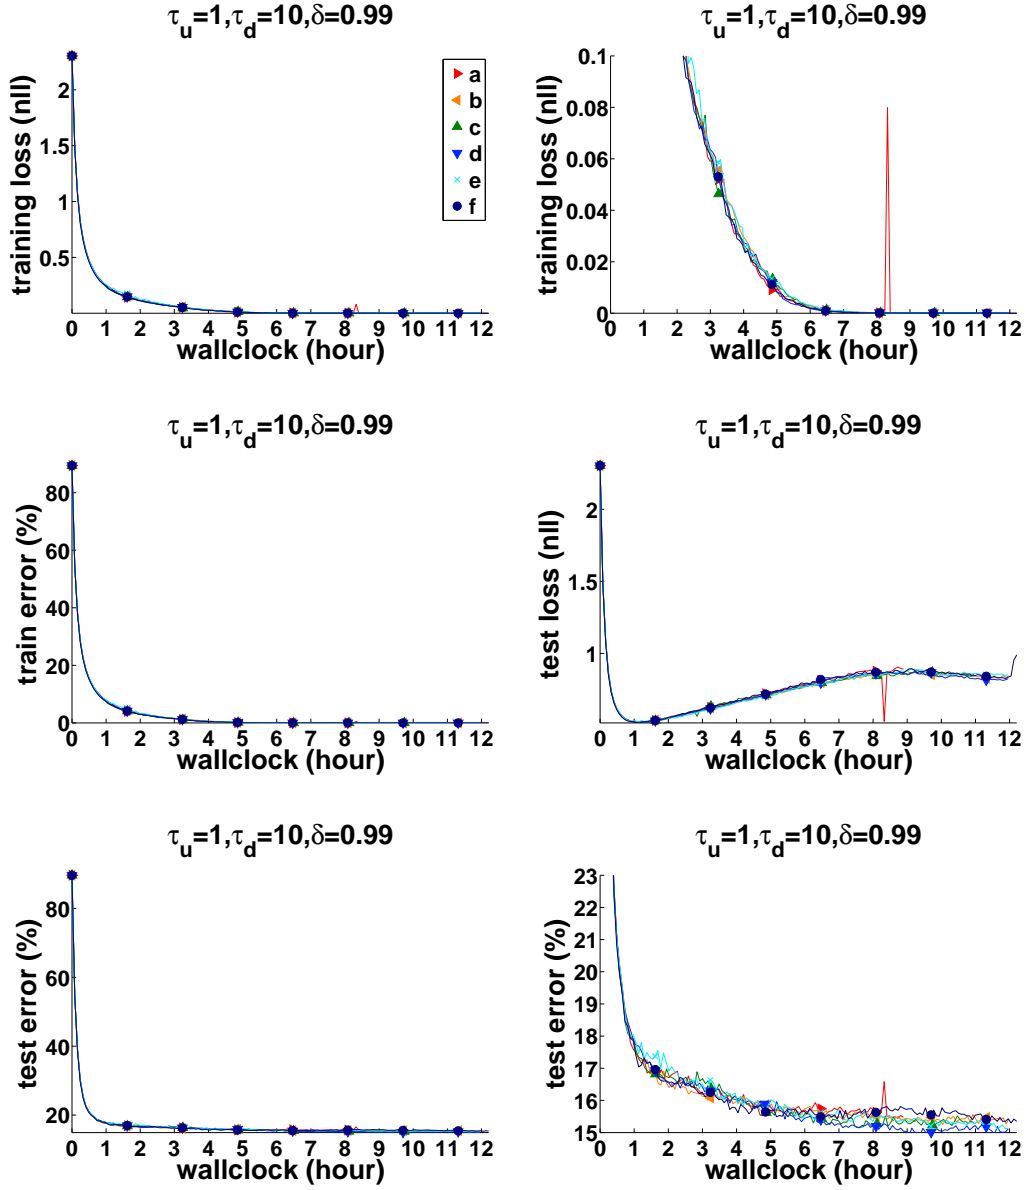


Figure 6.10: *EASGD Tree* on *CIFAR-lowrank* using the second communication scheme with  $\tau_u = 1$  and  $\tau_d = 10$ . Training loss and error, test loss and error of the root node versus a wallclock time. We run this experiment six times independently (labeled with a,b,c,d,e,f) with a same random initialization.  $p = 256, d = 16, \eta = 5e-4, \alpha = 0.9/(d+1)$ . The momentum rate  $\delta = 0.99$ .

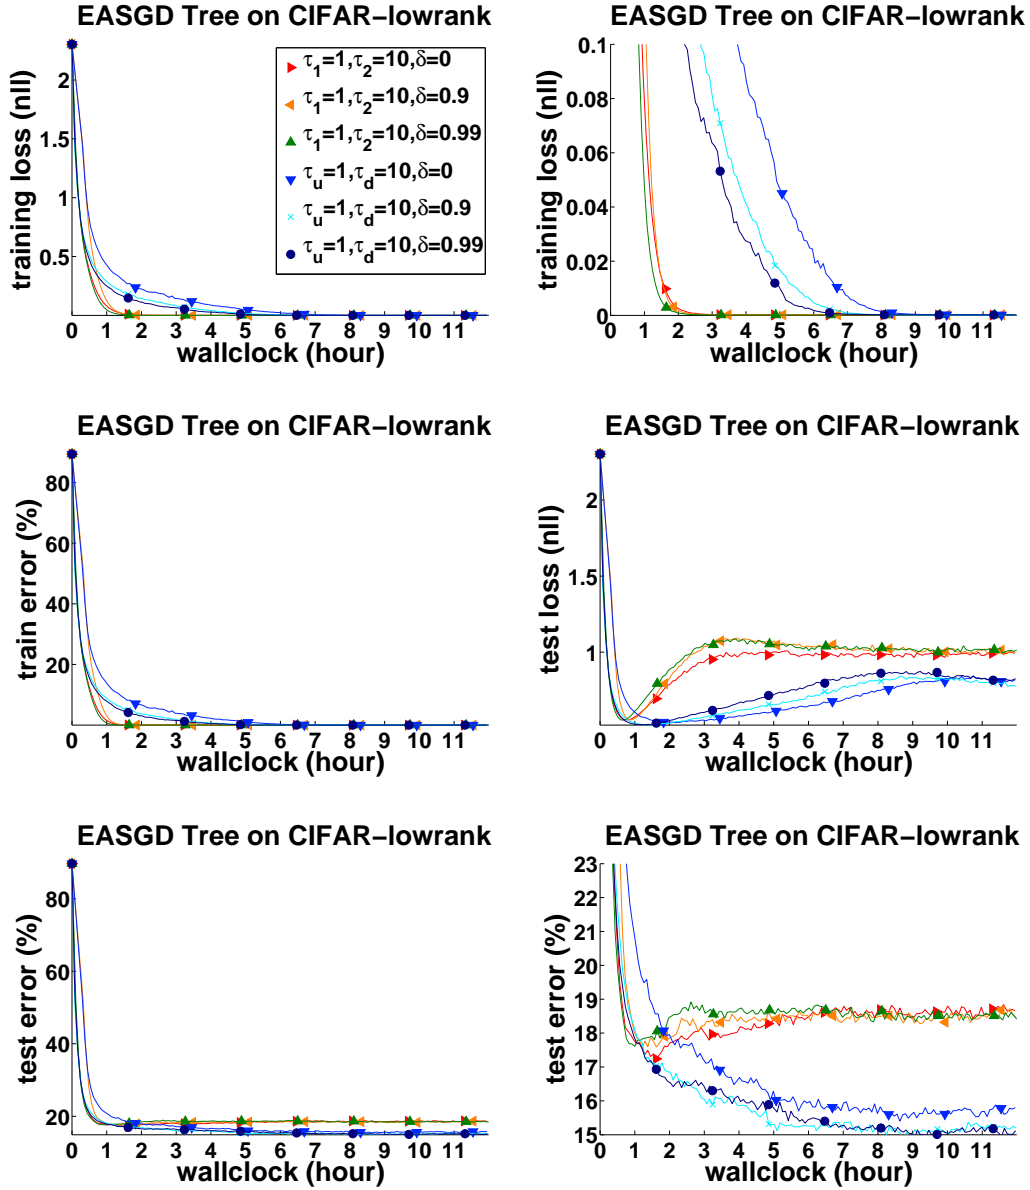


Figure 6.11: *EASGD Tree* on *CIFAR-lowrank* using the first and second communication scheme with momentum. Training loss and error, test loss and error of the root node versus a wallclock time. The curve e in Figure 6.5, the curve e in Figure 6.6, the curve a in Figure 6.7, the curve b in Figure 6.8, the curve b in Figure 6.9 and the curve d in Figure 6.10 are selected and plotted.



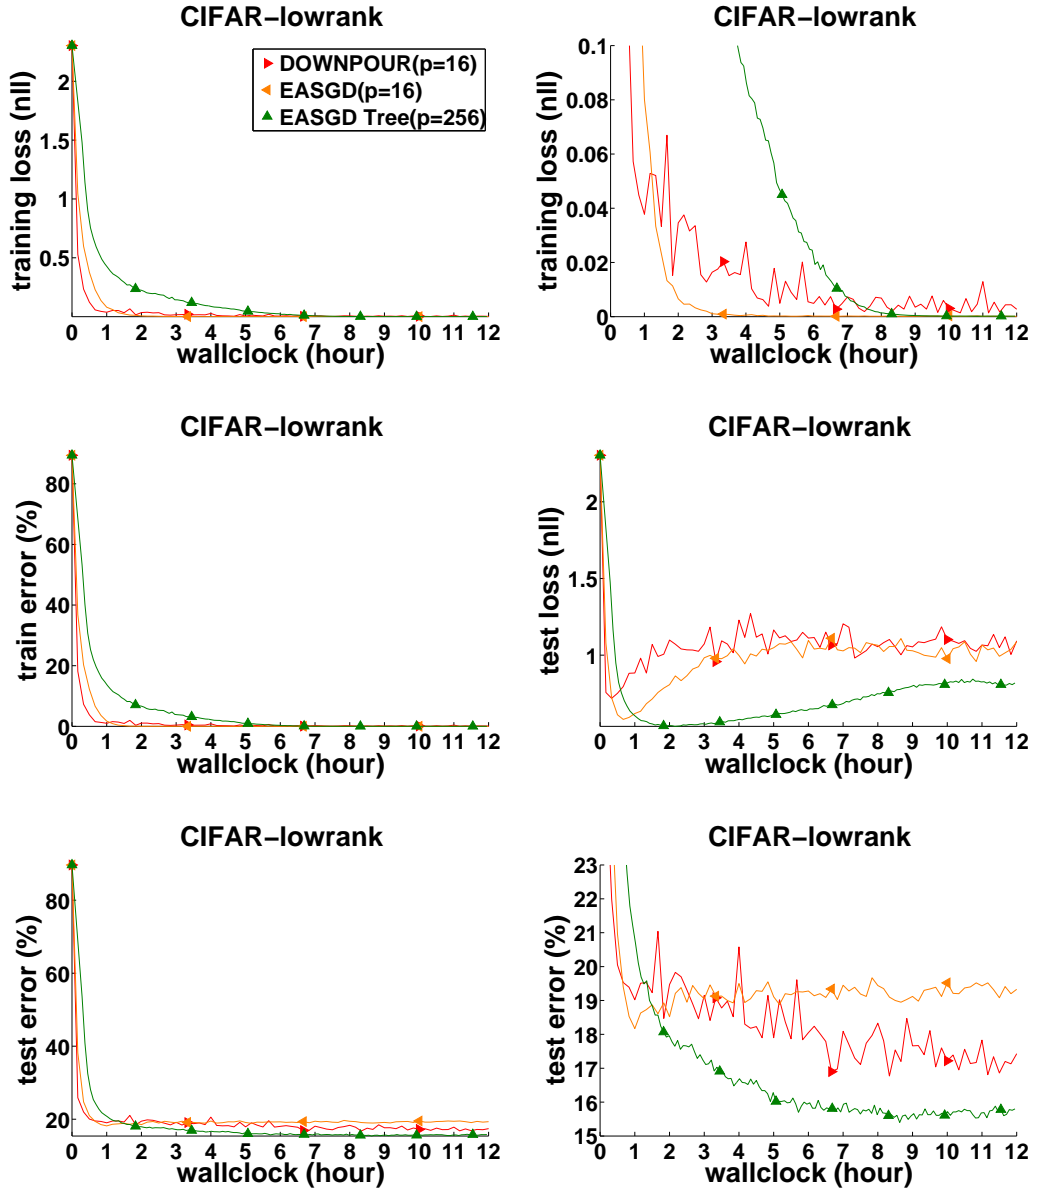


Figure 6.12: Best test performance of *DOWNPOUR* ( $p=16$ ), *EASGD* ( $p=16$ ), and *EASGD Tree* ( $p=256$ ) on *CIFAR-lowrank*. Training loss and error, test loss and error (of the center variable or the root node) versus a wallclock time. The momentum is not applied ( $\delta = 0$ ).

gives a smallest test error 14.91% in curve d. The most surprising observation is that the curve a has a large peak between hour 8 and 9. We have investigated the log and found that it is correlated with the network traffic.

We can now compare all the momentum results by selecting the best test performance curve in previous results. They are summarized in Figure 6.11. It's very clear that the two communication schemes give very different performance. The first scheme gives better training performance, while the second scheme gives better test performance. The momentum no longer plays a significant role to determine the test performance as in *EAMSGD*, but it is still helpful. To explain why still remains very open.

For completeness, we compare *EASGD*, *DOWNPOUR* with  $p = 16$ , and the *EASGD Tree* with  $p = 256$ . Figure 6.12 shows the best test performance for each method, selected among a wide range of the learning rates.

## 6.2 Unifying *EASGD* and *DOWNPOUR*

It's beneficial to write down the updates of the *DOWNPOUR* method (Algorithm 3 in Section 4.4) in a synchronous scenario. For simplicity, let's assume that the communication period  $\tau = 1$ . The global clock  $t$  is shared among the local workers. Every step *DOWNPOUR* performs the following updates to solve the original Problem 1.1 in Chapter 1:

$$\begin{aligned} x_{t+1}^i &= \tilde{x}_t - \eta \nabla f(\tilde{x}_t, \xi_t^i), \\ \tilde{x}_{t+1} &= \tilde{x}_t + \sum_{i=1}^p (x_{t+1}^i - \tilde{x}_t). \end{aligned} \tag{6.1}$$

Now compare with the *EASGD* updates (from Equation 2.4 in Chapter 2):

$$\begin{aligned} x_{t+1}^i &= x_t^i - \eta \nabla f(x_t^i, \xi_t^i) - \alpha (x_t^i - \tilde{x}_t), \\ \tilde{x}_{t+1} &= \tilde{x}_t + \frac{\beta}{p} \sum_{i=1}^p (x_t^i - \tilde{x}_t). \end{aligned} \tag{6.2}$$

We realize that it's possible to unify these two methods by drawing a parallel between

the classical *Jacobi* and *Gauss-Seidel* methods. To motivate this connection, let's look back at the convergence result of *EASGD*. Lemma 3.1.1 indicates that the divergence can happen if the condition 3.4 is not satisfied. One might wonder why there's such a stability condition. After all, *EASGD* performs only the gradient descent and the moving averages. Each of these operations is contractive on its own! A short answer is that *EASGD* is formalized in the *Jacobi* form rather than the *Gauss-Seidel* form in order to capture the delay of the parameter communication. We can transform the *EASGD*'s *Jacobi* form in Equation 6.2 to a *Gauss-Seidel* form as follows:

$$\begin{aligned} x_{t+1/2}^i &= x_t^i - \alpha(x_t^i - \tilde{x}_t), \\ x_{t+1}^i &= x_{t+1/2}^i - \eta \nabla f(x_{t+1/2}^i, \xi_{t+1/2}^i), \\ \tilde{x}_{t+1} &= \tilde{x}_t + \frac{\beta}{p} \sum_{i=1}^p (x_{t+1}^i - \tilde{x}_t). \end{aligned} \tag{6.3}$$

The first and the third step in Equation 6.3 is contractive if we choose  $\alpha \in (0, 2)$  and  $\beta \in (0, 2)$ . Thus the stability condition 3.4 is too much a constraint in this case. Notice that the *DOWNPOUR* method in Equation 6.1 is equivalent to the *EASGD*'s *Gauss-Seidel* form in Equation 6.3 with  $\alpha = 1$  and  $\beta = p$ . It is in this sense that *DOWNPOUR* and *EASGD* get connected! However,  $\beta = p$  can be very large, how can it be stable?

For stability analysis, we consider again the quadratic case where  $\nabla f(x, \xi) = x$ . We can reduce the Equation 6.3 by introducing  $\bar{x}_t = \frac{1}{p} \sum_{i=1}^p x_t^i$  as follows:

$$\begin{aligned} \bar{x}_{t+1/2} &= \bar{x}_t - \alpha(\bar{x}_t - \tilde{x}_t), \\ \bar{x}_{t+1} &= \bar{x}_{t+1/2} - \eta \bar{x}_{t+1/2}, \\ \tilde{x}_{t+1} &= \tilde{x}_t + \beta(\bar{x}_{t+1} - \tilde{x}_t). \end{aligned} \tag{6.4}$$

This is a linear system which can be written as  $(\bar{x}_{t+1}, \tilde{x}_{t+1})^T = M(\bar{x}_t, \tilde{x}_t)^T$ , where

$$M = \begin{pmatrix} 1 & 0 \\ \beta & 1 - \beta \end{pmatrix} \begin{pmatrix} 1 - \eta & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 - \alpha & \alpha \\ 0 & 1 \end{pmatrix}. \tag{6.5}$$

We evaluate the two eigenvalues  $z_1$  and  $z_2$  of this matrix  $M$ , which should satisfy the equation  $z^2 - 2bz + c = 0$ . One can check that

$$\begin{aligned} b &= \frac{1}{2}(2 - \beta - \eta - \alpha(1 - \beta)(1 - \eta)), \\ c &= (1 - \eta)(1 - \alpha)(1 - \beta). \end{aligned} \tag{6.6}$$

We need a basic lemma to study the stability, i.e. when the absolute value of  $z_1$  and  $z_2$  are smaller than one.

**Lemma 6.2.1.** *Let  $z_1$  and  $z_2$  be the two roots of the polynomial  $z^2 - 2bz + c = 0$  with real coefficients  $b \in \mathbb{R}$  and  $c \in \mathbb{R}$ , then the sufficient and necessary condition for  $|z_1| < 1$  and  $|z_2| < 1$  is*

$$-1 < c < 1, \quad c > 2b - 1, \quad c > -2b - 1. \tag{6.7}$$

*Proof.* There are two cases to consider. The first one is the real-valued case, i.e.  $b^2 > c$ . The second one is the imaginary case, i.e.  $b^2 \leq c$ . If  $b^2 \leq c$ , then  $|z_1| = |z_2| = \sqrt{c} < 1$  is equivalent to  $-1 < c < 1$ . If  $b^2 > c$  then we can suppose  $z_1 = b + \sqrt{b^2 - c}$  and  $z_2 = b - \sqrt{b^2 - c}$ . The condition  $-1 < z_2 < z_1 < 1$  is equivalent to  $c > 2b - 1$ ,  $b < 1$ ,  $c > -2b - 1$  and  $b > -1$ . Combining the condition from the two cases gives us the Equation 6.7.  $\square$

Applying the Lemma 6.7 to the above case in Equation 6.6, we obtain the stability condition for the reduced linear system 6.4 as below,

$$-1 < c < 1, \quad 0 < \beta\eta < 2(c + 1),$$

where  $c = (1 - \eta)(1 - \alpha)(1 - \beta)$ . It's easy to verify this since  $2b = c + 1 - \beta\eta$ .

We summarize the obtained stability condition (using Mathematica) in the following theorem.

**Theorem 6.2.1.** *Assume  $\eta > 0$ ,  $\alpha > 0$  and  $\beta > 0$ , the stability condition for the linear system 6.4 is equivalent to union of the following set of conditions:*

- $0 < \eta < 1, 0 < \beta < 1$  and  $0 < \alpha < c_2$ ,
- $0 < \eta < 1, \beta = 1$  and  $\alpha > 0$ ,
- $0 < \eta < 1, 1 < \beta \leq 2$  and  $0 < \alpha < c_3$ ,
- $0 < \eta < 1, 2 < \beta < 4/\eta$  and  $c_2 < \alpha < c_3$ ,
- $\eta = 1, 0 < \beta < 2$  and  $\alpha > 0$ ,
- $1 < \eta < 2, 0 < \beta < 1$  and  $0 < \alpha < c_3$ ,
- $1 < \eta < 2, \beta = 1$  and  $\alpha > 0$ ,
- $1 < \eta < 2, 1 < \beta < 2$  and  $0 < \alpha < c_2$ ,
- $2 \leq \eta \leq 4, 0 < \beta < 1$  and  $c_2 < \alpha < c_3$ ,
- $\eta > 4, 0 < \beta < 4/\eta$  and  $c_2 < \alpha < c_3$ ,

where

$$c_2 = \frac{4 - 2\beta - 2\eta + \beta\eta}{2 - 2\beta - 2\eta + 2\beta\eta}, \quad c_3 = \frac{-\beta - \eta + \beta\eta}{1 - \beta - \eta + \beta\eta}.$$

There is only one condition above which satisfies the *DOWNPOUR* method when  $p > 2$ . This condition is  $0 < \eta < 1, 2 < \beta < 4/\eta$  and  $c_2 < \alpha < c_3$ . It means that if  $\beta = p$  is very large, we need to use quite small learning rate such that  $0 < \eta < 4/p$ . Moreover,  $\alpha$  has to stay very close to 1 (i.e. forgetting all the local memory). This is because  $c_2$  tends to 1 as  $\beta = p$  tends to infinity (assume  $\beta\eta < 4$  always holds); and  $c_3$  tends to 1 as well. For example, this condition does not hold for  $\alpha = 0.9$  and  $\beta = 16$ , for any  $\eta > 0$ . This is rather surprising: the stability condition for  $\alpha = 1$  and  $\beta = 16$  is  $0 < \eta < 1/8$ . However, there's no  $\eta > 0$  for  $\alpha = 0.9$  and  $\beta = 16$  to be stable. This suggest *DOWNPOUR* is a very peculiar singleton in the class of the methods defined by Equation 6.3. *EASGD* has much more flexibilities as it operates in the region where  $\beta \in (0, 2)$ . In practice, we have observed that the case  $\beta = 1, 0 < \eta < 2, \alpha > 0$  works well. It is very dangerous to operate in the region of  $\eta > 2$  when  $\alpha$  is small.

## Chapter 7

# Conclusion

In this thesis, we focused on the problem of training large-scale deep learning models in a parallel and distributed computing environment.

Our starting point is a reformulation of the global variable consensus problem into an unconstrained optimization problem using a quadratic penalty between each local variable and the center variable. By reinterpreting the gradient of the quadratic penalty as an averaging process, we introduced the Elastic Averaging SGD (*EASGD*) method. The *EASGD* maintains the stochastic nature of the sequential *SGD* method through a weak coupling between the local variables and the center variable. Each local variable is optimized using an *SGD*-based method. The averaging process attracts the local variable and the center variable toward each other so that all of them move toward a local optimum. The center variable can slowly track the spatial average of the local variables, thus having a smaller variance. We then discussed how to extend the *EASGD* method to the asynchronous scenario, and how to accelerate it using the momentum.

We proceeded to study the convergence rate the *EASGD* method in terms of the center variable. In the synchronous scenario, we discussed the variance reduction effect of the *EASGD* method by increasing the number of local variables, for both quadratic and strongly-convex case. We also introduced a double averaging sequence in the spirit of

the *Polyak*'s averaging and showed that it is indeed asymptotically optimal. We then compared the stability region of the *EASGD* method with the *ADMM* method in the round-robin scheme, where we have found quite unusual instability region for the *ADMM* method.

Having studied the (local) convergence property of the *EASGD* method, we applied its asynchronous extension to train deep learning models. We focused on the *CIFAR-10* and *ImageNet ILSVRC 2013* dataset for image classification in a supervised learning setting. The (global) convergence behavior of the *EASGD* method was simulated starting from a random initialization. We found that the *EASGD* method accelerates the training of the baseline *SGD* method. Moreover, it is very communication efficient compared to the asynchronous *SGD* method *DOWNPOUR*. After combining with *Nesterov*'s momentum method, we obtained *EAMSGD* and it achieved even better test accuracy. We also discussed the tradeoff between the data communication and the parameter communication. In our cases, it turned out to save the total bandwidth a lot by using more data communication for less parameter communication (recall *EASGD* method needs to sample the whole dataset in order to avoid the bias of the stochastic gradient).

Based on the empirical results we have obtained, we asked what is the limit of the speedup by increasing the number of processors. We studied first an additive noise model to capture the asymptotic behavior of various stochastic optimization method, then we proposed a multiplicative noise model to capture the initial behavior of these methods. For the additive noise case, we first studied the *SGD* method with mini-batch. We saw that increasing the mini-batch size (i.e. number of processors) can reduce the asymptotic variance, but it is not able to increase the convergence speed. We then studied the momentum *SGD* in the *Nesterov*'s form, and showed that faster convergence speed is possible, but that will lead to larger asymptotic variance. We have also generalised the *EASGD* method by decoupling (de-symmetrize) the elastic averaging between the center variable and the local variables, and found that the *EASGD* method can behavior like a momentum *SGD* with a negative moving average rate in some optimal sense (i.e.

the center variable pushes the local variables instead of pulling). For the multiplicative noise case, we have found that *SGD* method with mini-batch can greatly improved the initial convergence speed by choosing the optimal (larger) learning rate. However, the momentum *SGD* is not as effective in this case. The *EASGD* method is also slower than *SGD* with mini-batch, but its stability region is still improved by increasing the number of processors. We showed that given an infinite number of processors, the improvement for stability region of *EASGD* depends on the spread of the Gamma distribution of the input data. The larger the spread, the more the improvement.

Back to our starting point, we studied an interesting non-convex case for the *EASGD* method based on the an unconstrained objective we have introduced. We have discussed the stability of the critical points. We found that when the quadratic penalty is too small (i.e. the coupling between the center variable and the local variable is too weak), the *EASGD* method can introduce a stable local optimum trapped by a saddle point.

Knowing the theoretical limitation of the speedup, we scale up the *EASGD* by using a tree structured network topology. We showed empirically that *EASGD Tree* can further accelerate the training speed, but no longer in a linear speedup region. This is somehow predicted by our analysis in the multiplicative noise model. The surprising observation is *EASGD Tree* can still yield better test accuracy when the local variables fluctuate further from the center variable. Moreover, inspired by the design of the *EASGD Tree* in the asynchronous scenario, we carefully distinguished the difference between the *Jacobi* form of *EASGD* with its *Gauss-Seidel* form, which in turn unified the *EASGD* and the *DOWNPOUR* method in the synchronous scenario.

We conclude with a few open problems from convex optimization, non-convex optimization, distributed and parallel computing, statistical and deep learning.

- The smoothing property of *EASGD* for non-smooth convex optimization. We have discussed the connection between the quadratic penalty term in Equation 1.2 and the Moreau-Yosida regularization. Can one expect a good rate of convergence?



- The connection between *EAMSGD* and *Katyusha* [3]. *Katyusha* is an accelerated variance reduction method. The speedup of its acceleration is dragged back by a center variable so as to achieve a smaller variance. This resembles the role of the center variable in *EAMSGD*.
- In the simulation of *EASGD Tree*, we have seen that network congestion can occur. Is there a good way to detect when the network becomes slow so as to avoid the potential congestion? This is an adaptive learning rate problem. It is not very simple as we have seen in Chapter 5 that the optimal moving rate can sometime be positive, but sometime be negative.
- The convergence analysis of the *EASGD Tree* is still missing. The basic tool [57] based on the mean-field method (i.e. treating each node the same) seems to be too macroscopic to capture the global (e.g. multi-scale variance reduction) behavior of the *EASGD Tree*.
- The description of the asynchronous behavior using vector-clock [53]. Although we have unified *EASGD* and *DOWNPOUR* in the synchronous scenario, it is still not clear how to measure their behavior in the asynchronous scenario and in the real-time workloads. The idea of using vector-clock from the distributed computing is worthing exploring in the future.
- How non-convex is the energy landscape? A related question may be what to do when *EASGD* is trapped by a saddle point. Also why there is a such a big difference between *ASGD* and *MVASGD* as we have studied in Figure 4.11?
- What is the role of the oscillation in deep learning to achieve better test performance? We have observed that both *EAMSGD* and *EASGD Tree* can achieve better test performance when the oscillation of the local variables is sufficiently large. Looking back to the Dropout [54] and DropConnect [58] regularization, is there any deeper reason in common [51]?

# Bibliography

- [1] A. Agarwal, O. Chapelle, M. Dudík, and J. Langford. A reliable effective terascale linear learning system. *The Journal of Machine Learning Research*, 15(1):1111–1133, 2014.
- [2] A. Agarwal and J. C. Duchi. Distributed delayed stochastic optimization. In *NIPS*, pages 873–881, 2011.
- [3] Z. Allen-Zhu. Katyusha: Accelerated variance reduction for faster sgd. *arXiv preprint arXiv:1603.05953*, 2016.
- [4] S. Azadi and S. Sra. Towards an optimal stochastic alternating direction method of multipliers. In *ICML*, 2014.
- [5] R. Bekkerman, M. Bilenko, and J. Langford. Scaling up machine learning: Parallel and distributed approaches. Cambridge University Press, 2011.
- [6] D. P. Bertsekas. A new class of incremental gradient methods for least squares problems. *SIAM Journal on Optimization*, 7(4):913–926, 1997.
- [7] D. P. Bertsekas and J. N. Tsitsiklis. *Parallel and Distributed Computation*. Prentice Hall, 1989.
- [8] V. Borkar. Asynchronous stochastic approximations. *SIAM Journal on Control and Optimization*, 36(3):840–851, 1998.

- [9] L. Bottou. Online algorithms and stochastic approximations. In *Online Learning and Neural Networks*. Cambridge University Press, 1998.
- [10] L. Bottou. Curiously fast convergence of some stochastic gradient descent algorithms. Unpublished open problem offered to the attendance of the SLDS 2009 conference, 2009.
- [11] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.*, 3(1):1–122, 2011.
- [12] N. Cesa-Bianchi, A. Conconi, and C. Gentile. On the generalization ability of on-line learning algorithms. *IEEE Transactions on Information Theory*, 50(9):2050–2057, 2004.
- [13] A. Choromanska, M. B. Henaff, M. Mathieu, G. B. Arous, and Y. LeCun. The loss surfaces of multilayer networks. In *AISTATS*, 2015.
- [14] A. Cotter, O. Shamir, N. Srebro, and K. Sridharan. Better mini-batch algorithms via accelerated gradient methods. In *Advances in neural information processing systems*, pages 1647–1655, 2011.
- [15] D. Das, S. Avancha, D. Mudigere, K. Vaidynathan, S. Sridharan, D. Kalamkar, B. Kaul, and P. Dubey. Distributed deep learning using synchronous stochastic gradient descent. *arXiv preprint arXiv:1602.06709*, 2016.
- [16] J. Dean, G. Corrado, R. Monga, K. Chen, M. Devin, Q. Le, M. Mao, M. Ranzato, A. Senior, P. Tucker, K. Yang, and A. Ng. Large scale distributed deep networks. In *NIPS*, 2012.
- [17] O. Dekel, R. Gilad-Bachrach, O. Shamir, and L. Xiao. Optimal distributed on-line prediction using mini-batches. *The Journal of Machine Learning Research*, 13(1):165–202, 2012.

- [18] O. Devolder, F. Glineur, and Y. Nesterov. First-order methods of smooth convex optimization with inexact oracle. *Mathematical Programming*, 146(1-2):37–75, 2014.
- [19] J. C. Duchi, A. Agarwal, and M. J. Wainwright. Dual averaging for distributed optimization: convergence analysis and network scaling. *Automatic control, IEEE Transactions on*, 57(3):592–606, 2012.
- [20] M. Feng, B. Xiang, and B. Zhou. Distributed deep learning for answer selection. *CoRR*, abs/1511.01158, 2015.
- [21] E. Ghadimi, H. R. Feyzmahdavian, and M. Johansson. Global convergence of the heavy-ball method for convex optimization. In *Control Conference (ECC), 2015 European*, pages 310–315. IEEE, 2015.
- [22] A. Halevy, P. Norvig, and F. Pereira. The unreasonable effectiveness of data. *Intelligent Systems, IEEE*, 24(2):8–12, 2009.
- [23] M. Hardt, B. Recht, and Y. Singer. Train faster, generalize better: Stability of stochastic gradient descent. *arXiv preprint arXiv:1509.01240*, 2015.
- [24] M. R. Hestenes. *Optimization theory: the finite dimensional case*. Wiley, 1975.
- [25] Q. Ho, J. Cipar, H. Cui, S. Lee, J. K. Kim, P. B. Gibbons, G. A. Gibson, G. Ganger, and E. P. Xing. More effective distributed ml via a stale synchronous parallel parameter server. In *Advances in Neural Information Processing Systems 26*, pages 1223–1231. Curran Associates, Inc., 2013.
- [26] H. Kim, J. Park, J. Jang, and S. Yoon. Deepspark: Spark-based deep learning supporting asynchronous updates and caffe compatibility. *arXiv preprint arXiv:1602.08191*, 2016.
- [27] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems 25*, pages 1106–1114, 2012.

- [28] G. Lan. An optimal method for stochastic composite optimization. *Mathematical Programming*, 133(1-2):365–397, 2012.
- [29] J. Langford, A. Smola, and M. Zinkevich. Slow learners are fast. In *NIPS*, 2009.
- [30] Y. A. LeCun, L. Bottou, G. B. Orr, and K.-R. Müller. Efficient backprop. In *Neural networks: Tricks of the trade*, pages 9–48. Springer, 2012.
- [31] M. Li, T. Zhang, Y. Chen, and A. J. Smola. Efficient mini-batch training for stochastic optimization. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 661–670. ACM, 2014.
- [32] Q. Li, C. Tai, and W. E. Dynamics of stochastic gradient algorithms. *CoRR*, abs/1511.06251, 2015.
- [33] H. Lin, J. Mairal, and Z. Harchaoui. A universal catalyst for first-order optimization. In *Advances in Neural Information Processing Systems*, pages 3366–3374, 2015.
- [34] P. Moritz, R. Nishihara, I. Stoica, and M. I. Jordan. Sparknet: Training deep networks in spark. *arXiv preprint arXiv:1511.06051*, 2015.
- [35] V. Nair and G. E. Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pages 807–814, 2010.
- [36] A. Nedić, D. Bertsekas, and V. Borkar. Distributed asynchronous incremental sub-gradient methods. In *Inherently Parallel Algorithms in Feasibility and Optimization and their Applications*, volume 8 of *Studies in Computational Mathematics*, pages 381 – 407, 2001.
- [37] A. Nedi and A. Ozdaglar. Cooperative distributed multi-agent optimization. In D. P. Palomar and Y. C. Eldar, editors, *Convex Optimization in Signal Processing and Communications*, pages 340–386. Cambridge University Press, 2009. Cambridge Books Online.

- [38] Y. Nesterov. *Introductory lectures on convex optimization*, volume 87. Springer Science & Business Media, 2004.
- [39] Y. Nesterov. Smooth minimization of non-smooth functions. *Math. Program.*, 103(1):127–152, 2005.
- [40] J. Nocedal and S. Wright. *Numerical Optimization, Second Edition*. Springer New York, 2006.
- [41] A. Olshevsky and J. N. Tsitsiklis. Convergence speed in distributed consensus and averaging. *SIAM review*, 53(4):747–772, 2011.
- [42] H. Ouyang, N. He, L. Tran, and A. Gray. Stochastic alternating direction method of multipliers. In *Proceedings of the 30th International Conference on Machine Learning*, pages 80–88, 2013.
- [43] T. Paine, H. Jin, J. Yang, Z. Lin, and T. Huang. Gpu asynchronous stochastic gradient descent to speed up neural network training. *arXiv preprint arXiv:1312.6186*, 2013.
- [44] B. T. Polyak. *Introduction to optimization*. Optimization Software New York, 1987.
- [45] B. T. Polyak and A. B. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855, 1992.
- [46] B. Recht, C. Re, S. J. Wright, and F. Niu. Hogwild: A Lock-Free Approach to Parallelizing Stochastic Gradient Descent. In *NIPS*, 2011.
- [47] Y. Saad. *Iterative methods for sparse linear systems*. SIAM, 2003.
- [48] F. Seide, H. Fu, J. Droppo, G. Li, and D. Yu. 1-bit stochastic gradient descent and application to data-parallel distributed training of speech dnns. In *Interspeech 2014*, September 2014.

- [49] P. Sermanet, D. Eigen, X. Zhang, M. Mathieu, R. Fergus, and Y. LeCun. Overfeat: Integrated recognition, localization and detection using convolutional networks. In *International Conference on Learning Representations (ICLR2014)*, 2014.
- [50] O. Shamir. Fundamental limits of online and distributed algorithms for statistical learning and estimation. In *NIPS*, 2014.
- [51] D. J. Sherratt. Oscillation helps to get division right. *Proceedings of the National Academy of Sciences*, page 201601792, 2016.
- [52] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.
- [53] M. Singhal and A. Kshemkalyani. An efficient implementation of vector clocks. *Information Processing Letters*, 43(1):47–52, 1992.
- [54] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15:1929–1958, 2014.
- [55] I. Sutskever, J. Martens, G. Dahl, and G. Hinton. On the importance of initialization and momentum in deep learning. In *ICML*, 2013.
- [56] C. Tai, T. Xiao, X. Wang, and W. E. Convolutional neural networks with low-rank regularization. In *ICLR*, 2016.
- [57] J. N. Tsitsiklis, D. P. Bertsekas, and M. Athans. Distributed asynchronous deterministic and stochastic gradient optimization algorithms. In *1984 American Control Conference*, pages 484–489, 1984.
- [58] L. Wan, M. D. Zeiler, S. Zhang, Y. LeCun, and R. Fergus. Regularization of neural networks using dropconnect. In *ICML*, 2013.

- [59] L. Xiao, S. Boyd, and S.-J. Kim. Distributed average consensus with least-mean-square deviation. *Journal of Parallel and Distributed Computing*, 67(1):33–46, 2007.
- [60] O. Yadan, K. Adams, Y. Taigman, and M. Ranzato. Multi-gpu training of convnets. In *Arxiv*, 2013.
- [61] R. Zhang and J. Kwok. Asynchronous distributed admm for consensus optimization. In T. Jebara and E. P. Xing, editors, *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 1701–1709. JMLR Workshop and Conference Proceedings, 2014.
- [62] W. Zhang, S. Gupta, X. Lian, and J. Liu. Staleness-aware async-sgd for distributed deep learning. *arXiv preprint arXiv:1511.05950*, 2015.
- [63] M. Zinkevich, M. Weimer, A. Smola, and L. Li. Parallelized stochastic gradient descent. In *NIPS*, 2010.