



City Research Online

City, University of London Institutional Repository

Citation: Fu, X., Sturtz, J., Alonso, E., Qingge, L. & Chaloo, R. (2024). Parallel Trajectory Training of Recurrent Neural Network Controllers with Levenberg–Marquardt and Forward Accumulation Through Time in Closed-loop Control Systems. *IEEE Transactions on Sustainable Computing*, 9(2), pp. 222-229. doi: 10.1109/tsusc.2023.3330573

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/31614/>

Link to published version: <https://doi.org/10.1109/tsusc.2023.3330573>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

City Research Online:

<http://openaccess.city.ac.uk/>

publications@city.ac.uk

Parallel Trajectory Training of Recurrent Neural Network Controllers with Levenberg–Marquardt and Forward Accumulation Through Time in Closed-loop Control Systems

Xingang Fu, *Senior Member, IEEE*, Jordan Sturtz, Eduardo Alonso, Letu Qingge, and [Rajab Chaloo](#)

Abstract—This paper introduces a novel parallel trajectory mechanism that combines Levenberg-Marquardt and Forward Accumulation Through Time algorithms to train a recurrent neural network controller in a closed-loop control system by distributing the calculation of trajectories across Central Processing Unit (CPU) cores/workers depending on the computing platforms, computing program languages, and software packages available. Without loss of generality, the recurrent neural network controller of a grid-connected converter for solar integration to a power system was selected as the benchmark test closed-loop control system. Two software packages were developed in Matlab and C++ to verify and demonstrate the efficiency of the proposed parallel training method. The training of the deep neural network controller was migrated from a single workstation to both cloud computing platforms and High-Performance Computing clusters. The training results show excellent speed-up performance, which significantly reduces the training time for a large number of trajectories with high sampling frequency, and further demonstrates the effectiveness and scalability of the proposed parallel mechanism.

Index Terms—Parallel Trajectory Training, Recurrent Neural Network Controller, Forward Accumulation Through Time, Levenberg–Marquardt, Cloud Computing, High-Performance Computing (HPC) Cluster

I. INTRODUCTION

The Levenberg-Marquardt (LM) algorithm provides a nice compromise between the speed of the second-order Newton’s method and the guaranteed convergence of first-order steepest descent method to solve nonlinear least squares problems [1], [2]. Thus it is particularly suitable for training small and medium-sized feed-forward Neural Networks (NNs) [3].

However, the computation loads of the LM algorithm are expensive due to the calculation needs of the Jacobian matrix,

This work was supported in part by the U.S. National Science Foundation under Grant CISE-MSI 2131214 and 2131175. (Corresponding author: Letu Qingge.)

Xingang Fu is with the Department of Electrical and Biomedical Engineering, The University of Nevada, Reno, Reno, NV 89557, USA (email: xfu@unr.edu).

Jordan Sturtz and Letu Qingge are with the Department of Computer Science, North Carolina A&T State University, Greensboro, NC 27411, USA (email: jasturtz@aggies.ncat.edu and lqingge@ncat.edu).

Eduardo Alonso is with the Artificial Intelligence Research Centre (CitAI), City, University of London, EC1V 0HB, UK (email: e.alonso@city.ac.uk).

Rajab Chaloo is with the Department of Electrical Engineering and Computer Science, Texas A&M Kingsville, Kingsville, TX 78363, USA (email: rajab.chaloo@tamuk.edu).

and researchers have explored ways to speed it up. For instance, the block-diagonal matrix has been proposed to approximate the Hessian matrix [4], [5], and, in [6], the forward difference method was used to approximate the Jacobian matrix by perturbing one parameter to produce a column, instead of calculating the Jacobian matrix directly.

In addition, several mechanisms have been proposed to parallelize the LM algorithm for NN training, appropriately distributing computational and space requirements. For example, the Single Program and Multiple Data (SPMD) strategy divides training data into groups and each group is distributed on one node in a cluster [7]. [8] also utilized the parallelization of data sets by calculating the objective functions simultaneously. Relatedly, [9] and [10] distributed the computing tasks/data points across parallel GPU multiprocessors to train the LM algorithm in parallel. In [11], the parallel selection of the damping parameter and multicore versions of the Basic-Linear-Algebra-Subprograms (BLAS) were used in the LM algorithm to increase computational efficiency. In any case, with or without parallelization, the application of feed-forward NNs is inherently limited due to their inability to identify and process sequential partners in large data sets.

Recurrent Neural Networks (RNNs) are potentially more powerful than feed-forward NNs thanks to their feedback connections and memory gates [3]. Many algorithms have been used for training RNNs such as Backpropagation Through Time (BPTT) [12], Real-Time Recurrent Learning (RTRL) [13], Extended Kalman Filters (EKF) [14], genetic algorithms [15], and Expectation Maximization (EM) [16], [17]. Notwithstanding their merits, they all suffer from serious drawbacks: the BPTT algorithm may cause gradient exploding and vanishing problems. The high computational cost of RTRL makes it only appropriate for online training of small RNNs. EKF is also computationally expensive since it requires many matrix calculations at each estimation. Evolutionary methods such as genetic algorithms have proved to be successful in training RNNs by formulating the RNN cost function as a nonlinear global optimization problem. However, they may get stuck in local minima and show a low speed of convergence. Finally, the application of EM to training neural networks is limited by the complicated calculations in the expectation step when the number of hidden neurons is large. Such deficiencies have been an impediment to the application of RNNs to real-life problems such as closed-loop control systems, that we take

as our benchmark.

Although some research has shown the potential of training RNNs using LM [18], [19], LM has not been used broadly for this purpose. The Forward Accumulation Through Time (FATT) algorithm was proposed to calculate the Jacobian matrix efficiently and combined with the LM algorithm to train an RNN controller applied to a power converter control system, which produced excellent performance [20]. However, training was based on a rather small number of trajectories (e.g., 10 trajectories) and a relatively low sampling frequency (e.g., 1000Hz) due to constraints in the computational power and the memory size of the single workstation used in the experiments.

To extend the LM training of RNNs to a large number of trajectories with high sampling frequency and accelerate the training process, this paper proposes a novel parallel trajectory training mechanism. The key contributions include 1) the introduction of a training mechanism tailored for RNNs with error integral terms in closed-loop control systems, a solar microinverter system in particular; 2) the development of a parallel LM and FATT algorithm designed for trajectory training of RNNs; 3) implementation of a parallel approach that efficiently distributes the FATT-based calculation of training trajectories across Central Processing Unit (CPU) cores/workers, 4) validation conducted through implementations in two programming languages: Matlab and C++; 5) comprehensive validation of the implementations and performance comparison on both cloud computing platforms and High-Performance Computing (HPC) clusters.

The rest of the paper is organized as follows. Section II introduces the RNN controllers in the benchmark test closed-loop control system. The parallel trajectory training algorithm of the RNN controllers is detailed in Section III. Training results using cloud platforms are presented in Section IV. Section V provides detailed implementation and training results on HPC clusters. Finally, the paper concludes with a summary of the main points in Section VI.

II. RNN CONTROLLERS IN A CLOSED-LOOP CONTROL SYSTEM FOR A SOLAR INVERTER

A. A Closed-loop Control System: a Solar Microinverter

Typically, solar inverters consist of two components: the DC-DC converter and the DC-AC inverter, as illustrated in the case of the Texas Instruments (TI) Microinverter in Fig. 1 [21]. The PhotoVoltaic (PV) solar panels attach to the DC-DC converter, while the DC-AC inverter maintains the voltage of the DC Bus at its rated value while feeding controlled AC current to the main power grid.

Fig. 2 further shows the schematic of a single-phase DC-AC inverter block with the LCL filters, in which a DC-link capacitor/DC bus is on the left, an LCL filter is placed in the middle, and a single-phase voltage source, representing the voltage at the Point of Common Coupling (PCC) of the AC power grid system is on the right [22].

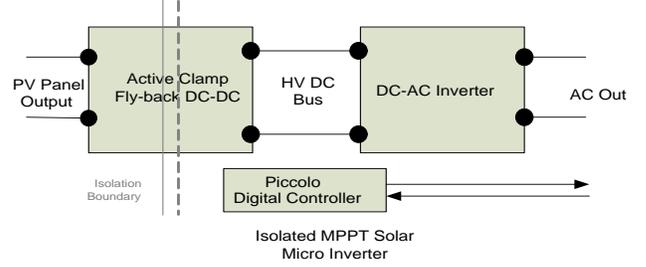


Fig. 1. TI Microinverter Block Diagram [21].

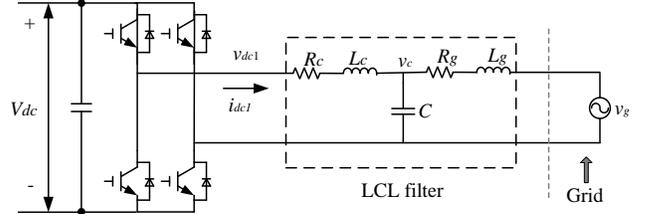


Fig. 2. The schematic of a single-phase DC-AC inverter block with LCL filters.

When using the d-q frame, the system state-space equation of the DC-AC inverter block can be described by equation (1) [23], which will be used for RNN training.

$$\frac{d}{dt} \begin{bmatrix} i_d \\ i_q \\ i_{dc_d1} \\ i_{dc_q1} \\ v_{cd} \\ v_{cq} \end{bmatrix} = \underbrace{\begin{bmatrix} -\frac{R_g}{L_g} & \omega_s & 0 & 0 & -\frac{1}{L_g} & 0 \\ -\omega_s & -\frac{R_g}{L_g} & 0 & 0 & 0 & -\frac{1}{L_g} \\ 0 & 0 & -\frac{R_c}{L_c} & \omega_s & \frac{1}{L_c} & 0 \\ 0 & 0 & -\omega_s & -\frac{R_c}{L_c} & 0 & \frac{1}{L_c} \\ \frac{1}{C} & 0 & -\frac{1}{C} & 0 & 0 & \omega_s \\ 0 & \frac{1}{C} & 0 & -\frac{1}{C} & -\omega_s & 0 \end{bmatrix}}_A \begin{bmatrix} i_d \\ i_q \\ i_{dc_d1} \\ i_{dc_q1} \\ v_{cd} \\ v_{cq} \end{bmatrix} + \underbrace{\begin{bmatrix} \frac{1}{L_g} & 0 & 0 & 0 & 0 & 0 \\ 0 & \frac{1}{L_g} & 0 & 0 & 0 & 0 \\ 0 & 0 & -\frac{1}{L_c} & 0 & 0 & 0 \\ 0 & 0 & 0 & -\frac{1}{L_c} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}}_B \begin{bmatrix} v_d \\ v_q \\ v_{d1} \\ v_{q1} \\ 0 \\ 0 \end{bmatrix} \quad (1)$$

where ω_s is the angular frequency of the grid voltage, and all other symbols are consistent with those shown in Fig. 2. i_d and i_q represent system states that need to be controlled. The controller outputs will be v_{d1} and v_{q1} . System parameters in equation (1) can be obtained from the user guide or datasheet of TI microinverter [21].

To train the RNN digital controller, the continuous state space model in equation (1) must first be converted into an equivalent discrete model using equation (2), either through a zero-order or a first-order hold discrete equivalent mechanism with a sampling time of T_s . For example, if the sampling frequency equals 10000Hz, then $T_s = 1/10000 = 0.1ms$.

$$\vec{i}_{dqs}(k+1) = A\vec{i}_{dqs}(k) + B\vec{u}_{dqs}(k) \quad (2)$$

in which, A stands for system matrix and B is the input matrix.

B. Trajectory Tracking

Fig. 5 demonstrates the reference trajectories for RNN training, which contains 6 trajectories: 3 for i_d and 3 for i_q . The reference trajectories were generated randomly within the system's controllable range. This range can be determined by the physical current/voltage ratings of solar inverters. For demonstration purposes, the range was set as $[-100, 100]$ in Fig. 5. The reference i_d and i_q values were set to change after certain time steps, e.g. 100, which is a tunable parameter. Normally, the control system needs time to reach its steady state, the 100 time steps turned out to be a well-balanced number in training the RNN controller for a solar inverter. The total time steps/trajectory length was set to a certain number, e.g. 1000 in Fig. 5, which is determined by training duration and sampling time. For example, if the sampling time $T_s = 1ms$ and the training duration $1s$ are used, the trajectory length will be $1s/T_s = 1000$. The total number of training trajectories can vary from 10 to several hundred. Utilizing a large number of trajectories will significantly slow down training, which is the reason for proposing the parallel trajectory training algorithm for overcoming this challenge.

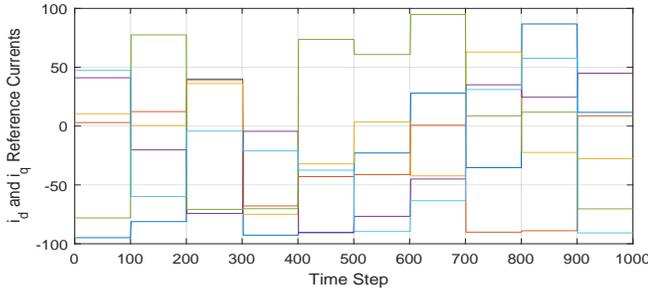


Fig. 5. The reference trajectories for RNN training.

Fig. 6 shows the ideal tracking performance after the RNN was well-trained. After training, an RNN controller will be able to regulate the system states i_d and i_q to follow the reference currents i_{d_ref} and i_{q_ref} .

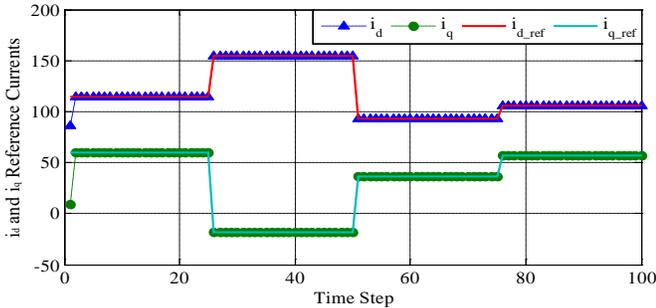


Fig. 6. The ideal tracking performance for RNN training.

C. LM Algorithm

If the performance error function is not a sum of squares, then the LM weight update equation is not directly applicable. To implement LM training, the cost function defined in (6) needs to be rewritten in a Sum-Of-Squares form. Consider the

cost function C_{dp} with $\gamma = 1$, $j = 1$ and $k = 1, \dots, N$, then it can be written in the form

$$C_{dp} = \sum_{k=1}^N U(\vec{e}_{dq}(k)) \xleftrightarrow{\text{def } V(k) = \sqrt{U(\vec{e}_{dq}(k))}} C_{dp} = \sum_{k=1}^N V^2(k) \quad (7)$$

and the gradient $\frac{\partial C_{dp}}{\partial \vec{w}}$ can be written in a matrix product form

$$\frac{\partial C_{dp}}{\partial \vec{w}} = \sum_{k=1}^N V(k) \frac{\partial V(k)}{\partial \vec{w}} = 2J_v(\vec{w})^T V \quad (8)$$

where the Jacobian matrix $J_v(\vec{w})$ is

$$J_v(w) = \begin{bmatrix} \frac{\partial V(1)}{\partial w_1} & \dots & \frac{\partial V(1)}{\partial w_M} \\ \vdots & \ddots & \vdots \\ \frac{\partial V(N)}{\partial w_1} & \dots & \frac{\partial V(N)}{\partial w_M} \end{bmatrix}, V = \begin{bmatrix} V(1) \\ \vdots \\ V(N) \end{bmatrix} \quad (9)$$

Therefore, the weight update using LM for an RNN controller can be expressed as

$$\Delta \vec{w} = -[J_v(\vec{w})^T J_v(\vec{w}) + \mu I]^{-1} J_v(\vec{w})^T V \quad (10)$$

D. FATT Algorithm

In order to calculate the Jacobian matrix $J_v(\vec{w})$ efficiently, FATT was used, which incorporates the procedures of unrolling the system, calculating the derivatives of the Jacobian matrix, and calculating the DP cost into one single process for each training epoch [20]. Fig. 7 illustrates the process of unrolling the trajectory in the forward path, and Alg. 1 specifies FATT [18], where $\vec{\phi}(k) = \sum_{j=1}^k \vec{i}_{dq_s}(j)$ and $\frac{\partial \vec{\phi}_{dq}(k)}{\partial \vec{w}} = \sum_{j=1}^k \frac{\partial \vec{i}_{dq_s}(j)}{\partial \vec{w}}$.

E. Parallel Training Combination of LM and FATT Algorithms

Fig. 8 presents the proposed parallel training combination of LM and FATT algorithms for training an RNN controller. FATT* in Fig. 8 refers to a modified version of Alg. 1 that only calculates the DP cost by eliminating lines 5-9 and 14-15 to save computation time. The most time-consuming parts include the DP cost calculation and the Jacobian matrix for each trajectory, which is conducted by the FATT Algorithms. To solve this challenge, the basic idea is to parallelize them as follows: First, all training trajectories are divided into N groups, each with a size from 1 to a number smaller than the number of total trajectories. Then, the calculation of each group of trajectories is allocated to one Worker or Central Processing Unit (CPU) core. The detailed implementation will depend on the specific programming language, platforms, etc. For example, for the MATLAB implementation, the computing unit will be one MATLAB worker, which corresponds to one CPU core. For the C++ implementation, this single worker could correspond to one CPU core or thread. For a fair comparison, one single CPU core was also used in the C++ implementation. The implementation and comparison of both cases will be detailed in Sections IV and V.

Fig. 8 also illustrates how the algorithm dynamically adjusts μ . When μ increases, training is closer to a gradient descent algorithm with a small learning rate, whereas when μ decreases, training approaches the Gauss-Newton method, which

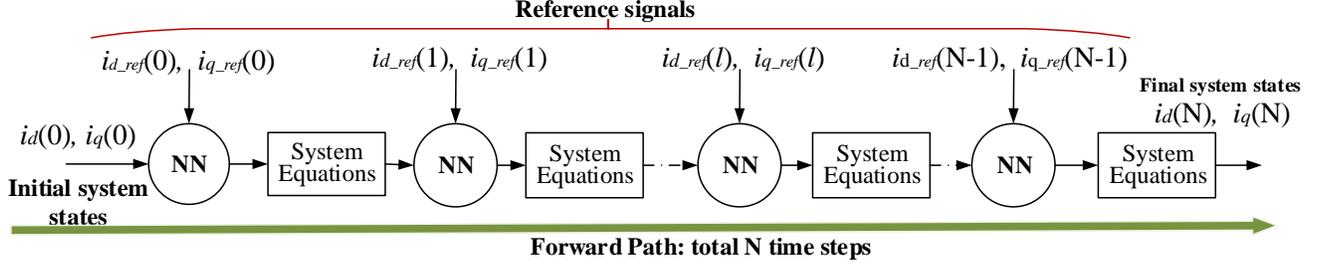


Fig. 7. Unrolling the forward path in the FATT algorithm for training an NN controller in a closed-control loop system.

Algorithm 1 FATT algorithm to calculate the Jacobian matrix and to accumulate DP cost for one trajectory.

- 1: $C \leftarrow 0, \vec{e}_{dq}(0) \leftarrow 0, \vec{s}_{dq}(0) \leftarrow 0, \frac{\partial \vec{i}_{dq}(0)}{\partial \vec{w}} \leftarrow 0, \frac{\partial \vec{\phi}_{dq}(0)}{\partial \vec{w}} \leftarrow 0$
- 2: {Calculate the Jacobian matrix $J_v(\vec{w})$ }
- 3: **for** $k = 0$ to $N - 1$ **do**
- 4: $\vec{v}_{dq1}(k) \leftarrow k_{PWM} R(\vec{e}_{dq}(k), \vec{s}_{dq}(k), \vec{w})$
- 5: $\frac{\partial \vec{s}_{dq}(k)}{\partial \vec{w}} \leftarrow T_s \left[\frac{\partial \vec{\phi}_{dq}(k)}{\partial \vec{w}} - \frac{1}{2} \frac{\partial \vec{i}_{dq}(k)}{\partial \vec{w}} \right]$
- 6: $\frac{\partial \vec{v}_{dq1}(k)}{\partial \vec{w}} \leftarrow k_{PWM} \left[\frac{\partial R(k)}{\partial \vec{w}} + \frac{\partial R(k)}{\partial \vec{e}_{dq}(k)} \frac{\partial \vec{i}_{dq}(k)}{\partial \vec{w}} + \frac{\partial R(k)}{\partial \vec{s}_{dq}(k)} \frac{\partial \vec{s}_{dq}(k)}{\partial \vec{w}} \right]$
- 7: $\frac{\partial \vec{i}_{dq}(k+1)}{\partial \vec{w}} \leftarrow A \frac{\partial \vec{i}_{dq}(k)}{\partial \vec{w}} + B \frac{\partial \vec{v}_{dq1}(k+1)}{\partial \vec{w}}$
- 8: $\frac{\partial \vec{i}_{dq}(k+1)}{\partial \vec{w}} \leftarrow$ the first two terms of $\frac{\partial \vec{i}_{dq}(k+1)}{\partial \vec{w}}$
- 9: $\frac{\partial \vec{\phi}_{dq}(k+1)}{\partial \vec{w}} \leftarrow \frac{\partial \vec{\phi}_{dq}(k)}{\partial \vec{w}} + \frac{\partial \vec{i}_{dq}(k+1)}{\partial \vec{w}}$
- 10: $\vec{i}_{dq}(k+1) \leftarrow A \vec{i}_{dq}(k) + B \vec{v}_{dq1}(k)$
- 11: $\vec{e}_{dq}(k+1) \leftarrow \vec{i}_{dq}(k+1) - \vec{i}_{dq_ref}(k+1)$
- 12: $\vec{s}_{dq}(k+1) \leftarrow \vec{s}_{dq}(k) + \frac{T_s}{2} [\vec{e}_{dq}(k) + \vec{e}_{dq}(k+1)]$
- 13: $C \leftarrow C + U(\vec{e}_{dq}(k+1))$ {accumulate DP cost}
- 14: $\frac{\partial \vec{V}(k+1)}{\partial \vec{w}} \leftarrow \frac{\partial \vec{V}(k+1)}{\partial \vec{e}_{dq}(k+1)} \frac{\partial \vec{i}_{dq}(k+1)}{\partial \vec{w}}$
- 15: the $(k+1)$ th row of $J(\vec{w}) \leftarrow \frac{\partial \vec{V}(k+1)}{\partial \vec{w}}$
- 16: **end for**
- 17: {On exit, the Jacobian matrix $J_v(\vec{w})$ is finished for one trajectory.}

provides faster convergence than gradient descent. There are three stopping conditions used for training: 1) when the training epoch reaches a maximum acceptable value Epoch_{max} ; 2) when μ is larger than μ_{max} ; and 3) when the gradient is smaller than the predefined minimum acceptable value $\|\partial C_{dp} / \partial \vec{w}\|_{min}$.

IV. AMAZON EC2 CLOUD TRAINING PERFORMANCE

A. Amazon EC2 cloud

Amazon EC2 cloud was utilized as the test cloud platform. The cloud cluster was configured and connected to Amazon EC2. For the Amazon EC2 cluster, the region is selected as US East (N. Virginia). The worker machine type is General Purpose (m5.24xlarge, 48 core), which supports 48 workers per machine. The headnode machine type is Standard (c5d.xlarge, 2core,1x100NVMe). It is noted that the m5.24xlarge type

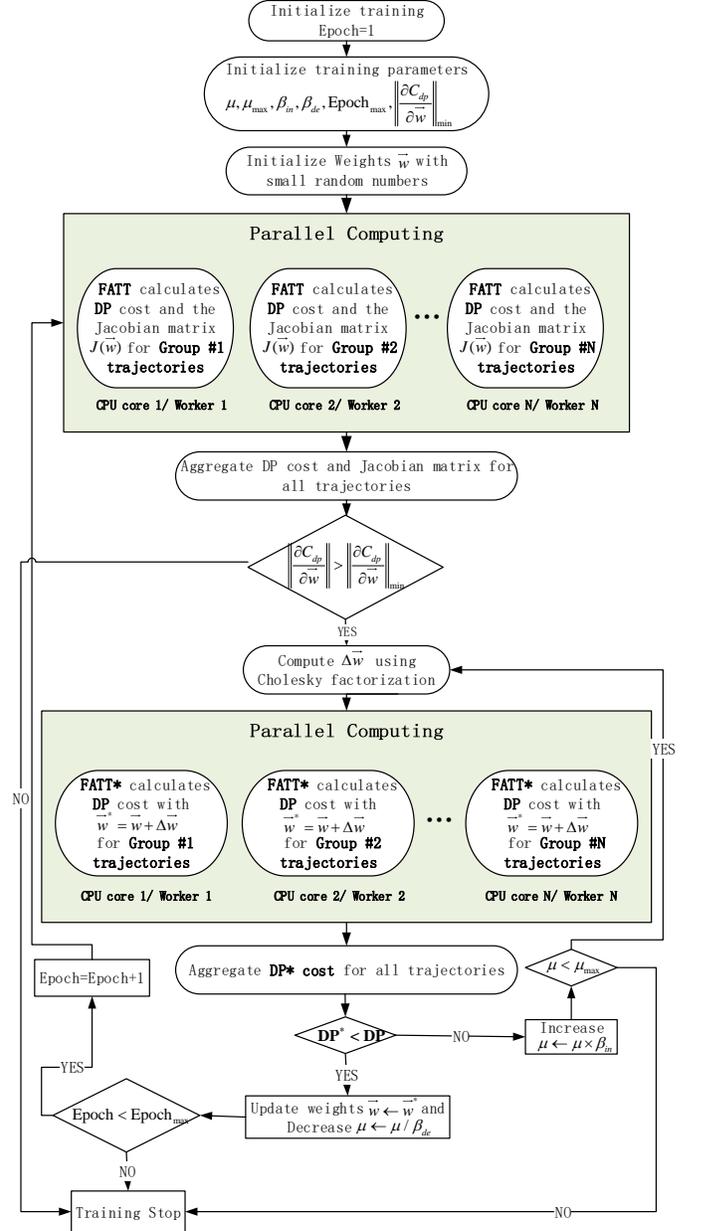


Fig. 8. The proposed parallel LM+FATT trajectory training algorithm for an RNN controller. μ_{max} stands for maximum μ , and β_{de} and β_{in} signify the decreasing and increasing factors, respectively.

instance uses up to 3.1 GHz Intel Xeon Platinum Processor,

however, the specific CPU processor number is not provided by Amazon EC2 [30].

B. MATLAB Implementation and Speedup Performance

Matlab [31] was first used to develop the training program and validate the proposed parallel training algorithm. Matlab Online running version R2022b was used. For the two parallel computing parts described in Fig. 8, two parfor-loop structures were used to execute for-loop iterations in parallel on workers in a parallel pool. To execute the parfor-loop, Matlab starts a parallel pool with one worker per physical CPU core, not CPU thread. A parallel pool with 48 workers was used.

To test whether training is successful or not, the current tracking performance in the closed-loop control was performed. Fig. 9 shows the current tracking performance after the RNN controller was well-trained. The current i_d and i_q can track their reference currents i_{d_ref} and i_{q_ref} very well only with some slight oscillations at the beginning of the reference change.

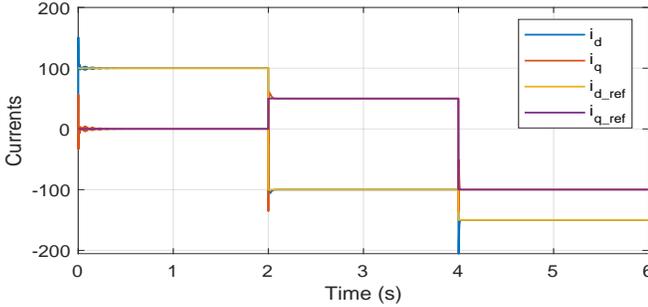


Fig. 9. The tracking performance in a closed-loop after successful RNN training.

Fig. 10 shows the running time comparison on different numbers of Matlab workers and trajectories. The horizontal axis represents the number of CPU cores/Matlab workers, while the vertical axis stands for the average running time. To test the performance of the proposed approach, all inputs across all runs and any randomness inherent in the code were fixed. The test program was run with trajectories in $\{10, 20, 30, \dots, 100\}$ and with Matlab workers from $\{1, 2, 3, \dots, 48\}$. When the number of Matlab workers equals 1, no parallel computing was used. For each running case, the program was repeated 10 times, and the average running time was calculated to remove possible variations of the running time in each run for a fair comparison. Regardless of the number of trajectories used in the experiment, a consistent trend emerges: as the number of CPU cores/Matlab workers increases to a certain threshold, the average running time stabilizes and cannot be further reduced.

The speedup is defined as the ratio of serial execution time over parallel execution time, which is $speedup = \text{Serial Execution Time} / \text{Parallel Execution Time}$. Fig. 11 shows the speedup comparison across different numbers of Matlab workers. The horizontal axis represents the number of CPU cores/Matlab workers and the vertical axis stands for the speedup. Figs. 10 and 11 clearly show an excellent,

although nonlinear speedup performance. When the number of Matlab workers surpasses the number of trajectories, no speedup benefits can be further achieved and the running time is slightly longer due to communication loads between Matlab workers. The trend was observed in all number of trajectories. Compared to nonparallel training, the parallel approach achieves a least 4 times speedup. When the number of trajectories increases, the speedup becomes even more significant.

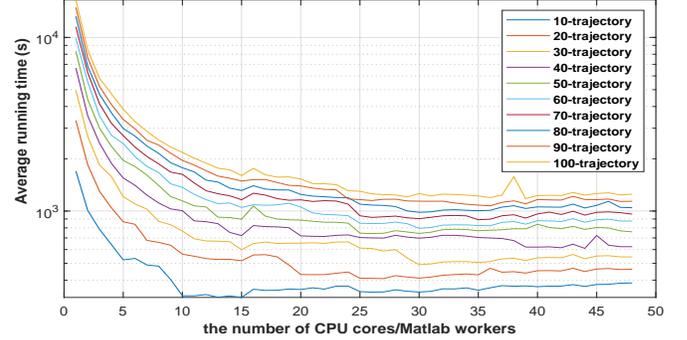


Fig. 10. Average running time across Matlab workers on Amazon EC2 Cloud.

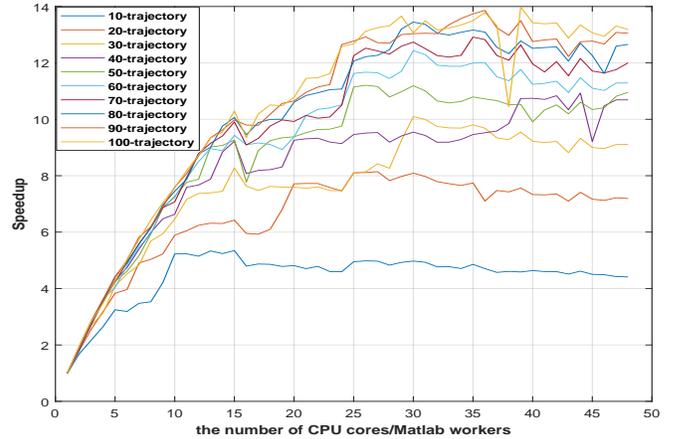


Fig. 11. Speedup across Matlab workers on Amazon EC2 Cloud.

V. HPC CLUSTER TRAINING PERFORMANCE

A. HPC Cluster Platform

The selected HPC platform is a cluster consisting of four Non-Uniform Memory Access (NUMA) system [32] compute nodes running the AlmaLinux distribution of the Linux operating system [33]. The CPU of each compute node is Intel(R) Xeon(R) Platinum 8180M CPU. Each compute node contains 2 sockets, one 28-core CPU per socket, and 2 threads per CPU core, providing a total of $2 * 28 = 56$ CPU cores and $56 * 2 = 112$ threads. See Fig. 12 for an illustration of the HPC architecture.

B. MATLAB Implementation and Speedup Performance

The HPC is configured to use Slurm to manage job scheduling [34]. The Matlab cluster configuration supports up to

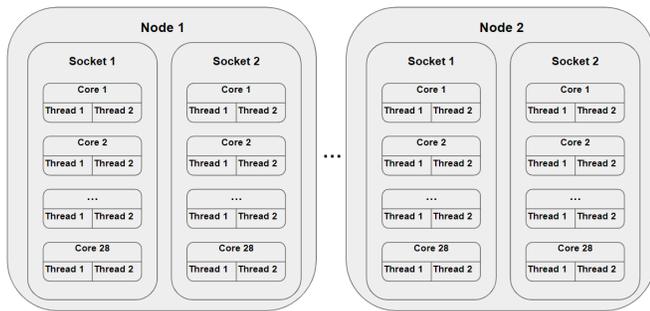


Fig. 12. The schematic of the HPC Cluster architecture.

$28 \times 2 = 56$ computational CPU cores. For a fair comparison with results running on Amazon EC2 Cloud, the maximum number of Matlab workers was also set to 48.

Fig. 13 shows the running time comparison across different numbers of Matlab workers and trajectories. Fig. 14 shows the speedup comparison across different numbers of Matlab workers. Figs. 13 and 14 show slightly better speedup performance running on the HPC cluster than Figs. 10 and 11 running on Amazon EC2 Cloud. For example, the maximum speedup running on the HPC cluster is 6 times compared to only 5 times speedup on the Amazon EC2 Cloud for 10-trajectory training.

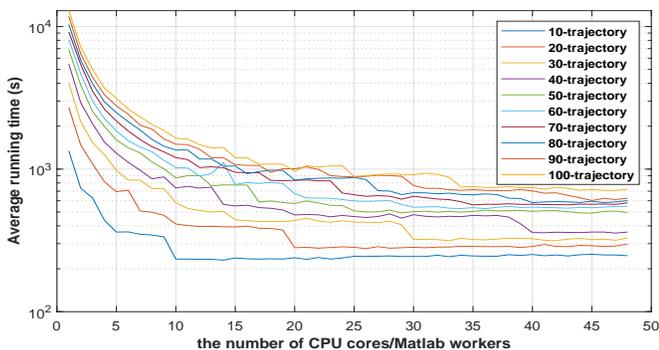


Fig. 13. Average running time across Matlab workers on the HPC cluster.

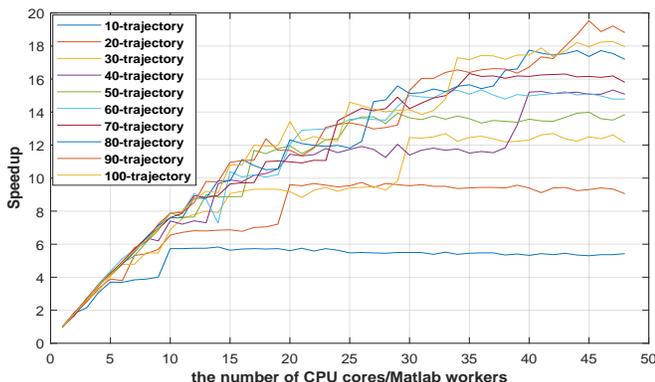


Fig. 14. Speedup across Matlab workers on the HPC cluster.

C. C++ Implementation and Speedup Performance

The c++17 standard [35] and the compiler g++ [36] were used to develop the training program. An important third-party package used in the training program is the C++ package Armadillo, which is a fast linear algebra library to perform linear algebra computations [37], [38]. Armadillo relies on BLAS [39] and Linear Algebra PACKage (LAPACK) [40]. BLAS is a specification of low-level routines for performing basic linear algebra operations. The training program used OpenBLAS, which is an open-source implementation in Fortran of this software specification [41]. The configuration of an implementation of BLAS is paramount since it can be tuned to particular architectures to improve performance. LAPACK is a software library of more complex linear algebra operations also written in Fortran.

The OpenMPI implementation of the Message Passing Interface (MPI) standard [42] was utilized to parallelize the workload among the available worker nodes. An advantage of the Open MPI framework is that programs are scalable to larger and more powerful clusters. To parallelize the workload, the MPI program has a sequential part and parallel parts as illustrated in Fig. 8. The master process in the MPI program handles the sequential parts and sends messages to each worker process to manage its parallel part of the workload.

Fig. 15 shows the running time comparison across different numbers of workers. The horizontal axis represents the number of CPU cores/workers, while the vertical axis stands for the average running time. Fig. 16 shows the speedup comparison on different numbers of workers. The horizontal axis represents the number of CPU cores/workers and the vertical axis stands for the speedup. Compared with Figs. 13 and 14, the parallel version of the C++ program (Figs. 15 and 16) runs faster than the Matlab version on the HPC cluster. For example, the maximum speedup of the C++ version is around 8 times compared to the 6 times speedup of the corresponding Matlab version for 10-trajectory training. Overall, among the three implementations: the C++ version on the HPC cluster, the Matlab version on the HPC cluster, and the Matlab version on the cloud, the C++ version on the HPC cluster achieves the best results. The next one is the Matlab version on the HPC cluster, while the Matlab version on the cloud is the slowest one.

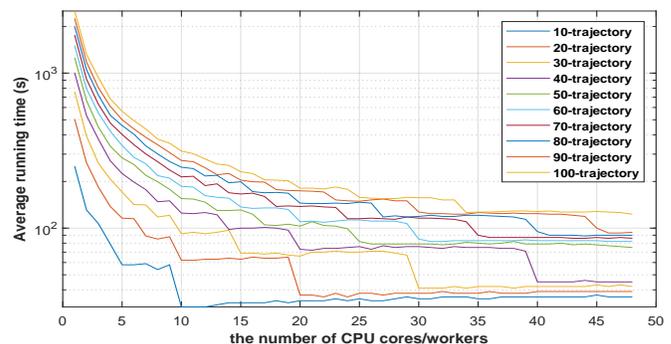


Fig. 15. Average running time across workers on the HPC cluster.

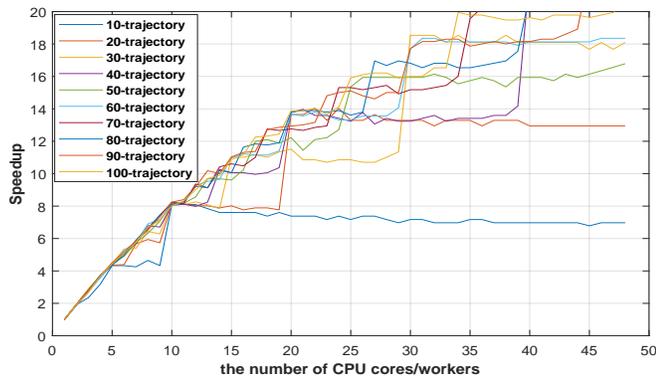


Fig. 16. Speedup across workers on the HPC cluster.

VI. CONCLUSION

This paper investigated how to use parallel computing to accelerate the LM algorithm for training an RNN controller in a closed-loop control system. The proposed parallel trajectory training algorithm incorporates LM and FATT algorithms. The training programs were implemented using both Matlab and C++. The developed programs were tested on two computing platforms, namely the Amazon EC2 Cloud and an HPC cluster. Performance comparison results show that the parallel training algorithm can provide a significant speedup compared to its non-parallelized counterparts. Specifically, the C++ implementation achieves better speedup than the corresponding MATLAB implementation. The program running on HPC clusters can further yield even better speedup than that on cloud platforms. The significant speedup performance makes it suitable to train RNN controllers with a large number of trajectories and long-duration trajectories.

REFERENCES

- [1] K. Levenberg. "A Method for the Solution of Certain Non-Linear Problems in Least Squares". *Quart. Appl. Math.*, 2: 164-168 (1944).
- [2] D.W. Marquardt. "An algorithm for least-squares estimation of nonlinear parameters," *Journal of the Society for Industrial and Applied Mathematics*, 11(2):431-441, 1963.
- [3] M. T. Hagan, H. B. Demuth, and M. H. Beale, *Neural Network Design*, Boston: PWS, 2002, ch.12, pp 19-23.
- [4] L Chan and C. Szeto, "Training Recurrent Network with Block-Diagonal Approximated Levenberg-Marquardt Algorithm", *Proc. of the International Joint Conference on Neural Networks (IJCNN'99)*, Washington, DC, Jul. 10-16, 1999.
- [5] Y. Kwak. "An accelerated Levenberg-Marquardt algorithm for feedforward network," *Journal of the Korean Data & Information Science Society*, 23(5), 1027-1035, 2012.
- [6] T. Tawara, "Levenberg-Marquardt with Sparse Block Matrices on the GPU". [Online]. Available: <https://on-demand.gputechconf.com/gtc/2012/presentations/S0231-Levenberg-Marquardt-Using-Block-Sparse-Matrices-on-CUDA.pdf>
- [7] N. N. R. Ranga Suri, D. Deodhare, and P. Nagabhushan, "Parallel Levenberg-Marquardt-based Neural Network Training on Linux Clusters - A Case Study", *Proc. of 3rd Indian conference on computer vision, graphics and image processing (ICVGIP 2002)*, N.P.: Allied Publishers.
- [8] J. Cao, K. A. Novstrup, A. Goyal, S. P. Midkiff, and J. M. Caruthers, "A parallel levenberg-marquardt algorithm", *ICS '09: Proc. of the 23rd international conference on Supercomputing* Jun. 2009 PP 450-459, <https://doi.org/10.1145/1542275.1542338>.
- [9] A. Przybylski, B. Thiel, J. Keller-Findeisen, B. Stock and M. Bates, "Gpufit: An open-source toolkit for GPU-accelerated curve fitting". *Scientific Reports* 7, 15722 (2017). <https://doi.org/10.1038/s41598-017-15313-9>.
- [10] X. Zhu and D. Zhang, "Efficient Parallel Levenberg-Marquardt Model Fitting towards Real-Time Automated Parametric Imaging Microscopy", *PLOS ONE* 8(10): e76665, 2013. <https://doi.org/10.1371/journal.pone.0076665>
- [11] Y. Lin, D. O'Malley, and V. V. Vesselinov, "A computationally efficient parallel Levenberg-Marquardt algorithm for highly parameterized inverse model analyses", *Water Resources Research*, Vol. 52, Issue 9, Sep. 2016, PP 6948-6977, <https://doi.org/10.1002/2016WR019028>
- [12] S. Li, M. Fairbank, D. C. Wunsch, and E. Alonso, "Vector Control of a Grid-Connected Rectifier/Inverter Using an Artificial Neural Network," *Proc. of 2012 International Joint Conference on Neural Network*, Brisbane Australia, Jun. 2012.
- [13] R. J. Williams and D. Zipser, "A learning algorithm for continually running fully recurrent neural networks," *Neural Computation*, Vol.1 Issue 2, Summer 1989, pp. 270-280.
- [14] H. Jaeger, "Tutorial on training recurrent neural networks, covering BPPT, RTRL, EKF and the "echo state network" approach," *GMD Report 159*, German National Research Center for Information Technology, 2002.
- [15] F. Gomez, J. Schmidhuber, R. Miikkulainen, and M. Mitchell, 2008. "Accelerated Neural Evolution through Cooperatively Coevolved Synapses". *Journal of Machine Learning Research*, 9 (5).
- [16] Ma, Sheng, and Chuanyi Ji. "A unified approach on fast training of feedforward and recurrent networks using EM algorithm." *IEEE Trans. Signal Process.* 46, no. 8 (1998): 2270-2274.
- [17] S.K. Ng and G.J. McLachlan, 2004. Using the EM algorithm to train neural networks: misconceptions and a new algorithm for multiclass classification. *IEEE Trans. Neural Netw.*, 15(3), pp.738-749.
- [18] X. Fu, S. Li, and I. Jaithwa, "Implement optimal vector control for LCL-filter-based grid-connected converters by using recurrent neural networks," *IEEE Trans. Ind. Electron.*, vol. 62, no. 7, pp. 4443-4454, Jul. 2015.
- [19] Y.Tanoto, W. Ongsakul, and C. O.P. Marpaung, "Levenberg-Marquardt Recurrent Networks for Long Term Electricity Peak Load Forecasting," *TELKOMNIKA (Telecommunication Computing Electronics and Control)*, Vol 9, No.2, August 2011, pp 257-266.
- [20] X. Fu, S. Li, M. Fairbank, D. C. Wunsch, and E. Alonso, "Training recurrent neural networks with the Levenberg-Marquardt algorithm for optimal control of a grid-connected converter", *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 26, no. 9, pp. 1900-1912, Sep. 2015.
- [21] Texas Instruments, "Digitally controlled solar micro inverter design using C2000 Piccolo microcontroller user's guide". [Online]. Available: <https://www.ti.com/lit/ug/tidu405b/tidu405b.pdf>
- [22] W. Wu, Y. Sun, Z. Lin, T. Tang, F. Blaabjerg, and H. S. Chung, "A new LCL-filter with in-series parallel resonant circuit for single-phase grid-tied inverter," *IEEE Trans. Ind. Electron.*, vol. 61, no. 9, pp. 4640-4644, Sep. 2014.
- [23] R. Teodorescu, M. Liserre, and P. Rodriguez, *Grid Converters for Photovoltaic and Wind Power Systems*. Chichester, West Sussex, United Kingdom: John Wiley and Sons, 2011.
- [24] J. Sturtz, X. Fu, C. D. Hingu and L. Qingge, "A Novel Weight Dropout Approach to Accelerate the Neural Network Controller Embedded Implementation on FPGA for a Solar Inverter," *2023 IEEE International Conference on Smart Computing (SMARTCOMP)*, Nashville, TN, USA, 2023, pp. 157-163.
- [25] W. Waithaka, X. Fu, A. Hadi, R. Chaloo, and S. Li, "DSP Implementation of a Novel Recurrent Neural Network Controller into a TI Solar Microinverter", *Proc. of 2021 IEEE PES General Meeting*, Jul. 26 - Jul. 29, 2021
- [26] X. Fu, S. Li, D. C. Wunsch, and E. Alonso, "Local Stability and Convergence Analysis of Neural Network Controllers With Error Integral Inputs," *IEEE Trans. Neural Netw. Learn. Syst.* (Early Access), 14 Oct. 2021, <https://ieeexplore.ieee.org/abstract/document/9573504>.
- [27] F. Wang, H. Zhang, and D. Liu, "Adaptive dynamic programming: An introduction," *IEEE Comput. Intell. Mag.*, vol. 4, no. 2, pp. 39-47, May 2009.
- [28] R. E. Bellman, *Dynamic Programming*. Princeton, New Jersey, USA: Princeton University Press, 1957.
- [29] D. V. Prokhorov and D. C. Wunsch, "Adaptive critic designs," *IEEE Trans. Neural Netw.*, vol. 8, no. 5, pp. 997-1007, Sep. 1997.
- [30] Amazon EC2 instance Types. [Online]. Available: <https://aws.amazon.com/ec2/instance-types/>
- [31] Matlab. Available: <https://www.mathworks.com/products/matlab.html>
- [32] C. Lameter, "NUMA (Non-Uniform Memory Access): An Overview: NUMA becomes more common because memory controllers get close to execution units on microprocessors", *Queue*, vol. 11, no. 7, pp. 40-51, July 2013.

- [33] AlmaLinux OS Foundation, Almalinux. [Online]. Available: <https://almalinux.org/>
- [34] Slurm: Workload Manager. [Online]. Available: <https://slurm.schedmd.com/overview.html>
- [35] C++17. [Online]. Available: <https://en.cppreference.com/w/cpp/17>.
- [36] Free Software Foundation, Inc. *g++(1) - Linux man page*. [Online]. Available: <https://linux.die.net/man/1/g++>.
- [37] C. Sanderson, and R. Curtin. "Armadillo: a template-based C++ library for linear algebra," *Journal of Open Source Software*, vol. 1, pp. 26, 2016.
- [38] C. Sanderson, and R. Curtin. "A User-Friendly Hybrid Sparse Matrix Class in C++," *Lecture Notes in Computer Science (LNCS)*, vol. 10931, pp. 422–430, , 2018.
- [39] BLAS. [Online]. Available: <https://netlib.org/blas/>.
- [40] LAPACK. [Online]. Available: <https://netlib.org/lapack/>.
- [41] OpenBLAS. [Online]. Available: <https://github.com/xianyi/OpenBLAS>.
- [42] Open MPI: Open Source High Performance Computing. [Online]. Available: <https://www.open-mpi.org/>.