

# Layer-Parallel Training with GPU Concurrency of Deep Residual Neural Networks via Nonlinear Multigrid

Andrew Kirby, Siddharth Samsi, Michael Jones, Albert Reuther, Jeremy Kepner, Vijay Gadepally  
Massachusetts Institute of Technology Lincoln Laboratory Supercomputing Center  
Lexington, MA, USA

**Abstract**—A Multigrid Full Approximation Storage algorithm for solving Deep Residual Networks is developed to enable neural network parallelized layer-wise training and concurrent computational kernel execution on GPUs. This work demonstrates a 10.2x speedup over traditional layer-wise model parallelism techniques using the same number of compute units.

**Index Terms**—Deep learning, residual networks, multigrid.

## I. INTRODUCTION

Improvements in computing power have driven resurgence and significant advancements in Machine Learning and, specifically, Deep Learning [1], [2]. To realize these improvements, the transition from the Dennard-scaling era of increasing CPU clock frequencies through the miniaturization of transistors to the Multicore era [3] has compelled parallel algorithmic development. For Supervised Learning, the primary parallel methodology employed is data parallelism which trains neural network models by evaluating multiple data samples simultaneously on separate hardware components, computing their respective model parameter gradients with respect to the input data sample, and collectively averaging the gradients which are then used to update the model parameters.

However, in the post-Moore's-Law era composed of increasingly parallel hardware (within and between compute units), data parallelism suffers from global collective communication patterns, which limit the long-term performance capabilities. Further, data parallelism does not address model size limitations or sequential evaluation of the neural network using forward or backward propagation. An orthogonal approach, which can be paired with data parallelism, to address model-size limitation is model parallelism, which partitions a model across multiple compute units. However, traditional model parallelism approaches also fail to address the sequential evaluation of neural networks, thereby being constrained in performance by Amdahl's Law<sup>1</sup>. To circumvent Amdahl's

Research was sponsored by the United States Air Force Research Laboratory and was accomplished under Cooperative Agreement Number FA8750-19-2-1000. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the United States Air Force or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

<sup>1</sup>Amdahl's Law states the maximum speedup of an algorithm is limited by the fraction of sequential procedure in a program.

Law, we develop an asynchronous layer-parallel algorithm for Graphical Processing Units (GPU) enabling constant asymptotic execution time independent of the number of neural network layers.

Beyond hardware trends, a number of applications such as image classification, video recognition, natural language processing have relied on increasingly deep neural networks. For example, Figure 1 describes winning entries from the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) [4]. In this and other similar applications, deeper networks can often provide greater performance (e.g., accuracy). However, training such networks requires significant increase in computational power required. For example, the authors of [5] indicate a 3.4 month doubling time in computing requirements. A significant number of deep neural network architectures fall into the class of deep residual networks.

## II. DEEP RESIDUAL NETWORKS

Deep Residual Networks have been demonstrated to have superior performance for supervised learning techniques applied to real-world applications such as image classification problems [6] and various other technologies [7], [8], and have resolved network training issues due to *vanishing gradients* by enacting a skip connection allowing gradient information to accumulate and bypass multiplicative operators.

Residual neural networks (ResNet) are composed of residual blocks stacked in  $N$ -layers of the following structure:

$$\mathbf{u}^{n+1} = \mathbf{u}^n + hF(\mathbf{u}^n; \boldsymbol{\theta}^n), \quad \text{for } n = 0, \dots, N-1 \quad (1)$$

with  $\mathbf{u}, \boldsymbol{\theta} \in \mathbb{R}^q$  being the layer state and parameters, respectively, and  $q$  is the layer width,  $F$  is the feature transformation operator (e.g. activation of convolution and bias), and  $h$  is a scaling factor.

A ResNet can be naturally described mathematically as an Initial Value Problem (IVP) such that

$$\frac{\partial \mathbf{u}(t)}{\partial t} = F(\mathbf{u}(t); \boldsymbol{\theta}(t)), \quad \forall t \in T \quad (2)$$

where  $T \subseteq \mathbb{R}^+ \cup \{0\}$ . Equation (2) can be discretized by the explicit Forward Euler Method [9], [10] leading to the formulation in (1) where  $h$  represents the discrete time step size.

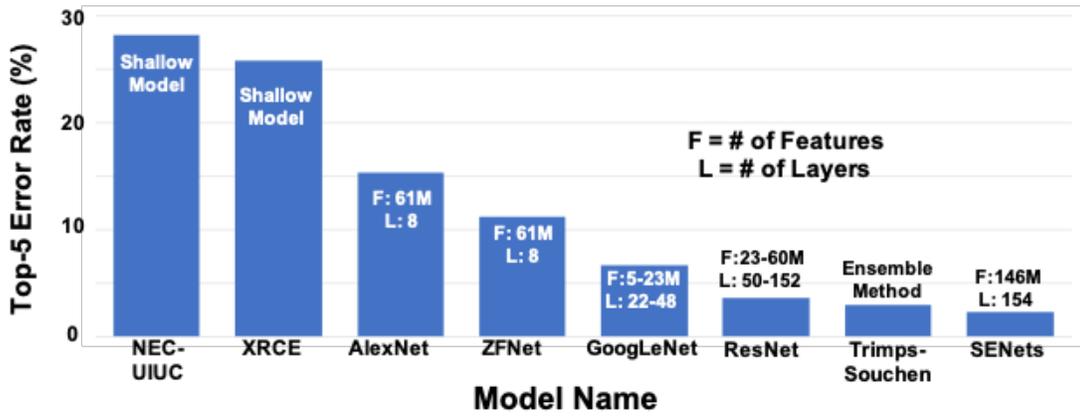


Fig. 1. Analysis of ILSVRC Winners from 2010-2017. The y-axis corresponds to Top-5 error rate (%) of winning entries. An observable trend in image classification and other applications is the correlation between performance and increased network depth.

### III. MULTIGRID METHODS

Multigrid (MG) methods are iterative algorithms that accelerate the solution convergence of a system of equations. The technique is traditionally applied in scientific and engineering applications such as problems found in Computational Fluid Dynamics [11], [12], and, more recently, in Deep Learning including the acceleration of training optimization [13] and layer parallelism on CPUs [14]. This work follows [14] and extends the algorithm to GPUs and enables concurrent kernel execution through NVIDIA's CUDA Streams [15] and Deep Neural Network (CuDNN) [16].

#### A. Multigrid Principles

The principle idea of Multigrid is to use multiple *levels* of resolution to iteratively solve an error prescribed by the residual equation governed by the system. Suppose we wish to solve for the unknown  $u$  in the following:

$$Lu = f \quad (3)$$

$$r := f - L\bar{u} \quad (4)$$

$L$  is a functional (linear or nonlinear) operating on the solution  $u$  which is equal to a known source  $f$ . Given an approximate solution  $\bar{u}$ , the residual  $r$  is defined by (4); note that if  $\bar{u} = u$ , then  $r = 0$ . The construction of the MG algorithm is independent of the set of equations or discretization setting. These principles will be used to demonstrate the construction of multigrid algorithm for solving the neural network states via forward propagation. Consider the solution of the discrete problem:

$$L_h u_h = f_h \quad (5)$$

where the subscripts refer to the discretization of the continuous formulation (3) on a *level* of step size  $h$ . Let  $\bar{u}_h$  be the current estimate of  $u_h$  obtained by approximately solving (3). Since  $\bar{u}_h$  does not satisfy the equation exactly, then the non-zero residual  $r_h$  is given as:

$$r_h = f_h - L_h \bar{u}_h \quad (6)$$

The objective of Multigrid is to solve for the solution correction  $\delta u_h$  such that the exact solution is given by:

$$u_h = \bar{u}_h + \delta u_h \quad (7)$$

Subtracting (5) from (6) yields the following:

$$L_h u_h - L_h \bar{u}_h = r_h \quad (8)$$

If  $L_h$  is a linear operator, then substituting (7) into (8) results in the correction equation:

$$L_h \delta u_h = r_h \quad (9)$$

Assuming the errors on level  $h$  are high-frequency, the correction  $\delta u_h$  can be computed more efficiently on a coarser level  $H$  by solving the equation

$$L_H \delta u_H = I_h^H r_h \quad (10)$$

where  $I_h^H$  denotes a restriction operator which interpolates the residual from the fine level  $h$  to the coarse level  $H$ . If (10) contains low computational work, it may be solved exactly by direct or iterative methods. Alternatively, the above procedure may be applied recursively on coarser levels to yield an approximate solution to the correction  $\delta u_H$  itself. Once the correction is computed, an update is employed on the fine level to correct the solution:

$$\bar{u}_h^{new} = \bar{u}_h^{old} + I_H^h \delta u_H \quad (11)$$

where  $I_H^h$  is a prolongation operator which interpolates the coarse level correction to the fine level. Following the fine-level update, the solution may be solved approximately again. This process is known as relaxation which may be repeated until an acceptable level of error tolerance is achieved. Note that when the fine level solution is achieved, then  $r_h = 0$  and (9) is solved trivially  $\delta u_H = 0$ . Thus no further corrections to the fine level solution are produced.

### B. Full Approximation Storage (FAS)

Suppose the operator in (8) is nonlinear, then the equation can no longer be replaced by  $L_h \delta u_h$ , thus modifications are needed to accommodate this nonlinear behavior. This is done by introducing a new coarse level variable  $\bar{u}_H$ :

$$\bar{u}_H = \bar{I}_h^H \bar{u}_h + \delta u_H \quad (12)$$

where  $\bar{I}_h^H$  represents the restriction operator which transfers the fine level solution to the coarse level. The analogue nonlinear equation to (10) is as follows:

$$L_H \bar{u}_H - L_H \bar{I}_h^H \bar{u}_h = I_h^H r_h \quad (13)$$

Recall the middle and right terms of (13) are known quantities leaving  $\bar{u}_H$  as the only unknown variable. Thus it is useful to introduce a new variable  $S_H$  such that

$$L_H \bar{u}_H = S_H \quad (14)$$

$$\text{where } S_H := L_H \bar{I}_h^H \bar{u}_h + I_h^H r_h \quad (15)$$

On the coarse level, the set of equations defined by (14) are analogous to (5) so similar solution techniques can be applied to solve for the coarse level correction. Lastly, once the coarse level correction  $\bar{u}_H^{new}$  is solved by (14), the fine level solution variables are updated as follows:

$$\bar{u}_h^{new} = \bar{u}_h^{old} + I_h^H \left( \bar{u}_H^{new} - \bar{I}_h^H \bar{u}_h^{old} \right) \quad (16)$$

which, utilizing (12), is written finally as

$$\bar{u}_h^{new} = \bar{u}_h^{old} + I_h^H \delta u_H \quad (17)$$

### C. Multigrid for Residual Network Forward Propagation

Residual Networks correspond to a time-discretized Ordinary Differential Equation IVP, and our goal is to remove the serial propagation of information by forward and backward propagation through network layers. We apply the iterative FAS MG scheme to the layers of the network. The application of Multigrid to the time dimension in Computational Physics is known as the Multigrid Reduction In Time (MGRIT) algorithm [17], [18] and has been demonstrated to achieve significant speedups when sufficient computational resources are utilized [19]–[21]. This algorithm allows the layers to be partitioned into local blocks, which can be solved concurrently and in parallel, thus, breaking Amdahl's Law limitation within the training process of neural networks. We will cast Residual Networks as the nonlinear system in (3) to formulate the MG approach.

Let  $\bar{u}_h$  denote the approximate layer state to the true solution  $\mathbf{u}_h$ , and let  $\mathbf{U}_h = (\mathbf{u}_h^0, \dots, \mathbf{u}_h^{N-1})$  denote the array of layer states in the neural network composed of  $N$  residual blocks on the *level* with step size  $h$ . The forward propagation algorithm corresponding to (1) and (3) can be written as:

$$\underbrace{\begin{bmatrix} \mathbf{u}_h^0 \\ \mathbf{u}_h^1 - hF(\mathbf{u}_h^0, \boldsymbol{\theta}_h^0) \\ \vdots \\ \mathbf{u}_h^{N-1} - hF(\mathbf{u}_h^{N-2}, \boldsymbol{\theta}_h^{N-2}) \end{bmatrix}}_{L_h(\mathbf{U}_h, \boldsymbol{\theta}_h)} = \underbrace{\begin{bmatrix} \mathbf{F}_{in} \mathbf{y} \\ 0 \\ \vdots \\ 0 \end{bmatrix}}_{\mathbf{f}_h} \quad (18)$$

where  $\mathbf{F}_{in}$  corresponds to the network operator and  $\mathbf{y}$  representing the data sample. Let  $\bar{\mathbf{U}}_h$  denote the approximation of  $\mathbf{U}_h$  such that  $\mathbf{U}_h = \bar{\mathbf{U}}_h + \delta \mathbf{U}_h$  and define the notation  $L_h(\bar{\mathbf{U}}, \boldsymbol{\theta}) := L(\bar{\mathbf{U}}_h, \boldsymbol{\theta}_h)$ . We express the solution error via the residual equation corresponding the level with step size  $h$ :

$$\mathbf{R}_h := \mathbf{f}_h - L_h(\bar{\mathbf{U}}, \boldsymbol{\theta}) \quad (19)$$

$$= L_h(\mathbf{U}, \boldsymbol{\theta}) - L_h(\bar{\mathbf{U}}, \boldsymbol{\theta}) \quad (20)$$

$$= L_h(\bar{\mathbf{U}} + \delta \mathbf{U}, \boldsymbol{\theta}) - L_h(\bar{\mathbf{U}}, \boldsymbol{\theta}) \quad (21)$$

Equation (21) is rearranged and a variable substitution is applied such that the correction equation is given as follows:

$$L_h(\mathbf{V}, \boldsymbol{\theta}) = L_h(\bar{\mathbf{U}}, \boldsymbol{\theta}) + \mathbf{R}_h \quad (22)$$

where  $\mathbf{V} = \bar{\mathbf{U}} + \delta \mathbf{U}$ . The residual equation in (22) can be solved on a coarse level with step size  $H$  for  $\mathbf{V}$  resulting in the solution correction  $\delta \mathbf{U} = \mathbf{V} - \bar{\mathbf{U}}$ .

To construct the multilevel hierarchy of the neural network, we define a coarsening factor  $c$  such that every  $c^{th}$  layer within the level is additionally assigned as a member to the coarser level as demonstrated in Figure 2. The coarsening process can be applied repeatedly to make successive coarser levels containing fewer layers. To solve (22) on the coarse level  $H$  for Residual Networks, the fine level residual  $\mathbf{R}_h$ , approximate solution  $\bar{\mathbf{U}}_h$ , and network parameters  $\boldsymbol{\theta}_h$  are restricted to the coarse level by injection (copy) for every  $c^{th}$ -layer. These are denoted as C-Layers as displayed in Figure 2. That is, the restricted solution on the coarse level is given as follows:

$$\bar{\mathbf{U}}_H = \left( \bar{u}_H^0, \bar{u}_H^1, \dots, \bar{u}_H^{N_H-1} \right), \text{ where } \bar{u}_H^n = \bar{u}_h^{nc} \quad (23)$$

The coarse level analogous expression of (22) reads

$$L_H(\mathbf{V}_H, \boldsymbol{\theta}_H) = L_H(\bar{\mathbf{U}}_H, \boldsymbol{\theta}_H) + \mathbf{R}_H =: \mathbf{S}_H \quad (24)$$

and we define the coarse-level nonlinear operator  $L_H$  as

$$L_H(\mathbf{U}_H, \boldsymbol{\theta}_H) = \begin{bmatrix} \mathbf{u}_H^0 \\ \mathbf{u}_H^1 - HF(\mathbf{u}_H^0, \boldsymbol{\theta}_H^0) \\ \vdots \\ \mathbf{u}_H^{N_H-1} - HF(\mathbf{u}_H^{N_H-2}, \boldsymbol{\theta}_H^{N_H-2}) \end{bmatrix} \quad (25)$$

where the coarse level feature-transformation operator  $F_H$  is scaled by  $H = ch$ . This is used to approximate the information transfer between C-Layers. In a two-level MG method,  $\mathbf{V}_H$  is solved directly in (24) providing the coarse-level error  $\delta \mathbf{U}_H = \mathbf{V}_H - \bar{\mathbf{U}}_H$ . The fine-level solutions in the C-Layers are updated via  $\bar{u}_h^{nc} \leftarrow \bar{u}_h^{nc} + \delta \mathbf{u}_H^n$ .

Lastly, as shown in Figure 3, the F-Layers within a layer block are updated directly via F-relaxation using sequential forward propagation. The transfer of information between layer blocks and partitions is achieved by C-relaxation via sequential forward propagation using the last F-Layer in the preceding layer block.

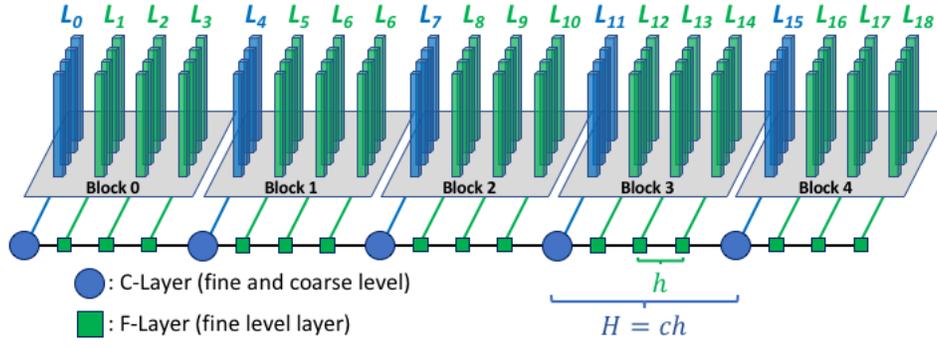


Fig. 2. Neural network layer partitioning with coarsening factor  $c = 4$ . Each layer block is composed of a  $c$  layers.

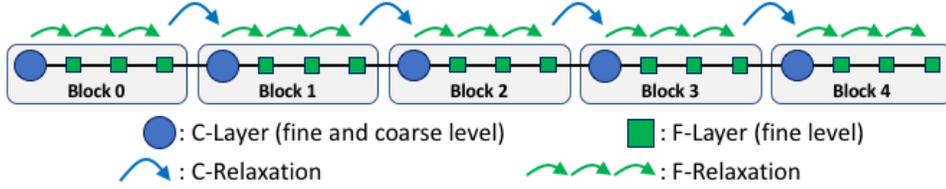


Fig. 3. C-relaxation: forward propagation from preceding F-Layer to C-Layer. F-relaxation: forward propagation from C-Layer to F-Layers in block.

#### Algorithm 1: ResNet Multigrid FAS Scheme

1. FCF-Relaxation to  $L_h(\mathbf{U}_h, \boldsymbol{\theta}_h) = f_h$
2. Restrict residual  $\mathbf{R}_h$  and approximate solution  $\bar{\mathbf{U}}_h$ :  
 $\mathbf{R}_H^n \leftarrow \mathbf{R}_h^n, \bar{\mathbf{U}}_H^n \leftarrow \bar{\mathbf{U}}_h^{nc}$  for  $n = 0, \dots, N_H - 1$
3. Solve  $L_H(\mathbf{V}_H, \boldsymbol{\theta}_H) = L_H(\bar{\mathbf{U}}_H, \boldsymbol{\theta}_H) + \mathbf{R}_H$
4. Solve coarse-level error:  $\delta \mathbf{U}_H = \mathbf{V}_H - \bar{\mathbf{U}}_H$
5. Update fine-level approximate solution:  
 $\bar{\mathbf{u}}_h^{nc} \leftarrow \bar{\mathbf{u}}_h^{nc} + \delta \mathbf{u}_H^n$  for  $n = 0, \dots, N_H - 1$
6. Check convergence:  
**if**  $\|\mathbf{R}_h\| \leq \text{tol}$  **then**  
| Return;  
**end**  
Proceed to step 1;

#### D. Implementation

Algorithm 1 outlines the solution procedure using the MG FAS scheme for a two-level network approach. The implementation of the algorithm for forward-propagation was performed in C++ wrapping NVIDIA’s CuDNN [16] library kernels. CuDNN’s API allows the use of CUDA Streams enabling asynchronous and concurrent execution of kernels. Each layer block shown in Figure 3 is assigned a unique CUDA Stream allowing threads blocks to be solved concurrently. To achieve concurrent execution, kernels are launched in parallel using OpenMP assigning one CPU thread per layer block. Additionally, the usage of multiple GPUs for model-parallelism is implemented using the Message Passing Interface (MPI) [22]. The layer blocks are distributed into contiguous model partitions across GPUs, and distributed memory communication is performed during the C-relaxation phase for abutted layer blocks residing on different GPUs.

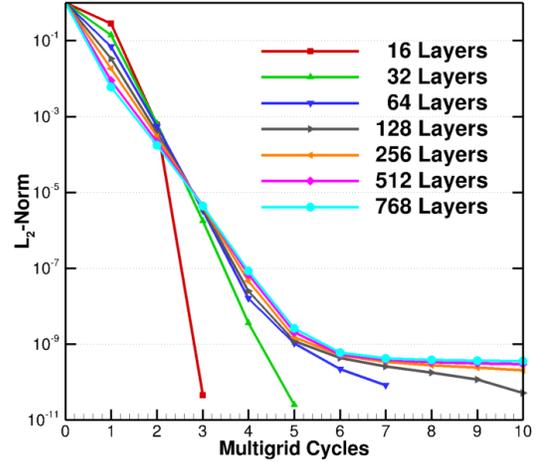


Fig. 4. Residual convergence rates for varying network layer depths demonstrating network layer-size independence.

#### IV. NUMERICAL EXPERIMENTS

Herein, we demonstrate the algorithm using the MIT Supercloud supercomputer TX-GAIA [23], which is composed of 448 compute nodes, each node contains two 2.5GHz Intel Xeon Gold 6248 (20-core) processors with 384 GB of RAM. Every node contains two NVIDIA Tesla 32GB V100 GPUs, both connected to the first CPU only. These results utilize a 25 Gb/s Ethernet interconnect using Mellanox ConnectX-5 adapters connected to a single non-blocking Arista DCS-7516 Ethernet core switch, noting that NVLink is not present in the system. Note that Multigrid is independent of architecture, and can be employed on legacy hardware such as CPUs or FPGAs. MG can be used in conjunction with data-parallelism for the training of neural networks, but we restrict our experiments

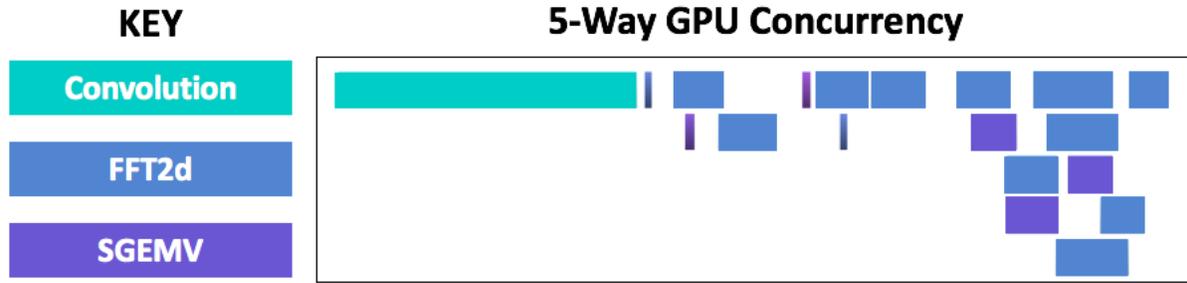


Fig. 5. Compute execution time line within one NVIDIA V100 GPU demonstrating kernel concurrency.

to exclusively explore performance metrics of the algorithm as a model-parallelism approach. We utilize the MNIST [24] data set for testing, which is composed of 60,000 hand-written digits encoded as  $28 \times 28$  pixel single channel (grey-scale) images.

#### A. Verification and Layer-Independent Convergence

To ensure algorithmic correctness, we perform a convergence study on residual networks of various layer-depth counts. Figure 4 demonstrates the  $L_2$ -norm convergence histories of the residual vector  $\mathbf{R}_h$  found in (19). Note that as more Multigrid Cycles are utilized,  $\|\mathbf{R}_h\|_{L_2} \rightarrow 0$  implies  $\bar{\mathbf{U}} \rightarrow \mathbf{U}$  meaning the approximate solution converges to the analytic solution normally obtained by sequential forward propagation.

Further, Figure 4 shows the norm values for all networks composed of different layer counts converge with similar behavior down to  $\|\mathbf{R}_h\|_{L_2} \leq 10^{-9}$ . This indicates an important property that the solution convergence is independent of the network depth. Thus, given enough computational resources, this property will provide approximate constant time to solution evaluation of arbitrarily-deep networks.

Secondarily, an approach known as *early stopping* from *One-Shot* methods [25] can be applied such that the solution convergence is stopped after a fixed number of iterations. This provides an approximate solution which can be used to generate approximate neural network gradients for the training optimization process. For this work, we find that **two** cycles suffice, providing accurate state estimates resulting in approximately the same Top-1 error rates after each epoch of training for the MNIST data set.

#### B. GPU Concurrency

The aforementioned implementation utilizes NVIDIA's Cuda Streams to enable concurrent kernel execution. Using NVIDIA's Profiler (nvprof), Figure 5 shows an excerpt of the kernel execution time line for a single GPU computing a MG cycle composed of convolution layers. As seen in the figure, 5-way kernel concurrency is achieved. However, the number of registers within the GPU prevents multiple convolution kernels from executing simultaneously, thus limiting the performance for the exposed parallelism. At present, the concurrent execution capability is not fully realized due to this limitation but will allow for higher throughput in future architectures.

#### C. Multi-GPU Parallel Inference

Next, we test single image inference via the Multigrid Layer-Parallelism approach on a 4,096 layer Residual Neural Network containing 3,248,524 parameters. The network is composed of an opening convolution layer with a  $7 \times 7$  kernel, 50 output channels, one layer of padding, single striding, and a Rectified Linear Unit (ReLU) activation function. This layer is followed by 4,092 residual convolution layers of the same dimensions and activation functions, and terminated by a fully connected layer mapping the output to a 10-element softmax function paired with cross-entropy loss.

Figure 6a shows the performance profile comparison of the Multigrid algorithm compared to the serial forward propagation approach. When using a single GPU, Multigrid is approximately four times slower; this is expected as MG is an iterative method requiring multiple evaluations of the network. However, as we distribute the model over multiple GPUs, particularly greater than three GPUs, the inference time is reduced achieving a speedup over the sequential approach. Using four GPUs, MG is 1.25x faster and when using 24 GPUs, it is 4x faster. We note that as the model is partitioned more finely, more distributed memory communication is introduced causing the performance to be limited. This is particularly true for heterogeneous architectures composed of CPUs as hosts and GPUs as accelerators.

#### D. Multi-GPU Parallel Training: 3 Million Parameters

To study the effect of communication, we examine the training phase of the same neural network used for the inference study. Figure 6b shows the performance comparisons of the serial method, Multigrid, and the traditional sequential partitioned model (PM) layer-parallel method<sup>2</sup>. When four or more GPUs are utilized, MG is the superior algorithm providing up to 3.5x speedup over serial and 5x speedup over the PM method. Figure 6c demonstrates the timing decomposition elucidating the effect of communication on parallel strong scaling. Computation is evenly distributed across each compute device as more devices are added, and the parallel execution in log-log scale shows near-perfect parallel strong scaling. In contrast, as more GPUs are used to partition

<sup>2</sup>The model is partitioned, layer-wise, over multiple GPUs; each GPU contains a subset of the network's layers, and the evaluation of the network is serialized across the GPUs.

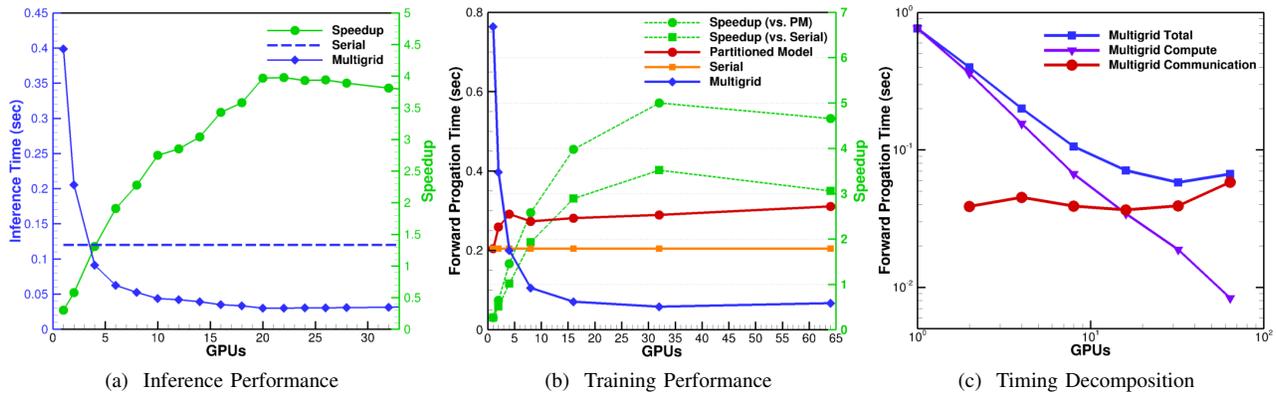


Fig. 6. Parallel strong scaling of forward propagation for a 4,096-layer Residual Neural Network composed of 3,248,534 parameters.

the network, the communication becomes the performance bottleneck. In the case of 64 GPUs, 97% of the evaluation time is consumed by the communication.

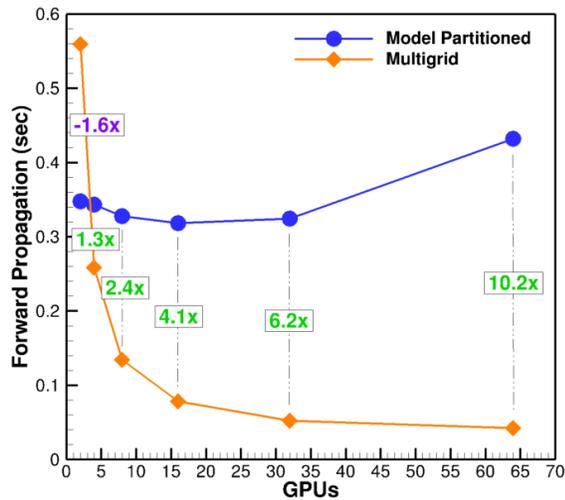


Fig. 7. 4,115-layer Residual Network composed of 2,071,328,150 parameters. Model Partitioned is serial execution distributed across multiple devices.

### E. Multi-GPU Parallel Training: 2 Billion Parameters

The 3.25 million parameter model highlighted the effects of the ratio of computation to communication with respect to the achievable performance via Multigrid Layer-Parallelism. To exploit this further, we derive a computation-intensive residual network composed of 4,115 layers: an opening convolution layer [ $7 \times 7$  kernel, 20 output channels, padding=1, single stride, ReLU], 16 repeated sequence blocks containing: one fully connected layer with matching input and output dimensions, and 256 residual convolution layers matching the opening convolution layer specifications. Lastly, the network is computed with a fully connected layer to a softmax function with cross-entropy loss. This model contains a total of **2,071,328,150** parameters, which is on the order of the Megatron-LM [26] model containing 8.3 billion parameters.

The 18 fully connected layers greatly increase the floating-point operation (FLOP) counts, therefore increasing the arith-

metic complexity defined as the ratio of FLOPs performed to the bytes transferred for the computation. Additionally, the number of model parameters requires some form of model parallelism as it is too large to reside on a single GPU. Thus, we will compare Multigrid to the traditional *Layer-Wise Parallelism*, which partitions the model across compute devices and performs serial network evaluation.

Figure 7 displays the parallel strong scaling performance comparing the Multigrid algorithm to the layer-wise parallelism method, denoted as Model Partitioned. While each method utilizes the same number of GPUs, it is clear that Multigrid outperforms its counterpart when using at least four GPUs. A speedup of 1.3x is achieved when four GPUs are utilized and a speedup of 10.2x is achieved when 64 GPUs are utilized. In the case of using four GPUs for MG, the computation to communication ratio is 92.8%, highlighting experimentally, the high arithmetic intensity. This ratio drops to 34.5% at 64 GPUs implying the performance becomes limited due to communication exhibiting low arithmetic intensity.

## V. CONCLUSION

This work exposed more parallelism within supervised Deep Learning, which can be paired with data-parallelism approaches. For model training and inference, Multigrid demonstrates avenues of increased throughput for compatible network architectures exhibiting high arithmetic intensity relative to the computing architecture utilized. We demonstrated performance gains up to 10.2x using NVIDIA V100 GPUs, a heterogeneous architecture tailored for high-compute and low-data movement. The algorithm is favorable towards architectures with improved communication interconnects between devices.

Lastly, the current state of hardware demonstrates limited kernel-execution concurrency, specifically in the context of convolution operators. Thus, the performance of this algorithm and implementation has not been fully realized. This algorithm will improve in future computer architectures containing massive thread-based hardware parallelism indicative of future exascale-era architectures [27]. Additionally, this algorithm can be implemented in conjunction with data-parallel techniques for multiplicative-compounding parallelism as required for next-generation supercomputing systems.

## ACKNOWLEDGMENT

The authors acknowledge The authors wish to acknowledge the following individuals for their contributions and support: Dave Martinez, Steve Rejto, Marc Zissman along with William Arcand, David Bestor, William Bergeron, Chansup Byun, Matthew Hubbell, Michael Houle, Anna Klein, Peter Michaleas, Lauren Miliechin, Julie Mullen, Andrew Prout, Antonio Rosa, and Charles Yee.

## REFERENCES

- [1] T. Ben-Nun and T. Hoefler, "Demystifying parallel and distributed deep learning: An in-depth concurrency analysis," *ACM Computing Surveys (CSUR)* 52, no. 4 (2019): 1-43.
- [2] A. Reuther, P. Michaleas, M. Jones, V. Gadepally, S. Samsi, and J. Kepner, "Survey and benchmarking of machine learning accelerators," *arXiv preprint arXiv:1908.11348* (2019).
- [3] C. E. Leiserson, N. C. Thompson, J. S. Emer, B. C. Kuszmaul, B. W. Lampson, D. Sanchez, and T. B. Schardl, "There's plenty of room at the Top: What will drive computer performance after Moore's law?," *Science* 368.6495, June 2020. DOI: 10.1126/science.aam9744.
- [4] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and Li Fei-Fei, "Imagenet large scale visual recognition challenge," *International journal of computer vision* 115, no. 3 (2015): 211-252.
- [5] D. Amodei, and D. Hernandez, "AI and Compute," <https://openai.com/blog/ai-and-compute/>, (2018).
- [6] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770-778. 2016.
- [7] V. Gadepally, J. Goodwin, J. Kepner, A. Reuther, H. Reynolds, S. Samsi, J. Su, and D. Martinez, "AI Enabling Technologies: A Survey," *arXiv preprint arXiv: 1905.03592* (2019).
- [8] S. Samsi, C. J. Mattioli, and M. S. Veillette, "Distributed Deep Learning for Precipitation Nowcasting," *arXiv preprint arXiv:1908.10964* (2019).
- [9] Y. Lu, A. Zhong, Q. Li, and B. Dong, "Beyond finite layer neural networks: Bridging deep architectures and numerical differential equations," *arXiv preprint arXiv: 1710.10121* (2017).
- [10] T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, "Neural ordinary differential equations," In *Advances in neural information processing systems*, pp. 6571-6583. 2018.
- [11] D. J. Mavriplis, and A. Jameson, "Multigrid solution of the Navier-Stokes equations on triangular meshes," *AIAA journal* 28, no. 8 (1990): 1415-1425.
- [12] D. J. Mavriplis, "Multigrid Techniques for Unstructured Mesh," *Numerical Methods for Fluid Dynamics V 5* (1995): 129.
- [13] L. Gaedke-Merzhäuser, A. Kopaničáková, and R. Krause, "Multilevel Minimization for Deep Residual Networks," *arXiv preprint arXiv: 2004.06196* (2020).
- [14] S. Günther, L. Ruthotto, J. B. Schroder, E. C. Cyr, and N. R. Gauger, "Layer-parallel training of deep residual neural networks," *SIAM Journal on Mathematics of Data Science* 2, no. 1 (2020): 1-23.
- [15] J. Nickolls, I. Buck, M. Garland, and K. Skadron, "Scalable parallel programming with CUDA," *Queue* 6, no. 2 (2008): 40-53.
- [16] S. Chetlur, C. Woolley, P. Vandermersch, J. Cohen, J. Tran, B. Catanzaro, and E. Shelhamer, "cudnn: Efficient primitives for deep learning," *arXiv preprint arXiv:1410.0759* (2014).
- [17] R. D. Falgout, S. Friedhoff, T. V. Kolev, S. P. MacLachlan, and J. B. Schroder, "Parallel time integration with multigrid," *SIAM J. Sci. Comput.*, 36(6):C635-C661, 2014. LLNL-JRNL645325.
- [18] V. A. Dobrev, T. Kolev, N. A. Petersson, and J. B. Schroder, "Two-level convergence theory for multigrid reduction in time (MGRIT)," *SIAM Journal on Scientific Computing* 39, no. 5 (2017): S501-S527.
- [19] R. D. Falgout, S. Friedhoff, T. V. Kolev, S. P. MacLachlan, and J. B. Schroder, "Parallel time integration with multigrid," *SIAM J. Sci. Comput.*, 36(6):C635-C661, 2014. LLNL-JRNL645325.
- [20] R. D. Falgout, S. Friedhoff, T. V. Kolev, S. P. MacLachlan, J. B. Schroder, and S. Vandewalle, "Multigrid methods with space-time concurrency," *Computing and Visualization in Science*, 18(4-5):123-143, 2017.
- [21] S. Günther, N. R. Gauger, and J. B. Schroder, "A non-intrusive parallel-in-time approach for simultaneous optimization with unsteady PDEs," *Optimization Methods and Software* 34, no. 6 (2019): 1306-1321.
- [22] W. Gropp, E. Lusk, A. Skjellum, *Using MPI: portable parallel programming with the message-passing interface*. Vol. 1. MIT press, 1999.
- [23] A. Reuther, J. Kepner, C. Byun, S. Samsi, W. Arcand, D. Bestor, B. Bergeron *et al.*, "Interactive supercomputing on 40,000 cores for machine learning and data analysis," In *2018 IEEE High Performance Extreme Computing Conference (HPEC)*, pp. 1-6. IEEE, 2018.
- [24] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE* 86, no. 11 (1998): 2278-2324.
- [25] T. Bosse, N. R. Gauger, A. Griewank, S. Günther, and V. Schulz, "One-shot approaches to design optimization," In *Trends in PDE Constrained Optimization*, pp. 43-66. Birkhäuser, Cham, 2014.
- [26] M. Shoeybi, M. Patwary, R. Puri, P. LeGresley, J. Casper, and B. Catanzaro, "Megatron-lm: Training multi-billion parameter language models using gpu model parallelism," *arXiv preprint arXiv:1909.08053* (2019).
- [27] P. Messina, "The exascale computing project," *Computing in Science & Engineering* 19, no. 3 (2017): 63-67.