Arithmetic-Intensity-Guided Fault Tolerance for Neural Network Inference on GPUs

Jack Kosaian
Carnegie Mellon University
jkosaian@cs.cmu.edu

K. V. Rashmi
Carnegie Mellon University
rvinayak@cs.cmu.edu

ABSTRACT

Neural networks (NNs) are increasingly employed in safety-critical domains and in environments prone to unreliability (e.g., soft errors), such as on spacecraft. Therefore, it is critical to impart fault tolerance to NN inference. Algorithm-based fault tolerance (ABFT) is emerging as an efficient approach for fault tolerance in NNs.

We propose an adaptive approach to ABFT for NN inference that exploits untapped opportunities in emerging deployment scenarios. GPUs have high compute-to-memory-bandwidth ratios, while NN layers have a wide range of arithmetic intensities. This leaves some layers compute bound and others memory-bandwidth bound, but current approaches to ABFT do not consider these differences. We first investigate ABFT schemes best suited for each of these scenarios. We then propose *intensity-guided ABFT*, an adaptive, arithmetic-intensity-guided approach that selects the most efficient ABFT scheme for each NN layer. Intensity-guided ABFT reduces execution-time overhead by 1.09–5.3× across many NNs compared to traditional approaches to ABFT.

CCS CONCEPTS

• **Computer systems organization** → **Redundancy; Reliability.**

KEYWORDS

fault tolerance, neural networks, arithmetic intensity

ACM Reference Format:

Jack Kosaian and K. V. Rashmi. 2021. Arithmetic-Intensity-Guided Fault Tolerance for Neural Network Inference on GPUs. In *The International Conference for High Performance Computing, Networking, Storage and Analysis (SC '21)*, November 14–19, 2021, St. Louis, MO, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3458817.3476184>

1 INTRODUCTION

Neural networks (NNs) are widely deployed for applications such as content recommendation [66], medical diagnosis [1], autonomous navigation [10], and space imaging [35]. These applications desire NNs to reliably make correct predictions: mispredictions in content recommendation can lead to revenue loss [57, 90], while those in safety-critical applications can result in loss of life [3].

One cause of unreliability in NNs is soft errors: transient errors that occur in processing logic and memory in computing systems

that can result in erroneous execution (e.g., $2 + 2 = 5$) [41]. The erroneous execution resulting from a soft error is referred to as a fault. There are many causes of soft errors, such as atmospheric radiation, voltage scaling, hardware wearout, and manufacturing error [37, 64, 85]. Recent works have shown the potentially-catastrophic effects of soft errors on NNs through fault injection [30, 31, 55, 61] and neutron beam experiments [38]: faults resulting from soft errors can cause mispredictions in NNs at a rate that violates automotive safety standards [3, 38, 55]. Furthermore, the rate at which soft errors occur increases with altitude [41, 64] and in space [25], posing a challenge to the trend of deploying NNs on low-cost hardware on spacecraft [35, 54, 81].

Therefore, applications that demand high reliability must employ some means of tolerating faults. However, tolerating faults caused by soft errors requires performing redundant execution (e.g., replication and comparison). For fault tolerance to be practical, it is critical that redundant execution operate with low overhead in terms of execution time and cost.

In this work, we focus on software-based approaches for detecting faults that occur in processing logic during NN inference on GPUs. We focus on detection, rather than correction, as detecting a catastrophic event is often more important to an application than quickly proceeding after such an event [43]. We focus on GPUs because they are commonly used for NN inference in both cluster and edge settings, including in emerging space applications [35, 81]. We focus on faults that occur in processing logic, rather than in the memory hierarchy, as many modern systems contain ECC-protected memory hierarchies [18]. In contrast, processing logic is not as amenable to lightweight hardware fault tolerance [21].

Algorithm-based fault tolerance (ABFT¹) is emerging as a promising approach toward imparting efficient software-based fault tolerance to NN inference [43, 57, 67, 89]. ABFT adds redundant computations employing carefully-designed mathematical structures, and exploits the invariants so introduced to detect faults. This approach enables ABFT to achieve significantly lower execution-time overhead than replication-based approaches. For this reason, ABFT is a common approach for fault tolerance in traditional HPC computations, such as matrix multiplication (e.g., [23, 24, 46, 85]), LU decomposition [83], sorting [58], and other iterative methods [28].

Figure 1 shows a toy example of ABFT-protected matrix multiplication between matrices A and B of size 2×2 to produce output matrix C . ABFT constructs a *column checksum vector* by summing each column of matrix A and a *row checksum vector* by summing each row of matrix B . It is straightforward to see that the result of taking the dot product of these checksum vectors should, in the absence of a fault, equal the summation of all entries of output

This work is licensed under a Creative Commons Attribution International 4.0 License.

SC '21, November 14–19, 2021, St. Louis, MO, USA
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8442-1/21/11.
<https://doi.org/10.1145/3458817.3476184>

¹While we focus on detecting errors, an approach sometimes termed “algorithm-based error detection” (ABED), we use the more common terminology ABFT for familiarity.

Figure 1: Toy example of ABFT with $M = N = K = 2$.

matrix C , which we refer to as the output summation. Correspondingly, comparison between the checksum dot-product result and the output summation can detect a single fault in C .

Multiple recent works have explored leveraging ABFT to impart fault tolerance to NNs [43, 57, 67, 89]. Since existing ABFT techniques support only linear computations, these approaches use ABFT for the linear operations of NNs (e.g., fully-connected and convolutional layers, which are often executed as matrix multiplications), and replicate nonlinear operations (e.g., activation functions). We similarly focus on using ABFT for linear layers implemented as matrix multiplications in this work, and use the terminology “linear layer” to refer to fully-connected and convolutional layers.

Key to efficient operation in any approach to redundant execution is identifying and exploiting underutilized resources. If the computation-to-be-protected underutilizes certain compute units, redundant execution can potentially be performed on those units without adding much execution-time overhead. However, existing approaches to ABFT typically only assume that computations being protected are compute bound, and thus aim to minimize the amount of redundant computation they perform.

In this work, we first present a case for challenging this assumption based on trends in GPU hardware and NN design: The introduction of processing units optimized for NNs (e.g., Tensor Cores [13]) has led to an unprecedented increase in FLOPs/sec in inference-optimized GPUs. However, such GPUs have had a far less profound growth in memory bandwidth. This results in inference-optimized GPUs having high compute-to-memory-bandwidth ratios (CMRs). High CMRs require kernels to have high arithmetic intensity to keep computational units highly utilized. However, many convolutional and fully-connected layers in NNs have *low arithmetic intensity*. Furthermore, many efforts toward reducing NN latency, such as efficient NN design [78], model specialization [45, 49, 50, 65, 75], and pruning [22], further reduce arithmetic intensity.

These trends result in many linear layers in NNs that have arithmetic intensity far lower than the CMR on GPUs, rendering such layers memory-bandwidth² bound, rather than compute bound. As a result, such layers are unable to keep computational units highly utilized, opening opportunities for redundant execution to be performed for free. However, current approaches to ABFT for NN inference, which are well-suited for compute-bound linear layers, cannot exploit this opportunity to squeeze in redundant execution alongside the computation being protected.

To better exploit this nascent opportunity, we (1) investigate ABFT schemes, which we refer to as thread-level ABFT, that exploit the unused computation cycles of the linear layer under protection on inference-optimized GPUs, and (2) propose a new, adaptive

approach to ABFT, called intensity-guided ABFT, that selects among thread-level ABFT and traditional approaches to ABFT on a per-layer basis, using the layer’s arithmetic intensity as a guide.

To design an approach to ABFT that can exploit the unused computation cycles of linear layers on modern inference-optimized GPUs, the key approach we leverage is to perform ABFT at the smallest unit of the parallel subproblem performed by the matrix multiplication for a layer. As illustrated in Figure 2, high-performance matrix multiplication on GPUs involves decomposing the overall matrix multiplication into a hierarchy of subproblems across threadblocks, warps, and, at the smallest level, threads. Existing approaches to ABFT for NN inference on GPUs, which we term “global ABFT,” generate checksums over the full input matrices to minimize the amount of redundant computation performed in checksum dot products. In contrast, we leverage an ABFT scheme in which each thread performs ABFT over the small matrix multiplication subproblem it is responsible for. We refer to this approach as *thread-level ABFT*. Under thread-level ABFT, each thread computes ABFT checksums and dot products on the fly in tandem with its computation of the original matrix multiplication, and performs its own thread-local checksum equality check.

The approach taken in thread-level ABFT may at first appear counterintuitive, as it performs more redundant computation than global ABFT: thread-level ABFT performs ABFT over many small, thread-local matrix multiplications, whereas global ABFT performs ABFT over one large matrix multiplication. In fact, thread-level ABFT results in multiple threads each computing identical checksums (e.g., in Figure 1, identical column checksums for threads that compute elements in the same rows in C). However, we show that, through careful design decisions, this approach is effective in exploiting the gaps in compute utilization of bandwidth-bound linear layers. This approach also eliminates any additional loads/stores, which would compete with the matrix multiplication itself for memory bandwidth, which is the bottleneck resource. The net result is low execution-time overhead for bandwidth-bound linear layers.

As described above, thread-level ABFT primarily benefits linear layers that are bandwidth bound. In contrast, it is not well-suited for compute-bound linear layers, for which global ABFT suffices. As we show in §3, NNs contain *both bandwidth- and compute-bound linear layers*, making one-size-fits-all approaches inefficient.

Therefore, we propose *intensity-guided ABFT*, an adaptive ABFT approach that selects among global ABFT and thread-level ABFT for each linear layer of a NN depending on which approach offers the lowest execution-time overhead, letting the arithmetic intensity of the layer and CMR of the device guide such selection.

We implement and evaluate intensity-guided ABFT atop CUTLASS [8], a high-performance library from NVIDIA for matrix multiplications on GPUs. We evaluate execution-time overhead on the inference-optimized NVIDIA T4 GPU when using Tensor Cores. We consider eight popular convolutional NNs (CNNs), two NNs used within recommendation models (DLRM) [66], and four CNNs developed through model specialization and used for video analytics [50]. Compared to an optimized global ABFT approach [43], intensity-guided ABFT reduces execution-time overhead by up to 2.75× for popular CNNs, up to 4.55× for DLRMs, and up to 5.3× for specialized CNNs. These results show the promise of taking

²We refer to memory-bandwidth-bound layers as “bandwidth-bound” for short.

an arithmetic-intensity-guided approach to ABFT to impart low-overhead fault tolerance to NN inference.

The code used in this paper is available at <https://github.com/Thesys-lab/arithmetic-intensity-guided-abft>.

2 BACKGROUND

In this section, we provide background on matrix multiplication on GPUs, the need for fault tolerance in NN inference, and how ABFT is performed and optimized for NN inference.

2.1 Efficient matrix multiplication on GPUs

As described in §1, our focus is on redundant execution for the convolutional and fully-connected layers of NNs, which we refer to as “linear layers.” For the remainder of this paper, we describe these operations as matrix multiplications, as high-performance implementations of these layers are often achieved through matrix multiplications [8]. However, the approaches we propose can apply to other implementations as well.

Within this setting, we denote a linear layer as the multiplication of matrix A of size $M \times K$ by matrix B of size $K \times N$ to produce an output matrix C of size $M \times N$. Matrix A contains the inputs to the layer (e.g., activations from the previous layer). Matrix B contains the learned weights of this layer. Weights (matrix B) are known a priori, while activations (matrix A) are known only during computation. Output C contains the output of the layer, which will be fed to the next layer, typically after being operated on by an activation function (e.g., ReLU).

GPU terminology. We use NVIDIA’s terminology [2] in describing the architectural components and programming abstractions of GPUs. A GPU consists of a number of streaming multiprocessors (SMs), each of which has many cores on which computation is performed along with a register file and shared memory region. Computation is executed on GPUs in kernels consisting of many threads. Threads are grouped into threadblocks, with all threads in a threadblock executing on the same SM and able to communicate with one another via shared memory. Groups of 32 threads within a threadblock execute in lockstep as a so-called warp. Each thread executes on an individual core, except when using Tensor Cores, new processing units that enable warp-wide collaborative execution of matrix multiplications [13].

Hierarchical matrix multiplication. High-performance implementations of matrix multiplication on GPUs decompose the problem solved by the kernel into a number of sub-matrix multiplications solved by threadblocks, warps, and threads. Figure 2 shows an example of this decomposition: each threadblock is responsible for computing a subset of C , which it decomposes into subsets to be computed by warps of the threadblock, each of which in turn decomposes the problem into subsets to be computed by individual threads. We denote the portions of A and B used by a thread as A_t and B_t , respectively. A_t is of size $M_t \times K$ and B_t is of size $K \times N_t$.

As our focus is on NN inference, we focus on low-precision (e.g., FP16) matrix multiplications on Tensor Cores, which are heavily used for accelerating inference. Our description follows the use of such operations in CUTLASS. We focus in particular on the FP16 m16n8k8 Tensor Core operation, though our discussion and proposed solutions apply to other Tensor Core operations as well.

Each m16n8k8 Tensor Core operation is a warp-wide operation that multiplies a 16×8 matrix A_{tc} by an 8×8 matrix B_{tc} and accumulates results into a 16×8 output matrix C_{tc} (we use subscript “tc” to denote Tensor Core operands/outputs) [12]. Each thread in the warp provides four elements of A_{tc} and two elements of B_{tc} to the operation, and obtains four output elements of C_{tc} from the operation. We refer to one such m16n8k8 matrix-multiply-accumulation operation as an “MMA,” following NVIDIA’s terminology [12].

CUTLASS leverages MMAs within the hierarchical matrix multiplication framework described above. Each thread walks down the K dimension of the problem and loads an $M_t \times 2$ chunk of A_t and a $2 \times N_t$ chunk of B_t . These loaded chunks are then used in $\frac{M_t N_t}{2}$ MMAs, each of which uses two rows of the loaded chunk of A_t and one column from the loaded chunk of B_t from each thread, as shown in Figure 3. The results of these operations are accumulated into the thread’s $M_t N_t$ registers that store the partial accumulation of the thread’s matrix multiplication output. CUTLASS uses standard optimizations to overlap loading the next chunks of A_t and B_t while the current MMAs are performed (e.g., double buffering).

2.2 Need for fault tolerance

As described in §1, our focus in this work is on detecting faults resulting from transient soft errors on GPUs. Handling soft errors has long been a concern for HPC systems [37, 41, 44] due to their large scale and the criticality of the workloads that utilize them. Beyond these settings, the increasing trend of leveraging NNs in cyber-physical systems, such as autonomous vehicles [72], and in harsh operating environments, such as in spacecraft [20, 35, 81], has bolstered the need for fault tolerance solutions for NNs. For example, autonomous vehicles leverage GPUs and must tolerate faults to meet strict reliability requirements [3, 72]. Furthermore, servers equipped with general-purpose GPUs have recently been sent to the International Space Station to perform scientific computations [20, 81], and which leverage software-based fault tolerance to handle the harsh operating environment therein.

Combining the longstanding need for reliability in HPC applications with the growing need for reliability in safety-critical and edge deployments of GPUs, efficient approaches to fault tolerance are necessary both in the present and for the future.

2.3 Fault model

We next shift our focus to fault detection for matrix multiplication. To set the stage, we first describe the fault model we consider.

We focus on detecting a single faulty output value in matrix C . We focus on detection, rather than correction, as being able to detect a catastrophic event is often more important than being able to quickly continue after such an event [43]. Following prior work [29, 67, 89], we focus on detecting a single fault because the execution of one layer in a NN is short enough that the likelihood of more than one soft error occurring during execution is low.

We focus on faults occurring due to soft errors in the processing logic of a GPU. We do not focus on faults in the memory hierarchy, such as in global memory, caches, shared memory, register files, and busses, as these components are more easily protected by ECC [18]. In contrast hardware fault tolerance for processing units is more expensive, typically requiring dual-modular-redundant circuits [21].

Figure 2: Hierarchical matrix multiplication. Shaded regions show inputs/outputs used in the next level of the hierarchy.

Figure 3: One step along the K dimension for a thread using **m16n8k8** Tensor Core operations (MMAs). A total of $\frac{M_t N_t}{2}$ MMAs are performed per step, one for each combination of two consecutive rows of A_t and one column of B_t .

We also assume that control logic on the GPU is protected. This fault model is in line with prior work [29, 56].

2.4 ABFT for matrix multiplication

ABFT falls under a class of techniques called “redundant execution” in which additional computation is performed on top of the computation-to-be-protected for the purpose of fault tolerance. As described in §1, ABFT adds redundant computations employing carefully-designed mathematical structures, and exploits the invariants so introduced to detect errors while performing less redundant computation than replication-based approaches [46].

ABFT for matrix multiplication typically operates by (1) generating a $1 \times K$ column checksum vector of matrix A and a $K \times 1$ row checksum vector of matrix B , (2) performing the dot product between the column checksum vector and row checksum vector, (3) summing all entries of the output matrix C , and (4) comparing the values generated in (2) and (3) above. Approaches to ABFT typically generate a single column checksum for the entire input matrix A (and similarly for B) [43, 89]. We thus term such approaches “global ABFT.” Global ABFT results in the minimum additional dot-product computations required for fault detection in matrix multiplication, making it well-suited for compute-bound matrix multiplications.

While we focus on detecting a single fault, ABFT also supports detecting multiple faults. To do so, ABFT generates multiple checksum columns and rows based on independent linear combinations of columns/rows. In this scenario, multiple output checksums are also generated based on these linear combinations and compared to checksum dot products. The approaches to ABFT that we propose in this work can also handle higher fault rates in this way.

2.5 Optimizing global ABFT for NN inference

Recent works leverage global ABFT to protect the linear layers of NNs, and add multiple NN-specific optimizations [43, 57, 89], which we describe next. Recall that, for NN inference, matrix A contains input activations and B contains layer weights. We therefore refer to the column checksum of matrix A as the “activation checksum” and the row checksum of matrix B as the “weight checksum.”

Offline construction of weight checksum. Since operand B contains the layer’s weights, which remain the same for every inference request, the weight checksum of each linear layer in a NN can be constructed once offline and reused for every inference request [43, 57, 89]. The same does not hold for the activation checksum of operand A , because its contents change for each inference request.

Checksum fusion. Recent work [43] proposes to fuse the generation of the output summation used in the ABFT check to the end of the matrix multiplication kernel itself. Kernel fusion significantly reduces the amount of data that must be read from memory to form the output summation, which speeds up checksum generation. As the next layer’s input A is generated by the current layer, the current layer can also fuse the generation of the next layer’s activation checksum to the end of its matrix multiplication kernel (after the activation function has been applied) [43].

Flow of ABFT in NN inference. With the above optimizations, the workflow of an ABFT-protected linear layer is as follows: (1) perform matrix multiplication to generate output C , (2) perform fused output summation generation, (3) apply the layer’s activation function to C , (4) perform fused next-layer activation checksum generation, (5) launch a kernel that performs the ABFT dot product for the current layer and compares the results to the output checksum generated in Step 3. Steps 1–4 must take place sequentially, while Step 5 can take place in parallel with the next layer of the NN. Step 5 occurs in a separate kernel because it involves a global reduction over the partial checksums generated by threadblocks.

By minimizing redundant computation, global ABFT offers low execution-time overhead for compute-bound linear layers. However, we next identify trends in GPU hardware and NNs that lead to many linear layers being memory-bandwidth-bound. This opens new opportunities for efficient redundant execution that current approaches to ABFT for NN inference are unable to exploit.

3 NEW OPPORTUNITIES FOR EFFICIENT REDUNDANT EXECUTION

Critical to reducing execution-time overhead for any approach to redundant execution is discovering opportunities to exploit unused resources. In this section, we identify trends in GPU hardware and

NN design that create new, currently unexploited opportunities for efficient redundant execution in NN inference.

3.1 Resource bottlenecks for GPU kernels

GPU kernels are typically either bound by computational throughput or by memory bandwidth. A popular model to determine whether a kernel is compute or memory-bandwidth bound is comparing the the *arithmetic intensity* of the kernel to the *compute-to-memory-bandwidth ratio* (CMR) of the device [9, 82]. Under this model, a kernel is compute bound if the theoretical amount of time it spends performing computation is greater than the theoretical amount of time it spends loading/storing data from/to memory:

$$\frac{\text{FLOPs}}{\text{Compute Bandwidth}} > \frac{\text{Bytes}}{\text{Memory Bandwidth}}$$

Here, “FLOPs” is the number of arithmetic operations performed by the kernel, “Bytes” is the amount of data it transfers to/from memory, “Compute Bandwidth” is the GPU’s peak FLOPs/sec, and “Memory Bandwidth” is the GPU’s memory bandwidth (bytes/sec). Rearranging this inequality to pair properties of the kernel on the left-hand side and properties of the GPU on the right-hand gives:

$$\frac{\text{FLOPs}}{\text{Bytes}} > \frac{\text{Compute Bandwidth}}{\text{Memory Bandwidth}} \quad (1)$$

The left-hand ratio of Equation 1 is the kernel’s arithmetic intensity: the ratio between the FLOPs the kernel performs and the bytes it transfers to/from memory. The right-hand ratio is the GPU’s CMR.

Takeaway. From the lens of this performance model, it is clear that the arithmetic intensity of a given kernel and CMR of a given GPU play key roles in determining opportunities for redundant execution to leverage unused resources. For example, a kernel with low arithmetic intensity running on a GPU with a high CMR will likely be bandwidth bound and underutilize compute units. This leaves opportunities for redundant execution to leverage such units without hampering the performance of the kernel itself.

We next examine trends in GPU hardware and NN design to identify opportunities for such efficient redundant execution.

3.2 Wide range of arithmetic intensities exhibited by NN layers

We first examine the arithmetic intensities of current NNs and their individual linear layers under various operational settings. In this analysis, we consider only “linear layers”, such as convolutional and fully-connected layers, which are often implemented as matrix multiplications. Other operations, such as activation functions, are typically fused to these linear layers and contribute far less to overall arithmetic intensity and execution time.

The “aggregate arithmetic intensity” of a NN as a whole is computed by summing the FLOPs performed across all linear layers, summing the bytes read/written across all linear layers, and dividing these quantities. This metric provides an estimate of whether the NN as a whole is more compute or memory-bandwidth bound.

Figure 4 shows the FP16 aggregate arithmetic intensities of eight widely-used CNNs from the popular PyTorch Torchvision

Figure 4: FP16 aggregate arithmetic intensity of CNNs operating on images of size 1080×1920 at batch size of one.

library [16].³ The figure shows a wide range of aggregate arithmetic intensities among such CNNs (from 71 to 220).

Furthermore, many domains leverage NNs that are significantly smaller than those described above, and thus have even lower aggregate arithmetic intensities. For example, NNs used for recommendation serving, such as Facebook’s popular DLRM [66], leverage small NNs consisting of a few fully-connected layers. Consequently, these NNs have low aggregate arithmetic intensities (e.g., 7 in FP16).

Figure 5 shows the arithmetic intensities of individual convolutional and fully-connected layers of ResNet-50. As illustrated, there is a wide variance of arithmetic intensities (1–511) among even various linear layers of the same NN (other NNs are similar).

Finally, arithmetic intensity also varies with settings of the applications in which NNs operate, such as the size of inputs to the NN. For example, increasing the batch size used in inference typically increases arithmetic intensity by amortizing the overhead of loading NN weights from memory. Thus, the many applications that use small batch sizes for low-latency inference are likely to have low arithmetic intensity [32, 88], while those that can aggressively batch inputs may have higher arithmetic intensity. For example, the FP16 aggregate arithmetic intensities of the NNs used in DLRM increase from 7 at batch size of 1 to 70–109 at batch size 256. For CNNs, the resolution of input images also affects arithmetic intensity for similar reasons, as operating over large images amortizes the cost of loading convolutional filters from memory. For example, the FP16 aggregate arithmetic intensity of ResNet-50 is 72 when operating over images of resolution 224×224 (the resolution typically used for ImageNet [71]), but increases to 122 when operating over images of resolution 1080×1920 (typically considered HD).

Takeaway. NNs exhibit wide variance in arithmetic intensity across NNs, across individual linear layers within a NN, and across application settings. This renders some NNs, some layers, and some application settings likely to underutilize computational resources.

3.3 Inference-optimized GPUs have high CMR

We now discuss trends in CMR, the right-hand ratio in Equation 1.

General-purpose GPUs have been a workhorse for NNs since the early 2010s [53]. Recent GPUs have further bolstered NN acceleration by adding hardware units specifically designed for the matrix multiplications found in NNs, such as NVIDIA’s Tensor Cores [13]. These hardware units offer unprecedented performance in terms of

³We replace the group convolutions in ShuffleNet and ResNext-50 with non-grouped convolutions to ease their conversion to matrix multiplications. The reported aggregate arithmetic intensities of these NNs are, thus, higher than they would be with grouped convolutions, which typically decrease arithmetic intensity.

Figure 5: FP16 arithmetic intensity of convolutional and fully-connected layers of ResNet-50 on HD images (resolution 1080×1920) with batch size of one.

FLOPs/sec, particularly when using low-precision arithmetic (e.g., FP16), as is common in NN inference.

For example, the inference-optimized T4 GPU offers 65 FP16 TFLOPs/sec [19], a marked increase from the 11 FP16 TFLOPs/sec offered by its predecessor, the P4 [15], which did not contain Tensor Cores. Such high performance is also offered in other server-grade GPUs, such as the V100 and A100 GPUs, which offer 125 and 312 FP16 TFLOPs/sec, respectively [4, 18]. This trend has also made its way to edge devices, as GPUs in the NVIDIA Jetson family now offer up to 32 INT8 TOPs/sec via Tensor Cores, whereas predecessors were bound to single-digit INT8 TOPs/sec [11].

While the FLOPs/sec offered by inference-optimized GPUs has drastically increased, memory bandwidth has not increased at the same rate. For example, while the T4 GPU increases FP16 FLOPs/sec by $5.9\times$ compared to the P4 GPU, it offers only a $1.7\times$ increase in memory bandwidth. Similar trends hold for other GPUs.

The net result of these trends in compute and memory bandwidth is a *significant increase in the CMR of GPUs*. For example, the FP16 CMR of the T4 GPU is 203, while that of the P4 was 58. Even GPUs with high-bandwidth memory (e.g., HBM2) have high CMRs (139 and 201 in FP16 for V100 and A100, respectively [4, 18]), as do edge GPUs (235 in INT8 for Jetson AGX Xavier [11]).

Takeaway. The introduction of specialized hardware units for matrix multiplications that drastically increase computational throughput, but a significantly slower increase in memory bandwidth, results in inference-optimized GPUs with high CMRs. This trend “raises the bar” for GPU kernels, *making them more likely to be memory-bandwidth bound and underutilize GPU compute units*.

3.4 Many NN optimization trends reduce arithmetic intensity

A secondary trend further exacerbates the growing bandwidth-bound nature of many NNs: designing small NNs to perform tasks with high throughput or low latency. The NNs shown in Figure 4 are large, general-purpose NNs designed to classify a wide variety of objects (e.g., from ImageNet [71]). There is a growing body of work on designing more efficient NN architectures that can accomplish the same task as such general-purpose NNs, but with a significantly smaller NN. There are many techniques along these lines, including efficient neural architecture search [78, 91], pruning [22], and model specialization [45, 49, 50, 65, 75]. These techniques often result in deploying NNs with lower aggregate arithmetic intensity than the general-purpose NNs shown in Figure 4.

For example, in model specialization for offline video analytics, a small, specialized CNN is designed to answer specific queries (e.g., find red trucks), and which consults a larger, general-purpose CNN only when unsure [45, 50, 75]. By targeting a focused query,

specialized CNNs can typically be made smaller and faster than general-purpose NNs, but exhibit far lower aggregate arithmetic intensity: the specialized CNNs from the widely-cited NoScope video analytics system [50] have FP16 aggregate arithmetic intensities of 15–53, even with large batch size.

Takeaway. Current trends in efficient NN design result in NNs that have lower arithmetic intensity, making current and future workloads likely to underutilize GPU compute units.

3.5 Takeaways and new opportunities

The previous sections have identified trends that lead to the conclusion that *the current and future landscape of NN inference will contain a significant number of memory-bandwidth bound linear layers*: §3.2 illustrated that current NNs, the linear layers within them, and their application settings exhibit a wide variance of arithmetic intensities (including many with low arithmetic intensity), and §3.4 described increasingly prevalent trends in NN design that often reduce arithmetic intensity. Coupling this with the dramatic growth in CMR for GPUs described in §3.3 drives home the conclusion that current and future NNs will contain bandwidth-bound linear layers that underutilize the computational capabilities of GPUs.

Such bandwidth-bound linear layers leave room open for redundant execution to fill gaps in compute utilization during matrix multiplication. However, current approaches to ABFT for NN inference are unable to exploit these fine-grained opportunities for efficient redundant execution. As described in §2.5, global ABFT operates at a much higher level (specifically, kernel level), and hence is unable to exploit compute underutilization that occurs at *finer granularity within the matrix multiplication operation*.

This calls for investigating approaches to redundant execution that can exploit the fine-grained compute underutilization exhibited by current and future matrix multiplication kernels in NN inference. Such an approach would complement global ABFT, which is well-suited for the compute-bound linear layers in NNs.

We next turn our focus toward investigating such an approach.

Key design principle. Driven by the opportunities outlined above, we use the following principle when considering approaches to redundant execution for memory-bandwidth-bound matrix multiplications: *avoid performing additional memory accesses whenever possible even if doing so comes at the expense of additional computation*. Adhering to this principle avoids competing with the matrix multiplication for its bottleneck resource, memory bandwidth.

4 THREAD-LEVEL REPLICATION?

A natural question that arises when considering options for redundant execution for bandwidth-bound linear layers is whether it is beneficial to use *thread-level replication*, rather than ABFT. After all, ABFT is primarily designed to reduce the number of redundant operations performed compared to replication, while spare compute cycles are plentiful in bandwidth-bound linear layers. Furthermore, thread-level replication easily satisfies the design principle stated in §3.5, by sharing loads with the original matrix multiplication.

We began our exploration of redundant execution for bandwidth-bound linear layers with replication for these very reasons, but ultimately found it to have higher execution-time overhead than ABFT, as we next describe. We focus on matrix multiplications

using m16n8k8 FP16 Tensor Core operations (MMAs) (described in §2.1). Recall that, in matrix multiplication using this operation, each thread participates in $\frac{M_t N_t}{2}$ MMAs on each iteration along the K dimension. For each MMA, a thread provides four elements from A_t , two elements from B_t , and receives four output elements.

We have considered two approaches to thread-level replication:

Traditional replication. The traditional approach to performing thread-level replication is to perform $\frac{M_t N_t}{2}$ additional MMAs per step down the K dimension, accumulate the results in a separate set of $M_t N_t$ registers, and compare these registers to the original $M_t N_t$ matrix multiplication output registers. However, we found that the $2\times$ increase in output register usage per thread in this approach limits the number of threadblocks that can be co-scheduled on a single SM (so-called “occupancy” [17]), and leads to significant slowdowns compared to the original matrix multiplication kernel.

Replicated MMA, single accumulation. Based on this limitation, we next explored replicating MMAs, but accumulating results to a single set of four output registers. Under this approach, one still performs $\frac{M_t N_t}{2}$ additional MMAs per step along the K dimension, but each redundant MMA accumulates results to the same set of four registers. By the end of the thread-level matrix multiplication, in the absence of a fault, the summation of these four registers will equal the summation of the thread’s “original” $M_t N_t$ output registers. Threads can use this invariant to detect faults.

We find that the limited additional register usage of this approach alleviates the occupancy-related slowdowns described above, and thus significantly reduces execution-time overhead compared to the traditional form of replication. However, as we will show in §6.5, doubling the number of MMAs performed results in higher execution-time overhead than ABFT.

We thus turn our focus to investigating ABFT schemes that can exploit the compute underutilization identified in §3.

5 ARITHMETIC-INTENSITY-GUIDED ABFT

In this section, we first investigate approaches to ABFT that can exploit the fine-grained compute underutilization of bandwidth-bound linear layers identified in §3. We then describe the design of an adaptive approach to ABFT that selects an ABFT scheme for each linear layer guided by the layer’s arithmetic intensity.

5.1 At which level should ABFT be performed?

The hierarchical decomposition of matrix multiplications described in §2.1 offers multiple levels at which ABFT can be performed: the kernel level (as in global ABFT), threadblock level, warp level, or thread level. However, performing ABFT at any level other than the thread level requires performing additional loads/stores to generate checksums. For example, performing ABFT at the level of a threadblock requires individual threads to cooperate to generate threadblock-wide checksums, which requires storing and loading thread-local partial checksums. Such additional loads and stores violate the design principle described in §3.5 and compete for bandwidth with the matrix multiplication itself.

In contrast, performing ABFT at the level of individual threads avoids additional loads/stores. Figure 6 compares one approach to thread-level ABFT with global ABFT at a high level. Concretely,

Figure 6: Global and two-sided thread-level ABFT.

thread-level ABFT involves threads in the matrix multiplication kernel performing their own, local ABFT calculations across their own, local sub-matrix multiplications. Thread-level ABFT eliminates additional loads/stores by (1) sharing the loads of operands that will be used for checksum generation with those that were already performed for thread-level matrix multiplication in a step along the K dimension, and (2) eliminating stores of partial checksums for use in threadblock- or warp-wide checksum generation.

Thus, we conclude that performing ABFT at the thread level is the appropriate fit for ABFT optimized for bandwidth-bound linear layers. This conclusion is heavily driven by the design principle established in §3.5 of avoiding additional loads/stores. In cases where this principle can be relaxed, performing ABFT at other levels of the matrix multiplication hierarchy may be appropriate. Even with the somewhat-extreme stance taken in thread-level ABFT, we will show in §6 that thread-level ABFT significantly reduces execution-time overhead for bandwidth-bound linear layers.

5.2 Design decisions for thread-level ABFT

Even having narrowed our focus to performing ABFT at thread level for bandwidth-bound linear layers, there remain multiple design decisions that affect performance, which we discuss next. Similar to §4, we focus on m16n8k8 Tensor Core operations (MMAs), which are described in detail in §2.1.

5.2.1 Online computation of weight checksums. Recall from §2.5 that optimized approaches to global ABFT for NNs typically compute the weight checksum of B once offline and load it upon every inference request. We do not employ this technique for thread-level ABFT, as doing so would require threads to load weight checksums from memory, violating the design principle described in §3.5. Thus, thread-level ABFT recomputes thread-local weight checksums alongside the thread-level matrix multiplication.

5.2.2 Balancing checksum generation and redundant MMAs. Adopting the ABFT approach described in §2.4 at thread level would involve performing the following for each step the thread takes along the K dimension: (1) computing a thread-level activation checksum from A_t , (2) computing a thread-level weight checksum from B_t , and (3) performing a single MMA over these checksums to generate ABFT output values. These steps are illustrated in the left-hand side of Figure 7, and are repeated for each iteration along the K dimension, accumulating into the same ABFT output registers. Once the thread has completed all iterations along the K dimension,

Figure 7: Comparison of two-sided and one-sided thread-level ABFT for a single step in the K dimension.

Table 1: Additional Tensor Core MMAs and checksum operations done by thread-level replication (Rep.), two-sided ABFT, and one-sided ABFT per step in the K dimension.

	Rep.	Two-sided	One-sided
Tensor Core MMAs	$M_t N_t / 2$	1	$M_t / 2$
Checksum ops.	0	$O(M_t + N_t)$	$O(N_t)$

it generates a thread-local output summation and compares it to the final ABFT output registers. We call this approach *two-sided thread-level ABFT*, as it generates checksums for both A_t and B_t .

Two-sided thread-level ABFT minimizes the number of redundant MMA operations performed by thread-level ABFT, as it performs only one extra MMA for every step along the K dimension. However, it maximizes the amount of computation performed in generating thread-local activation and weight checksums.

It is important to note that checksum generation involves summations that will execute on traditional arithmetic units on the GPU (e.g., using HADD2 PTX instructions), rather than on Tensor Cores. In contrast, redundant MMA operations will execute on Tensor Cores. Thus two-sided thread-level ABFT will more significantly utilize traditional arithmetic units than Tensor Cores because it performs $O(M_t + N_t)$ additional checksum generation operations but only one additional MMA per step along the K dimension.

Given that Tensor Cores are the drivers behind the math performed in the matrix multiplication for a linear layer, it is Tensor Cores that are heavily underutilized by bandwidth-bound linear layers, rather than traditional arithmetic units. Traditional arithmetic units are likely not as underutilized in bandwidth-bound linear layers, as they are also used by threads to carry out general control flow (e.g., updating loop counters) and to assist in loading/storing data (e.g., computing addresses). Thus, minimizing the number of additional MMAs performed in two-sided thread-level ABFT may insufficiently exploit underutilized Tensor Cores. At the same time, our experience with replication in §4 indicates that adding too many additional MMAs can also lead to high overhead.

To straddle this tradeoff between added operations to Tensor Cores and added operations to traditional arithmetic units, we leverage a *one-sided thread-level ABFT scheme*. Rather than computing checksums for both A_t and B_t and performing a single MMA across these checksums, one-sided thread-level ABFT instead generates a

checksum only for B_t and multiplies the entirety of A_t with this checksum.⁴ As illustrated in the right-hand side of Figure 7 this results in performing $\frac{M_t}{2}$ additional MMAs, and $O(N_t)$ checksum generation operations for each step along the K dimension.

As shown in Table 1, one-sided thread-level ABFT sits in the “sweet spot” between thread-level replication and two-sided thread-level ABFT in terms of additional MMAs and checksum operations performed. We illustrate in §6.5 that this enables one-sided thread-level ABFT to provide the lowest execution-time overhead among these approaches to thread-level redundant execution.

5.3 Per-layer, intensity-guided adaptation

As shown in §3, NNs have a mix of compute- and bandwidth-bound linear layers. The ABFT scheme with the lowest execution-time overhead for a given layer depends on the bottleneck of the layer, with global ABFT preferable for compute-bound layers and thread-level ABFT preferable for bandwidth-bound layers.

Rather than selecting one ABFT scheme to be applied to all linear layers of a NN, we propose *intensity-guided ABFT*, which selects among global ABFT and thread-level ABFT for each individual linear layer. Prior to deploying a NN, intensity-guided ABFT measures the execution-time overhead of each linear layer under global ABFT and thread-level ABFT, and chooses the scheme with the lowest overhead for that layer. As we show in §6, in conforming to the ideas presented in this paper, linear layers with higher arithmetic intensity typically benefit from global ABFT, while those with lower arithmetic intensity typically benefit from thread-level ABFT. Thus, intensity-guided ABFT uses arithmetic intensity as a guide in selecting the best ABFT scheme for each layer. Our evaluation in §6 shows that intensity-guided ABFT significantly reduces overhead compared to either global or thread-level ABFT alone.

Integration with pre-deployment optimizers. Intensity-guided ABFT fits alongside the popular approach of pre-deployment optimization in NN inference, as performed by frameworks like TensorRT [14], TVM [27], cuDNN [6], and CUTLASS. This process takes in a NN and an input size (e.g., image resolution, batch size) that will be used during inference and enumerates and executes all configurations of each layer in the NN (e.g., tile sizes, matrix layouts). The configuration with the lowest execution time for a layer is chosen for that layer and used for all inference requests during deployment. A pre-deployment optimizer using intensity-guided ABFT will include global ABFT and thread-level ABFT in its enumeration of configurations of a matrix multiplication. Intensity-guided ABFT chooses the fastest among these, which typically aligns with the arithmetic intensity of the layer, as we show in §6.

6 IMPLEMENTATION AND EVALUATION

We now evaluate the execution-time overhead of intensity-guided ABFT. The highlights of the evaluation are as follows:

- Across eight popular CNNs, two NNs used in DLRMs, and four specialized CNNs, intensity-guided ABFT reduces execution-time overhead compared to global ABFT by 1.09–5.3×.

⁴One can alternatively multiply a checksum of A_t with B_t . We have selected the converse due to ease of implementation in CUTLASS.

Figure 8: Execution-time overhead on all NNs considered. To avoid clutter, error bars are not plotted in this figure, but are plotted in Figures 9, 10, and 11 for each NN.

- Intensity-guided ABFT provides the largest reductions in execution-time overhead for NNs that have many linear layers with low arithmetic intensity, such as DLRLMs (up to 4.9× reduction) and specialized CNNs (up to 5.3× reduction).
- Even for NNs that have many linear layers with high arithmetic intensity, intensity-guided ABFT still significantly reduces execution-time overhead (e.g., 1.5× for Wide-ResNet-50). This shows the benefit of intensity-guided ABFT’s adaptive approach to ABFT, as even NNs that are primarily compute bound often have some linear layers with low arithmetic intensity.
- Intensity-guided ABFT provides similar benefits across various input resolutions (§6.4.1) and batch sizes (§6.4.2).
- The one-sided thread-level ABFT approach motivated in §5.2.2 significantly reduces execution-time overhead compared to two-sided thread-level ABFT and thread-level replication (§6.5).

6.1 Implementation

Recall that intensity-guided ABFT adapts to each linear layer in a NN by choosing between thread-level ABFT and global ABFT. We implement thread-level ABFT and global ABFT in CUDA/C++ atop CUTLASS [8], a high-performance, open-source matrix multiplication library developed by NVIDIA. For thread-level ABFT, we modify existing thread-level inner loops in CUTLASS to perform checksum generation, redundant MMAs, and final checksum comparison. We implement the global ABFT scheme based on the state-of-the-art approach from Hari et al. [43] (discussed in §2.5), using NVIDIA’s CUB library [5] when possible.

Recall from §5.3 that intensity-guided ABFT fits alongside common pre-deployment NN optimizers. We integrate intensity-guided ABFT into the pre-deployment workflow of the CUTLASS profiler, which selects the fastest matrix multiplication kernel and configuration (e.g., tile size, layout) for a given matrix multiplication size.

6.2 Evaluation setup

Baselines. Our main comparison is between intensity-guided ABFT and the state-of-the-art approach to global ABFT for NN inference on GPUs described in §2.5. We also evaluate one-sided thread-level ABFT alone (referred to as “thread-level ABFT”), and in §6.5 compare to two-sided thread-level ABFT and thread-level replication.

Metrics. Execution-time overhead is one of the primary metrics of interest for redundant execution. For each linear layer of a NN, we obtain the execution time of the original matrix multiplication without redundancy (T_o), as well as that of the redundant version (T_r) and report the percentage increase in execution time ($\frac{T_r - T_o}{T_o} * 100$). We report execution-time overhead for an entire NN

by summing the per-layer execution times and using these in the equation above. We include only linear layers, as these layers typically dominate the end-to-end execution time of a NN. Moreover, aggregating the execution times of each linear layer in this fashion is representative of overall execution-time overhead for the NN as a whole, because, for all of the NNs we consider, the subsequent layer of the NN cannot begin executing until the current layer has completed execution. We report the mean of 10 trials of 1000 runs after 100 warmup runs. Error bars show the maximum and minimum time overheads across trials. In many cases, error bars are imperceptible due to their tightness. We do not plot error bars in Figure 8 to avoid clutter and because all error bars for NNs in Figure 8 are plotted in Figures 9, 10, and 11. Note that while in some cases error bars may give the incorrect impression that intensity-guided ABFT performs worse than global ABFT, intensity-guided ABFT, by design, always performs at least as well as global ABFT.

We also report “aggregate arithmetic intensity” (defined in §3.2). In each figure, the FP16 aggregate arithmetic intensity is listed in parentheses below each model.

Evaluation setting. We evaluate on an NVIDIA T4 GPU [19], which is a state-of-the-art inference-optimized GPU, on an AWS g4dn.xlarge instance. The T4 offers 65 FP16 TFLOPs/sec and 320 GB/sec of memory bandwidth, giving it an FP16 CMR of 203. We use CUDA 11.0 and configure the clockrate of the GPU according to that used in CUTLASS [7]. We perform all experiments using FP16 datatypes and we use the m16n8k8 matrix multiplications targeting Tensor Cores described in §2.1. Note that it is standard to perform NN inference in low precision, such as FP16. We pad matrix dimensions M , N , and K to be multiples of eight when needed to operate with the m16n8k8 operation. We find CUTLASS’s m16n8k8 matrix multiplication with $M = N = K = 2048$ to achieve similar TFLOPs/sec to the highest reported on the T4 GPU [48].

Workloads. We consider workloads from multiple domains:

General-purpose CNNs. We consider eight widely-used CNNs from the popular PyTorch Torchvision library [16]: ResNet-50, VGG-16, AlexNet, SqueezeNet, ShuffleNet, DenseNet-161, ResNext-50, and Wide-ResNet-50. Each of these CNNs has 1000 output classes, as is standard for ImageNet. We primarily report performance when operating over HD images of size 1080×1920 with batch size of one, though we consider other image resolutions in §6.4.1.

Recommendation models. We consider Facebook’s DLRLM [66], which has two NNs consisting of fully-connected layers (also called multilayer perceptrons, or MLPs): MLP-Bottom, which has three hidden layers with 512, 256, and 64 nodes each, and MLP-Top which has two hidden layers with 512 and 256 nodes, and produces one output value. We primarily consider DLRLMs with batch size of one

Figure 9: Execution-time overhead on general-purpose CNNs with inputs of resolution 1080×1920 .

as this is the common case for low-latency, user-facing inference [32, 88]. For completeness, we also consider large batch size in §6.4.2.

Specialized CNNs. We also evaluate on NNs representative of ongoing efforts to deploy small NNs (described in §3.4). We consider four specialized CNNs used within the NoScope system [50]: Coral, Roundabout, Taipei, Amsterdam. These CNNs act as lightweight filters performing binary classification in front of large, general-purpose CNNs for high-throughput offline video analytics in cluster settings. These CNNs have 2–4 convolutional layers, each with 16–64 channels, at most two fully-connected layers, and operate over regions of video frames of size 50×50 pixels. As these CNNs are used for offline analytics, we use a large batch size of 64 for experiments.

Square matrix multiplications. We finally perform a more detailed comparison of one-sided thread-level ABFT and global ABFT, along with two-sided thread-level ABFT and thread-level replication on matrix multiplications with $M = N = K$ of various sizes (§6.5).

6.3 Summary of results

Figure 8 compares the execution-time overhead of global ABFT to that of intensity-guided ABFT on all NNs we consider (listed in order of increasing aggregate arithmetic intensity).⁵ Compared to global ABFT, intensity-guided ABFT reduces execution-time overhead by up to 5.3 $\times$. For example, for the Coral specialized CNN, intensity-guided ABFT reduces execution-time overhead from 17% to 4.6%. As expected, intensity-guided ABFT achieves the largest reduction in execution-time overhead for NNs with low aggregate arithmetic intensity, as these NNs contain more bandwidth-bound linear layers that benefit from thread-level ABFT. That said, intensity-guided ABFT reduces execution-time overhead considerably even for NNs with high aggregate arithmetic intensity. For example, intensity-guided ABFT reduces the execution-time overhead on Wide-ResNet-50 by 1.5 $\times$ compared to global ABFT (from 5.3% to 3.5%). Even though such NNs have high aggregate arithmetic intensity, they still contain bandwidth-bound linear layers, for which using thread-level ABFT over global ABFT reduces overhead.

6.4 Evaluation across various NN domains

6.4.1 General-purpose CNNs. Figure 9 shows the execution-time overhead for thread-level ABFT, global ABFT, and intensity-guided ABFT on eight popular general-purpose CNNs operating over HD images of size 1080×1920 at batch size one. Compared to global ABFT, intensity-guided ABFT reduces execution-time overhead by

Figure 10: Execution-time overheads on NNs from DLRM. Error bars are tight to the point of being imperceptible.

1.09–2.75 $\times$. As expected, thread-level ABFT obtains lower execution-time overhead than global ABFT for CNNs with low aggregate arithmetic intensity, while global ABFT has lower overhead for CNNs with higher aggregate arithmetic intensity. Intensity-guided ABFT obtains the lowest execution-time overhead across all the CNNs, motivating its per-layer, arithmetic intensity-guided approach.

Effect of image resolution. When operating on images of size 224×224 (the standard resolution in ImageNet), intensity-guided ABFT reduces execution-time overhead by 1.3–3.3 $\times$ compared to global ABFT. This larger reduction compared to operating on HD images stems from the lower aggregate arithmetic intensity of CNNs when operating on images with smaller resolution (described in §3.2). This leads to more linear layers being bandwidth-bound and benefiting from thread-level ABFT in intensity-guided ABFT.

6.4.2 Recommendation models (DLRM). We next consider the NNs used in Facebook’s DLRM. Figure 10 plots execution-time overheads on MLP-Bottom and MLP-Top. At batch size of one, which corresponds to low-latency deployments of DLRMs for user-facing services, both MLP-Bottom and MLP-Top have low aggregate arithmetic intensity. This results in intensity-guided ABFT reducing execution-time overhead compared to global ABFT by 4.55 $\times$ and 3.24 $\times$ for MLP-Bottom and MLP-Top, respectively. At a very large batch size of 2048, the aggregate arithmetic intensity of both MLP-Bottom and MLP-Top increase, but at different rates. The aggregate arithmetic intensity of MLP-Top increases from 7.7 to 175.8, resulting in the difference between global and thread-level ABFT decreasing. In contrast, the aggregate arithmetic intensity of MLP-Bottom grows only from 7.4 to 92, resulting in thread-level ABFT continuing to have lower overhead. In both cases, intensity-guided ABFT achieves the lowest overhead, illustrating the need for ABFT to consider the resource bottlenecks of each linear layer of a NN.

6.4.3 Specialized CNNs. Figure 11 shows the execution-time overheads on specialized CNNs from NoScope [50] at batch size 64. For these primarily bandwidth-bound CNNs with low aggregate arithmetic intensity, intensity-guided ABFT reduces execution-time

⁵Note that the execution-time overheads do not monotonically decrease with increasing aggregate arithmetic intensity since execution is performed layer-wise whereas the aggregate arithmetic intensity metric is not layer-wise.

Figure 11: Execution-time overheads on specialized CNNs.

Figure 12: Execution-time overhead on square matrix multiplications. Sizes left of the dashed line have arithmetic intensity below the T4’s FP16 CMR. The overhead for replication is above 70% for the final two sizes, and thus is cut off.

overhead by 1.6–5.3×. These results are particularly promising when considering the growing trends described in §3 of designing lightweight NNs, coupled with the increasing CMR of GPUs, which will likely result in more NNs being bandwidth-bound.

6.5 Evaluation of thread-level design decisions

To evaluate the design decisions made in leveraging thread-level ABFT for NN inference on GPUs, we now evaluate global ABFT and the various approaches to thread-level redundant execution described in §4 and §5. We perform such evaluation on square matrix multiplications (i.e., $M = N = K$) of varying size, allowing us to control arithmetic intensity and best illustrate the tradeoffs.

Figure 12 shows the execution-time overhead of each approach with $M = N = K$ ranging from 32 to 2048, corresponding to FP16 arithmetic intensities of 10 to 683. We first compare only the final version of thread-level ABFT we leverage (one-sided) to global ABFT. As expected, for matrix sizes with arithmetic intensity less than the FP16 CMR of the T4 (203), thread-level ABFT achieves an execution-time overhead up to 6.5× lower than that of global ABFT, while for matrix sizes with higher arithmetic intensity, global ABFT achieves overheads up to 14× lower than thread-level ABFT. It is clear that taking a one-size-fits-all approach to ABFT will lead to suboptimal performance on certain matrix sizes, motivating our adaptive approach in intensity-guided ABFT.

Figure 12 also shows that one-sided thread-level ABFT almost always exhibits lower execution-time overhead than two-sided thread-level ABFT and thread-level replication. This reinforces our decision to use one-sided ABFT for thread-level ABFT. The differences between replication and ABFT are particularly stark for larger sizes (512 and beyond), where the overhead of replication sharply spikes due to increasing competition for Tensor Cores.

7 DISCUSSION

7.1 Intensity-guided ABFT beyond NNs

While we have focused the design of intensity-guided ABFT for imparting fault tolerance to NN inference, intensity-guided ABFT is applicable to general matrix multiplication problems as well. We consider this particularly important as more traditional HPC applications begin exploring the use of NN hardware accelerators, such as Tensor Cores [33, 40, 42], and as NN hardware accelerators begin to add support for double- and single-precision floating point arithmetic [4], which are typically used in HPC applications.

7.2 Mathematical models for ABFT

In this work, intensity-guided ABFT leverages empirical profiling to make the final decision of which ABFT scheme to use for a given layer of a NN. An alternative is to leverage analytical models of ABFT with assumptions about compute and memory bandwidth to analytically determine which approach to ABFT will likely result in lower execution-time overhead. Intensity-guided ABFT could also leverage such models. We have chosen to use empirical profiling because this is the common practice used in optimizing NNs for inference by popular frameworks. Regardless of whether empirical profiling or analytical modeling is used, the core insights driving intensity-guided ABFT will remain relevant: layers with low arithmetic intensity relative to a GPU’s CMR are likely to benefit from thread-level ABFT, while layers with high arithmetic intensity relative to a GPU’s CMR are likely to benefit from global ABFT.

7.3 Input-size-dependent optimization

As noted in §3.2, the size of the input to a NN affects its arithmetic intensity, and thus the selection made by intensity-guided ABFT. Thus, one might wonder whether intensity-guided ABFT may be suboptimal if a deployment experiences changes in the size of inputs. This, however, is not a major concern because it is rare to have dynamically-sized inputs at inference time due to the common pre-processing step for NN inference of resizing inputs to a fixed format. The input-size-dependent heterogeneity of arithmetic intensity described in §3.2 will exist across deployments, but is unlikely within a single deployment. If multiple input sizes are expected within a deployment, one can handle this easily by performing separate ABFT selections for multiple input sizes and choosing among these at inference time depending on the size of an input.

8 RELATED WORK

Fault tolerance in general programs. There are various techniques for tolerating soft-error-induced faults in general programs:

One approach is hardware-based fault tolerance, such as through using ECC in memory, and radiation-hardened or redundant processing units [21, 34, 77]. While certain approaches to hardware-based fault tolerance are widely used, such as ECC-protected memory subsystems, hardware protection for processing units is less widely used due to its high overhead. Furthermore, hardware-based fault tolerance is inflexible to changes in the required fault tolerance of applications or the fault rate of operating environments. Thus, we focus on software-based fault tolerance.

Software-based fault tolerance for general programs is typically achieved through techniques like instruction duplication [63, 70], replication of threads/warps [36, 47, 80, 84], and compiler-driven re-execution [51, 60]. In contrast, we focus on using application-level features of NN inference to reduce the overhead of fault detection.

Fault tolerance in NNs. Recent works have illustrated the potentially-catastrophic effects of soft errors on NNs through fault injection tools [30, 31, 55, 61] and neutron beam experiments [38]. This has spurred many approaches for fault tolerance in NNs, such as leveraging the “inherent robustness” of NNs [68, 79, 87], training NNs to tolerate faults [52], anomalous activation suppression [29, 69], selective feature hardening [62], and learning to detect faults [59, 73, 74]. Our focus in this work is on leveraging ABFT to detect errors in NN inference. Compared to the approaches listed above, ABFT provides clearer fault-tolerance guarantees and does not require retraining a NN or understanding its behavior.

ABFT for NNs. Due to the heavy use of linear algebra in NNs, ABFT is a natural fit for fault tolerance in NNs, and a number of recent works have explored using ABFT for NNs [39, 43, 57, 67, 89]. Ozen et al. [67] leverage ABFT to protect convolutional and fully-connected layers and propose integration of ABFT into a systolic array architecture. Zhao et al. [89] propose a systematic workflow of ABFT checks for CNNs to provide a high degree of protection against faults with low execution-time overhead on CPUs. Li et al. [57] propose optimizations for ABFT in low-bitwidth DLRM inference on CPUs. Most closely related to our work is the work of Hari et al. [43], which proposes the optimized global ABFT scheme for GPUs (described in §2.5), and which forms a component of our proposed intensity-guided ABFT. Intensity-guided ABFT complements the work of Hari et al. [43] with ABFT schemes well-suited for bandwidth-bound linear layers, and by adaptively selecting between the two, using arithmetic intensity as a guide.

Compared to these works, the present work is unique in multiple aspects. First, the works listed above all focus on employing a single ABFT scheme across all linear layers of a NN. In contrast, we illustrate that different linear layers within a NN have varying resource bottlenecks that benefit from per-layer-optimization in the proposed intensity-guided ABFT. Second, to the best of our knowledge, our work is the first to analyze the growing bottleneck of memory bandwidth for NN inference in the context of exploiting it for efficient redundant execution. Careful analysis of this trend lends itself to developing optimizations that have been overlooked by prior works. Finally, to the best of our knowledge, this work presents the first thread-level approach to ABFT for NNs on GPUs.

ABFT in other domains. ABFT has been widely studied for imparting fault tolerance to linear algebra [23, 24, 46, 83, 85], iterative methods [26, 28], and other applications [58]. Our work differs from these works in its focus on the specific characteristics of NNs and GPUs, and its adaptive, intensity-guided approach of selecting ABFT schemes based on the resource bottlenecks of the problem.

Smith et al. [76] investigated fusing ABFT operations alongside matrix multiplications on CPUs. While similar to the approach to thread-level ABFT that we consider as part of intensity-guided ABFT, the techniques employed by Smith et al. [76] differ in that they do not perform ABFT at the level of the smallest unit of the parallel subproblem in the matrix multiplication. Thus, this approach generates checksums collaboratively across CPU threads

(although not globally), which requires additional loads and stores, albeit, at higher levels of the memory hierarchy. In contrast, we leverage thread-level ABFT specifically for bandwidth-bound linear layers in NNs on GPUs, and thus avoid performing any additional loads and stores (which would compete for the layer’s bottleneck resource). This results in thread-level ABFT performing ABFT at the smallest parallel sub-matrix multiplication solved (GPU thread level), requiring no coordination between threads. Furthermore, intensity-guided ABFT takes an adaptive approach to ABFT based on the resource bottleneck of a given matrix multiplication, whereas Smith et al. [76] use a one-size-fits-all approach.

Concurrent work with ours, FT-BLAS [86], proposes to choose between replication and ABFT in BLAS routines on CPUs depending on the BLAS level of an operation. Specifically, FT-BLAS [86] uses replication for operations in Levels 1 and 2 (vector-vector and matrix-vector operations), and ABFT for those in Level 3 (matrix-matrix operations). However, for a given operation in a BLAS level (e.g., for all matrix-matrix multiplications), FT-BLAS [86] takes a one-size-fits-all approach. In contrast, we show that the unique characteristics of NN inference on GPUs lead to *NNs containing a mix of compute- and bandwidth-bound matrix-matrix multiplications*, rendering one-size-fits-all approaches inefficient. Moreover, as shown in §6.5, leveraging replication even for bandwidth-bound matrix multiplications used in NNs on GPUs can lead to significant overhead, motivating intensity-guided ABFT’s approach of selecting between various ABFT schemes for each matrix multiplication.

9 CONCLUSION

We present intensity-guided ABFT, a new approach to ABFT for NN inference on GPUs that optimizes for the specific resource bottlenecks of individual layers of a NN. Through analysis of trends in NN design and GPU hardware, we present a case for a growing trend of compute underutilization in NN inference on GPUs, opening new opportunities for efficient redundant execution. However current approaches to ABFT for NNs are unable to exploit such fine-grained compute underutilization. We first carefully investigate a thread-level ABFT scheme to exploit such opportunities in bandwidth-bound linear layers of NNs on GPUs, complementing the use of traditional approaches to ABFT for compute-bound layers. Intensity-guided ABFT then adaptively selects among these ABFT schemes in a per-layer, arithmetic-intensity-driven fashion. This enables intensity-guided ABFT to reduce execution-time overhead by 1.09–5.3× across a number of popular and emerging NNs. Intensity-guided ABFT shows the promise of arithmetic-intensity-driven fault tolerance for current and future NNs. Finally, as intensity-guided ABFT protects general matrix multiplications, this approach may usher more efficient fault tolerance for broader HPC applications that are beginning to target NN hardware accelerators [33, 40, 42].

ACKNOWLEDGEMENTS

This work was funded in part by a National Science Foundation Graduate Research Fellowship (DGE-1745016 and DGE-1252522), in part by Amazon Web Services, and in part by the AIDA project (POCI-01-0247-FEDER-045907) co-financed by the European Regional Development Fund through the Operational Program for Competitiveness and Internationalisation 2020.

REFERENCES

- [1] CANDLE: Exascale Deep Learning and Simulation Enabled Precision Medicine for Cancer. <https://candle.cels.anl.gov/>. Last accessed 23 August 2021.
- [2] CUDA C++ Programming Guide. <https://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html>. Last accessed 23 August 2021.
- [3] ISO-26262 Road vehicles – Functional safety. <https://www.iso.org/standard/68383.html>. Last accessed 23 August 2021.
- [4] NVIDIA A100 GPU. <https://www.nvidia.com/en-us/data-center/a100/>. Last accessed 23 August 2021.
- [5] NVIDIA CUB. <https://nvlabs.github.io/cub/>. Last accessed 23 August 2021.
- [6] NVIDIA cuDNN. <https://developer.nvidia.com/cudnn>. Last accessed 23 August 2021.
- [7] NVIDIA CUTLASS T4 clock frequency setting <https://github.com/NVIDIA/cutlass/issues/154#issuecomment-745426099>. Last accessed 23 August 2021.
- [8] NVIDIA CUTLASS. <https://github.com/NVIDIA/cutlass>. Last accessed 23 August 2021.
- [9] NVIDIA Deep Learning Performance Guide. <https://docs.nvidia.com/deeplearning/performance/index.html>. Last accessed 23 August 2021.
- [10] NVIDIA DRIVE – Autonomous Vehicle Development Platforms. <https://developer.nvidia.com/drive>. Last accessed 23 August 2021.
- [11] NVIDIA Jetson. <https://developer.nvidia.com/embedded/jetson-modules>. Last accessed 23 August 2021.
- [12] NVIDIA Parallel Thread Execution ISA Version 7.2: Matrix Fragments for mma.m16n8k8 <https://docs.nvidia.com/cuda/parallel-thread-execution/index.html#warp-level-matrix-fragment-mma-1688>. Last accessed 23 August 2021.
- [13] NVIDIA Tensor Cores <https://www.nvidia.com/en-us/data-center/tensor-cores/>. Last accessed 23 August 2021.
- [14] NVIDIA TensorRT. <https://developer.nvidia.com/tensorrt>. Last accessed 23 August 2021.
- [15] NVIDIA Tesla P4 GPU Datasheet. <https://images.nvidia.com/content/pdf/tesla/184457-Tesla-P4-Datasheet-NV-Final-Letter-Web.pdf>. Last accessed 23 August 2021.
- [16] Torchvision Models. <https://pytorch.org/vision/stable/models.html>. Last accessed 23 August 2021.
- [17] CUDA Warps and Occupancy. https://on-demand.gputechconf.com/gtc-express/2011/presentations/cuda_webinars_WarpsAndOccupancy.pdf, 2011. Last accessed 23 August 2021.
- [18] NVIDIA Tesla V100 GPU Architecture. Tech. Rep. WP-08608-001_v1.1, 2017.
- [19] NVIDIA Turing GPU Architecture. Tech. Rep. WP-09183-001_v01, 2018.
- [20] Hewlett Packard Enterprise accelerates space exploration with first ever in-space commercial edge computing and artificial intelligence capabilities. <https://www.hpe.com/us/en/newsroom/press-release/2021/02/hewlett-packard-enterprise-accelerates-space-exploration-with-first-ever-in-space-commercial-edge-computing-and-artificial-intelligence-capabilities.html>, 2021. Last accessed 23 August 2021.
- [21] BARTLETT, W., AND SPAINHOWER, L. Commercial Fault Tolerance: A Tale of Two Systems. *IEEE Transactions on Dependable and Secure Computing* 1, 1 (2004), 87–96.
- [22] BLALOCK, D., ORTIZ, J. J. G., FRANKLE, J., AND GUTTAG, J. What is the State of Neural Network Pruning? In *The Third Conference on Systems and Machine Learning (MLSys 20)* (2020).
- [23] BOSILCA, G., DELMAS, R., DONGARRA, J., AND LANGOU, J. Algorithm-Based Fault Tolerance Applied to High Performance Computing. *Journal of Parallel and Distributed Computing* 69, 4 (2009), 410–416.
- [24] BRAUN, C., HALDER, S., AND WUNDERLICH, H. J. A-ABFT: Autonomous Algorithm-Based Fault Tolerance for Matrix Multiplications on Graphics Processing Units. In *2014 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 14)* (2014).
- [25] CAMPBELL, A., McDONALD, P., AND RAY, K. Single Event Upset Rates in Space. *IEEE Transactions on Nuclear Science* 39, 6 (1992), 1828–1835.
- [26] CHEN, J., LIANG, X., AND CHEN, Z. Online Algorithm-Based Fault Tolerance for Cholesky Decomposition on Heterogeneous Systems with GPUs. In *2016 IEEE International Parallel and Distributed Processing Symposium (IPDPS 16)* (2016).
- [27] CHEN, T., MOREAU, T., JIANG, Z., ZHENG, L., YAN, E., SHEN, H., COWAN, M., WANG, L., HU, Y., CEZE, L., GUESTRIN, C., AND KRISHNAMURTHY, A. TVM: An automated end-to-end optimizing compiler for deep learning. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)* (2018).
- [28] CHEN, Z. Online-ABFT: An Online Algorithm Based Fault Tolerance Scheme for Soft Error Detection in Iterative Methods. In *ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP 13)* (2013).
- [29] CHEN, Z., LI, G., AND PATTABIRAMAN, K. A Low-cost Fault Corrector for Deep Neural Networks through Range Restriction. *arXiv preprint arXiv:2003.13874v4* (2021).
- [30] CHEN, Z., LI, G., PATTABIRAMAN, K., AND DEBARDELEBEN, N. BinFI: An Efficient Fault Injector for Safety-Critical Machine Learning Systems. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC 19)* (2019).
- [31] CHEN, Z., NARAYANAN, N., FANG, B., LI, G., PATTABIRAMAN, K., AND DEBARDELEBEN, N. TensorFI: A Flexible Fault Injection Framework for TensorFlow Applications. In *2020 IEEE 31st International Symposium on Software Reliability Engineering (ISSRE 20)* (2020).
- [32] CHUNG, E., FOWERS, J., OVTCHAROV, K., PAPAMICHAEL, M., CAULFIELD, A., MASENGILL, T., LIU, M., LO, D., ALKALAY, S., HASELMAN, M., ET AL. Serving DNNs in Real Time at Datacenter Scale with Project Brainwave. *IEEE Micro* 38, 2 (2018), 8–20.
- [33] DAKKAK, A., LI, C., XIONG, J., GELADO, I., AND HWU, W.-M. Accelerating Reduction and Scan Using Tensor Core Units. In *Proceedings of the ACM International Conference on Supercomputing (ICS 19)* (2019).
- [34] DELL, T. J. A White Paper on the Benefits of Chipkill-Correct ECC for PC Server Main Memory. *IBM Microelectronics division 11* (1997), 1–23.
- [35] DENBY, B., AND LUCIA, B. Orbital Edge Computing: Nanosatellite Constellations as a New Class of Computer System. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 20)* (2020).
- [36] DIMITROV, M., MANTOR, M., AND ZHOU, H. Understanding Software Approaches for GPGPU Reliability. In *Proceedings of 2nd Workshop on General Purpose Processing on Graphics Processing Units (GPGPU 09)* (2009).
- [37] DIXIT, H. D., PENDHARKAR, S., BEADON, M., MASON, C., CHAKRAVARTHY, T., MUTHIAH, B., AND SANKAR, S. Silent Data Corruptions at Scale. *arXiv preprint arXiv:2102.11245* (2021).
- [38] DOS SANTOS, F. F., LUNARDI, C., OLIVEIRA, D., LIBANO, F., AND RECH, P. Reliability Evaluation of Mixed-Precision Architectures. In *2019 IEEE International Symposium on High Performance Computer Architecture (HPCA 19)* (2019).
- [39] DUTTA, S., BAI, Z., LOW, T. M., AND GROVER, P. CodeNet: Training Large Scale Neural Networks in Presence of Soft-Errors. *arXiv preprint arXiv:1903.01042* (2019).
- [40] FENG, B., WANG, Y., CHEN, G., ZHANG, W., XIE, Y., AND DING, Y. Egemmm-TC: Accelerating Scientific Computing on Tensor Cores with Extended Precision. In *Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP 21)* (2021).
- [41] GEIST, A. Supercomputing's Monster in the Closet. *IEEE Spectrum* 53, 3 (2016), 30–35.
- [42] HAIDAR, A., TOMOV, S., DONGARRA, J., AND HIGHAM, N. J. Harnessing GPU Tensor Cores for Fast FP16 Arithmetic to Speed up Mixed-Precision Iterative Refinement Solvers. In *International Conference for High Performance Computing, Networking, Storage and Analysis (SC 18)* (2018).
- [43] HARI, S. K. S., SULLIVAN, M., TSAI, T., AND KECKLER, S. W. Making Convolutions Resilient via Algorithm-Based Error Detection Techniques. *IEEE Transactions on Dependable and Secure Computing* (2021).
- [44] HOCHSCHILD, P. H., TURNER, P., MOGUL, J. C., GOVINDARAJU, R., RANGANATHAN, P., CULLER, D. E., AND VAHDAT, A. Cores that don't count. In *Proceedings of the 18th Workshop on Hot Topics in Operating System (HotOS 21)* (2021).
- [45] HSIEH, K., ANANTHANARAYANAN, G., BODIK, P., VENKATARAMAN, S., BAHL, P., PHILOPSE, M., GIBBONS, P. B., AND MUTLU, O. Focus: Querying Large Video Datasets with Low Latency and Low Cost. In *Proceedings of the 13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)* (2018).
- [46] HUANG, K.-H., AND ABRAHAM, J. A. Algorithm-Based Fault Tolerance for Matrix Operations. *IEEE Transactions on Computers* 100, 6 (1984), 518–528.
- [47] JEON, H., AND ANNAVARAM, M. Warped-DMR: Light-Weight Error Detection for GPGPU. In *2012 45th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO 12)* (2012).
- [48] JIA, Z., MAGGIONI, M., SMITH, J., AND SCARPAZZA, D. P. Dissecting the NVidia Turing T4 GPU via Microbenchmarking. *arXiv preprint arXiv:1903.07486* (2019).
- [49] KANG, D., BAILIS, P., AND ZAHARIA, M. BlazeIt: Optimizing Declarative Aggregation and Limit Queries for Neural Network-Based Video Analytics. *Proceedings of the VLDB Endowment* 13, 4 (2020).
- [50] KANG, D., EMMONS, J., ABUZALD, F., BAILIS, P., AND ZAHARIA, M. NoScope: Optimizing Neural Network Queries over Video at Scale. *Proceedings of the VLDB Endowment* 10, 11 (2017), 1586–1597.
- [51] KIM, H., ZENG, J., LIU, Q., ABDEL-MAJEED, M., LEE, J., AND JUNG, C. Compiler-Directed Soft Error Resilience for Lightweight GPU Register File Protection. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 20)* (2020).
- [52] KOPPULA, S., OROSA, L., YAĞLIKÇI, A. G., AZIZI, R., SHAHROODI, T., KANELLOPOULOS, K., AND MUTLU, O. EDEN: Enabling Energy-Efficient, High-Performance Deep Neural Network Inference Using Approximate DRAM. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO 19)* (2019).
- [53] KRIZHEVSKY, A., SUTSKEVER, I., AND HINTON, G. E. Imagenet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems (NIPS 12)* (2012).
- [54] LABEL, K. A. NASA and COTS Electronics: Past Approach and Successes-Future Considerations.
- [55] LI, G., HARI, S. K. S., SULLIVAN, M., TSAI, T., PATTABIRAMAN, K., EMER, J., AND KECKLER, S. W. Understanding Error Propagation in Deep Learning Neural

- Network (DNN) Accelerators and Applications. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC 17)* (2017).
- [56] LI, G., PATTABIRAMAN, K., HARI, S. K. S., SULLIVAN, M., AND TSAI, T. Modeling Soft-Error Propagation in Programs. In *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 18)* (2018).
- [57] LI, S., HUANG, J., TANG, P. T. P., KHUDIA, D., PARK, J., DIXIT, H. D., AND CHEN, Z. Efficient Soft-Error Detection for Low-Precision Deep Learning Recommendation Models. *arXiv preprint arXiv:2103.00130* (2021).
- [58] LI, S., LI, H., LIANG, X., CHEN, J., GIEM, E., OUYANG, K., ZHAO, K., DI, S., CAPPELLO, F., AND CHEN, Z. FT-iSort: Efficient Fault Tolerance for Introsort. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC 19)* (2019).
- [59] LI, Y., LI, M., LUO, B., TIAN, Y., AND XU, Q. DeepDyve: Dynamic Verification for Deep Neural Networks. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security (CCS 20)* (2020).
- [60] LIU, Q., JUNG, C., LEE, D., AND TIWARI, D. Compiler-Directed Lightweight Checkpointing for Fine-Grained Guaranteed Soft Error Recovery. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC 16)* (2016).
- [61] MAHMOUD, A., AGGARWAL, N., NOBBE, A., VICARTE, J. R. S., ADVE, S. V., FLETCHER, C. W., FROSIO, I., AND HARI, S. K. S. PyTorchFI: A Runtime Perturbation Tool for DNNs. In *2020 50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W 20)* (2020).
- [62] MAHMOUD, A., HARI, S. K. S., FLETCHER, C. W., ADVE, S. V., SAKR, C., SHANBHAG, N., MOLCHANOV, P., SULLIVAN, M. B., TSAI, T., AND KECKLER, S. W. HardDNN: Feature Map Vulnerability Evaluation in CNNs. *arXiv preprint arXiv:2002.09786* (2020).
- [63] MAHMOUD, A., HARI, S. K. S., SULLIVAN, M. B., TSAI, T., AND KECKLER, S. W. Optimizing Software-Directed Instruction Replication for GPU Error Detection. In *International Conference for High Performance Computing, Networking, Storage and Analysis (SC 18)* (2018).
- [64] MITTAL, S., AND VETTER, J. S. A Survey of Techniques for Modeling and Improving Reliability of Computing Systems. *IEEE Transactions on Parallel and Distributed Systems* 27, 4 (2015), 1226–1238.
- [65] MULLAPUDI, R. T., CHEN, S., ZHANG, K., RAMANAN, D., AND FATAHALIAN, K. Online Model Distillation for Efficient Video Inference. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV 19)* (2019), pp. 3573–3582.
- [66] NAUMOV, M., MUDIGERE, D., SHI, H.-J. M., HUANG, J., SUNDARAMAN, N., PARK, J., WANG, X., GUPTA, U., WU, C.-J., AZZOLINI, A. G., ET AL. Deep Learning Recommendation Model for Personalization and Recommendation Systems. *arXiv preprint arXiv:1906.00091* (2019).
- [67] OZEN, E., AND ORAIOGLU, A. Sanity-Check: Boosting the Reliability of Safety-Critical Deep Neural Network Applications. In *2019 IEEE 28th Asian Test Symposium (ATS 19)* (2019).
- [68] OZEN, E., AND ORAIOGLU, A. Concurrent Monitoring of Operational Health in Neural Networks Through Balanced Output Partitions. In *2020 25th Asia and South Pacific Design Automation Conference (ASP-DAC 20)* (2020).
- [69] OZEN, E., AND ORAIOGLU, A. Just Say Zero: Containing Critical Bit-Error Propagation in Deep Neural Networks with Anomalous Feature Suppression. In *2020 IEEE/ACM International Conference On Computer Aided Design (ICCAD 20)* (2020).
- [70] REIS, G. A., CHANG, J., VACHHARAJANI, N., RANGAN, R., AND AUGUST, D. I. SWIFT: Software Implemented Fault Tolerance. In *Proceedings of the International Symposium on Code Generation and Optimization (CGO 05)* (2005).
- [71] RUSSAKOVSKY, O., DENG, J., SU, H., KRAUSE, J., SATHEESH, S., MA, S., HUANG, Z., KARPATY, A., KHOSLA, A., BERNSTEIN, M., BERG, A. C., AND FEI-FEI, L. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)* 115, 3 (2015), 211–252.
- [72] SAXENA, N. The Road to Resilient Computing in Autonomous Driving is Paved with Redundancy. 2018 IEEE International Reliability Physics Symposium (IRPS) Keynote, 2018.
- [73] SCHORN, C., AND GAUERHOF, L. FACER: A Universal Framework for Detecting Anomalous Operation of Deep Neural Networks. In *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC 21)* (2020).
- [74] SCHORN, C., GUNTORO, A., AND ASCHEID, G. Efficient On-Line Error Detection and Mitigation for Deep Neural Network Accelerators. In *International Conference on Computer Safety, Reliability, and Security (SAFECOMP 18)* (2018).
- [75] SHEN, H., HAN, S., PHILIPPOSE, M., AND KRISHNAMURTHY, A. Fast Video Classification via Adaptive Cascading of Deep Models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 17)* (2017).
- [76] SMITH, T. M., VAN DE GEIJN, R. A., SMELYANSKIY, M., AND QUINTANA-ORTI, E. S. Toward ABFT for BLIS GEMM. In *Tech. Rep. TR-15-05*. The University of Texas at Austin, 2015.
- [77] SULLIVAN, M. B., HARI, S. K. S., ZIMMER, B., TSAI, T., AND KECKLER, S. W. Swapcodes: Error Codes for Hardware-Software Cooperative GPU Pipeline Error Detection. In *2018 51st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO 18)* (2018).
- [78] TAN, M., AND LE, Q. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In *Proceedings of the 36th International Conference on Machine Learning (ICML 19)* (2019).
- [79] TORRES-HUITZIL, C., AND GIRAU, B. Fault and Error Tolerance in Neural Networks: A Review. *IEEE Access* 5 (2017), 17322–17341.
- [80] WADDEN, J., LYASHEVSKY, A., GURUMURTHI, S., SRIDHARAN, V., AND SKADRON, K. Real-World Design and Evaluation of Compiler-Managed GPU Redundant Multithreading. In *2014 ACM/IEEE 41st Annual International Symposium on Computer Architecture (ISCA 14)* (2014).
- [81] WEISS, T. R. Microsoft, HPE Bringing AI, Edge, Cloud to Earth Orbit in Preparation for Mars Missions. <https://www.hpcwire.com/2021/02/12/microsoft-hpe-bringing-ai-edge-cloud-to-earth-orbit-in-preparation-for-mars-missions/>. Last accessed 23 August 2021.
- [82] WILLIAMS, S., WATERMAN, A., AND PATTERSON, D. Roofline: an Insightful Visual Performance Model for Multicore Architectures. *Communications of the ACM* 52, 4 (2009), 65–76.
- [83] WU, P., GUAN, Q., DEBARDELEBEN, N., BLANCHARD, S., TAO, D., LIANG, X., CHEN, J., AND CHEN, Z. Towards Practical Algorithm Based Fault Tolerance in Dense Linear Algebra. In *Proceedings of the 25th ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC 16)* (2016).
- [84] YANG, L., NIE, B., JOG, A., AND SMIRNI, E. Enabling Software Resilience in GPGPU Applications via Partial Thread Protection. In *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE 21)* (2021).
- [85] ZAMANI, H., LIU, Y., TRIPATHY, D., BHUYAN, L., AND CHEN, Z. GreenMM: Energy Efficient GPU Matrix Multiplication Through Undervolting. In *Proceedings of the ACM International Conference on Supercomputing (ICS 19)* (2019).
- [86] ZHAI, Y., GIEM, E., FAN, Q., ZHAO, K., LIU, J., AND CHEN, Z. FT-BLAS: A High Performance BLAS Implementation With Online Fault Tolerance. *arXiv preprint arXiv:2104.00897* (2021).
- [87] ZHANG, J., RANGINENI, K., GHODSI, Z., AND GARG, S. ThUndervolt: Enabling Aggressive Voltage Underscaling and Timing Error Resilience for Energy Efficient Deep Learning Accelerators. In *Proceedings of the 55th Annual Design Automation Conference (DAC 18)* (2018).
- [88] ZHANG, M., RAJBHANDARI, S., WANG, W., AND HE, Y. DeepCPU: Serving RNN-based Deep Learning Models 10x Faster. In *2018 USENIX Annual Technical Conference (USENIX ATC 18)* (2018).
- [89] ZHAO, K., DI, S., LI, S., LIANG, X., ZHAI, Y., CHEN, J., OUYANG, K., CAPPELLO, F., AND CHEN, Z. FT-CNN: Algorithm-Based Fault Tolerance for Convolutional Neural Networks. *IEEE Transactions on Parallel & Distributed Systems* 32, 07 (2021), 1677–1689.
- [90] ZHAO, W., ZHANG, J., XIE, D., QIAN, Y., JIA, R., AND LI, P. AIBox: CTR Prediction Model Training on a Single Node. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management (CIKM 2019)* (2019).
- [91] ZOPH, B., AND LE, Q. V. Neural Architecture Search with Reinforcement Learning. *arXiv preprint arXiv:1611.01578* (2016).

Appendix: Artifact Description/Artifact Evaluation

SUMMARY OF THE EXPERIMENTS REPORTED

We compared the execution-time overhead of arithmetic-intensity-guided ABFT and traditional approaches to ABFT on neural network inference workloads. Each ABFT version is implemented atop NVIDIA's open-source CUTLASS library (<https://github.com/NVIDIA/cutlass>).

We ran experiments on an NVIDIA T4 GPU on an AWS g4dn.xlarge instance. We ran the AWS instance atop the NVIDIA HPC SDK AMI and in a Docker container. We use CUDA 11.0 with NVIDIA driver version 450.51.06.

We evaluated on eight convolutional neural networks from the PyTorch torchvision library: AlexNet, ResNet-50, VGG-16, SqueezeNet, ShuffleNet, DenseNet-161, ResNext-50, and Wide-ResNet-50 (<https://pytorch.org/vision/stable/models.html>). We additionally evaluated on the multilayer perceptrons available used by Facebook's deep-learning-based recommendation models: MLP-Top and MLP-Bottom (<https://arxiv.org/abs/1906.00091>). We also evaluated on four convolutional neural networks from the NoScope project: Coral, Amsterdam, Taipei, and Roundabout (<https://www.vldb.org/pvldb/vol10/p1586-kang.pdf>).

We finally considered matrix-matrix multiplications of varying dimensions M , N , and K .

We decomposed each neural network into its individual convolutional/fully-connected layers and translated these layers into the matrix-matrix multiplications they use. We then profiled each matrix-matrix multiplication under the various ABFT schemes considered using the `cutlass_profiler` utility using 100 warmup iterations and 1000 profiling iterations.

Author-Created or Modified Artifacts:

Persistent ID: <https://github.com/Thesys-lab/arithmetic-intensity-guided-abft>
↪ `tic-intensity-guided-abft`
Artifact name: Implementation of
↪ arithmetic-intensity-guided ABFT atop NVIDIA's
↪ CUTLASS library

BASELINE EXPERIMENTAL SETUP, AND MODIFICATIONS MADE FOR THE PAPER

Relevant hardware details: NVIDIA T4 GPU, 4-core Intel(R) Xeon(R) Platinum 8259CL CPU @ 2.50GHz

Operating systems and versions: Ubuntu 18.04 running Linux kernel 5.30, NVIDIA HPC SDK AMI VM

Compilers and versions: NVCC release 11.0

Applications and versions: NVIDIA CUTLASS, commit `ccb697bac77fcc898e9c897b2c90aa5b60ac72fb`

Libraries and versions: CUDA 11.0

Key algorithms: matrix-matrix multiplication, algorithm-based fault tolerance