How Reduced Data Precision and Degree of Parallelism Impact the Reliability of Convolutional Neural Networks on FPGAs

F. LibanoG , P. RechG , *Senior Member, IEEE,* B. Neuman, J. Leavitt, M. WirthlinG , *Senior Member, IEEE,* and J. Brunhaver

Abstract-Convolutional neural networks (CNNs) are becoming attractive alternatives to traditional image-processing algorithms in self-driving vehicles for automotive, military, and aerospace applications. The high computational demand of state-of-the-art CNN architectures requires the use of hard- **ware acceleration on parallel devices. Field-programmable gate arrays (FPGAs) offer a great level of design flexibility, low power consumption, and are relatively low cost, which make them very good candidates for efficiently accelerating neural networks. Unfortunately, the configuration memories of SRAM-based FPGAs are sensitive to radiation-induced errors, which can compromise the circuit implemented on the programmable fabric and the overall reliability of the system. Through neutron beam experiments, we evaluate how lossless quantization processes and subsequent data precision reduction impact the area, performance, radiation sensitivity, and failure rate of neural networks on FPGAs. Our results show that an 8-bit integer design can deliver over six times more fault-free executions than a 32-bit floating-point implementation. Moreover, we discuss the tradeoffs associated with varying degrees of parallelism in a neural network accelerator. We show that, although increased parallelism increases radiation sensitivity, the performance gains generally outweigh it in terms of global failure rate.**

***lldex* Terms-Field-programmable gate array (FPGA), neural networks, parallelism, reduced precision, reliability.**

I. INTRODUCTION

IN RECENT years, the trillion dollar automotive industry has been very much focused on progressively adding technology to vehicles and on enabling self-driving capabilities fl]. Through the pairing of sensorial data collection with high-performance image processing , embedded computing systems used by companies like Mercedes-Benz [2]. Volvo [3]. and Tesla [4], are currently able to deliver semi-autonomous driving experiences. In fact, the push for autonomous vehicles is also very much true in military organizations (with the increased adoption of unmanned aerial vehicles (UAVs) [51) and in the space exploration sector (with the use of rovers and helicopters on Mars [6]. [7]). At the same time, convolutional neural networks (CNNs) have significantly evolved in terms of accuracy and became very attractive solutions for image processing and pattern recognition workloads [8]. Given that all of the aforementioned safety-critical applications depend on reliable computer vision [9]-111]. CNNs emerge as great alternatives to more old-fashioned algorithms.

Neural networks have a parallel structure, both in terms of neurons, and convolutional filters . Such inherent parallelism makes CNNs perfect candidates for efficient acceleration on parallel computing devices, such as application-specific integrated circuits (ASICs), graphics-processing units (GPUs) , and field-programmable gate arrays (FPGAs) . In a perfect world, we would be able to define an optimal hardware architecture, fabricate an ASIC, and never look back. However , the field of machine learning is ever evolving, which means that the requirements for accelerating the neural network topologie s of today might not be the same a few years in the future. Therefore, design flexibility and reprogrammability must be considered. GPUs currently offer parallel general-purpose arithmetic hardware and design flexibility at the software level. However , most state-of-the-art GPU architecture s are throughput-oriented [12] (rather than latency-oriented) and not power-efficient [13]. FPGAs, on the other hand , offer full hardware reprogrammability and lower power consumption [14]. which enables system patches (for architectural improvements, weight tuning, bug fixing , and more), along with the ability for deployment on power constrained missions (which is often the case in space applications) (15].

Un fortuna tely , FPGAs are very sens1t1ve to radiation-induced faults I 16]. More specifically, S RAM-ba sed FPGAs can experience single-event upset s (S EUs) in their configuration memory. which affects routing con nec tions as well as modi fie s the settings and content of LUTs. DSPs, BRAMs, and FFs. With a corrupted configuration memory . the design on the FPGA can start malfunc tioning

Manuscript received No vember 20, 2020; revise d December 27, 2020; accepted January 7. 2021. Date of publication January 11 , 2021; date of current version May 20, 2021. This work was supported in part by the Department of Energy of the United States , in part by the CAPES Foundation of the Mini st ry of Education, and in part by the CN Pq Research Counci I of the Ministry of Science and Technology.

F. Libano and J. Brunh aver are with the School of Ele ctrical. Computer and Energy Engineering (ECEE). Ar i zona State University (ASU), Tempe, AZ 85287 USA (e-mail: fli bano@asu.edu; jhrunhaver@ asu.edu).

P. Rech is with th e In s titut e of In formatics, Federal University of R io Grande do Sul (UF RGS). Porto Alegre 91 501- 9 70, Bra1.il (e-mail: prech@ inf.ufrgs.br).

B. Neuman is with the Los Alamos National Laborato ry 1 LANL), Los Alamos , NM X7544 USA (e-mail: bneu man@l anl.gov).

J. Leavitt and M. Wirthl in are with the Department of Electrical and Com - puter Engineer ing , Brigham Young Univ e rs ity (BYl J), Provo , UT 84602 USA (e-mail: jleavitt@ byu.edu; wi1thlin @by u.edu).

Co lor versions of one or more figure s in this article are avai lable at https://do i.or g/1 il. I 1ll9/f NS .21l21 . 31l50 707.

Digit al Object ldenti tier 10 1109/fNS.202 I . J 050707

00 I 8-9499 © 2021 IEEE. Pers on al use is permi tted, but republicatio n/redi stribution requires IEEE permission .
See http s://www. ieee.org/pu bli cation si rig hts/index. html for more infom1atio n.

and providing erroneous answer s. Furthermore , each type of logic resource available on the FPGA corresponds to a different percentage of the total configuration bits [17] and produces a different type of error/failure in the presence of SEUs. Therefore, the specific mapping of the circuit (i.e., the neural network accelerator) is an important component of overall resilience. More broadly speaking , it follows that reliability must always be taken into account when developing FPGA-based computational solutions for safety-critical applications .

Previous works have discussed the general reliability of CNNs executing in ASICs [18]. GPUs [19]. [20]. and FPGAs [21], [22]. Other studies have considered the tradeoffs in petformance and radiation sensitivity of using reduced and mixed precision in different computer architectures, while executing well-known benchmarks (including state-of-the-art neural networks for object detection) [23], [24]. Furthermore, the combined impact of accuracy (or lack thereof) and radiation on the overall failure rate of neural networks was discussed and modeled [25]. This article focuses on the reliability benefits of reduced precision in **CNNs,** whenever paired with proper quantization techniques (that allow for zero loss of accuracy). which is, to the best of our knowledge, a novel evaluation for FPGAs. We also evaluate architectural variations with majorly different degrees of parallelism and discuss the tradeoffs between area, speed, and reliability.

In order to deliver high accuracy, CNNs require high computational complexity. As an example, the number of arithmetic operations required for processing each input image in Google's state-of-the-art InceptionV3 architecture is over 5 billion [26]. On FPGAs, the need for higher computing power often translates to higher resource utilization, which in turn increases device cost, energy consumption, and sensitive area (i.e., radiation susceptibility) . In order to speed-up neural networks, designers have explored reducing the precision of data representation throughout the network . As a standard , state-of-the-art frameworks such as Google's TensorFlow use 32-bit floating point at training time, but recently developed quantization techniques have allowed the deployment of 16-bit floating-point and even 8-bit integer versions of same network, with little to no accuracy loss [27]. The first contribution of this article is an evaluation on how such data precision reduction impacts the overall reliability of CNNs. Our experimental data show that it is possible to decrease the failure rate of a neural network by over 70%, simply using proper quantization methods and a corresponding 8-bit integer representation.

Another important design decision for neural network hardware accelerators is the degree of parallelism. as the architect must define how many processing elements (PEs) are going to be instantiated in the programmable logic of the FPGA. Strictly from the performance standpoint, it seems that the answer would he *" as man y as possible ."* However , parallelism incr eases resource utilization, which in turn leads to increased radiation sensitivity , and, possibly , lower reliability. In addition , given the sizes of the current state-of-the-art neural network top olo gies, it is most often infeasible to implement *ALL* of the necessary PEs (providin g maximum parallelism). as the logic resources on FPGA devices are limited .

There fore, as the second contribution of this article, we analyze and discuss the area, performance , and reliability tradeoffs associated with varying degrees of parallelism for neural network accelerators. We show that radiation sensitivity can vary as much as two orders of magnitude , when comparing the two ends of the parallelism spectrum (i.e., minimum number of PEs versus maximum number of PEs), while the estimated failure rate is over 7 x lower in the faster, more parallel design.

The remainder of this article is organized as follows. Section II gives a more comprehensive background on neural networks, the available quantization tools, and the specifics of our case study **CNN.** Section **III** discusses the details of our FPGA designs, methodology, and radiation beam experiment. Section IV presents our results regarding data precision reduction and degree of parallelism . Section V goes over our main conclusions as well as intentions for future work.

II. BACKGROUND

A. Convolutional Neural Networks

Artificial neural networks (ANNs) are biologically inspired computational structures that are able to calculate their outputs by propagating data through sets of interconnected neurons, distributed across different layers. From a mathematical perspective. the output of a neural network is generated hy a series of matrix multiplications, which, in turn, are calculated through a series of multiply accumulated operations. A CNN is a special kind of ANN, mostly dedicated to image-processing tasks. The first few layers in a CNN are responsible for extracting features from the input image (such as edges and shadows) before feeding the layers of fully connected neurons (responsible for classification and decision-making). By using convolution and pooling operations, it is possible to reduce the amount of data to be processed by the neurons, ultimately making the computation more efficient. All filters and neurons have weights associated with them, which are learned during an extensive training process. Input data is iteratively presented to the model, and minor adjustments are made after each ste p. Once the training is complete, the network is able to compute solutions for novel inpu ts , based on its learned set of weigh ts.

B. Quantization of Neural Networks

In order to achieve high levels of accuracy, CNNs end up being very much computationally expensive. As safety-critical applications usually require low latency systems and algorithms, a number of simplification techniques have been developed for neural networks , such as weight trimming and quantization [28]. In this article, we focus on the latter. The main idea behind quantization is to reduce the precision in which the weights of a given model are represented , in order to speed-up computation and reduce resource uti lization . At training time, all industry-leading frameworks use 32-hit floating point as a default. For the specific case of TensorFlow, they provide a subset of tools called Tenso rFlow Lite [29], which allows developer s to quantize their model s to lower precisions, such as 16-bit floating point (IEEE 's half -precisi on) and 8-hit integer . Previ ous works have shown that careless

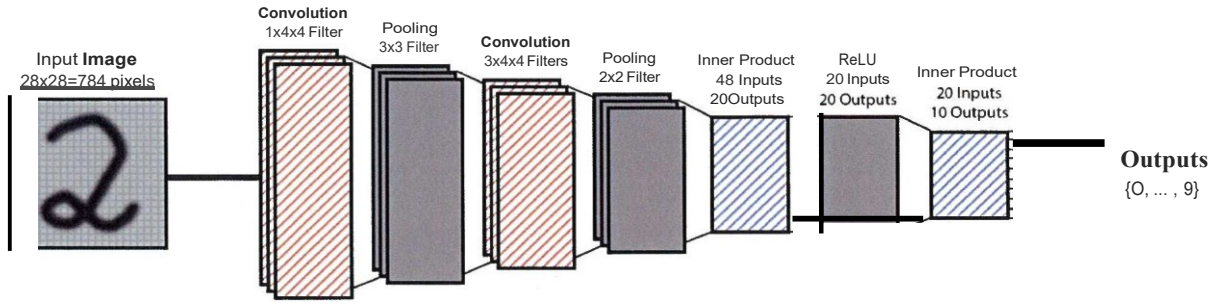

Fig. I. Topology of the MNlST CNN. Striped layers have weights associated with them (which go through the quantization process).

quantization procedures can often compromise accuracy (point at which the efforts in reducing radiation-induced .e rrors cease to pay oft) [25]. The clever thing about TensorFlow 's tools is that their quantization process uses a small subset of the training data for calibration purposes. The end result is that the quantized models show little to no accuracy loss when compared to their 32-bit float counterpart.

C. Related Work

As stated before, this article is not the first to evaluate the reliability of CNNs. Prior work has studied neural network resilience on ASICs. proposing and validating a framework for pre-silicon fault tolerance estimation, mostly based on topological characteristics [18]. Furthermore, GPU-focused papers have experimentally evaluated the reliability of state-of-the-art neural networks for object detection (with an emphasis on automotive applications), executing on different NVIDIA architectures. The authors have also tested an algorithm-based fault-tolerance (ABFf) technique for matrix multiplication, achieving significant error detection/correction rates [19]. [20]. Specifically on FPGAs, prior work has discussed the concept of error criticality and used extensive fault injection campaigns to identify differences in vulnerability across the layers in ANNs and CNNs . Additiona lly, a low-overhead selective hardening strategy was proposed, for scenarios in which traditional triple modular redundancy (TMR) is not possible due to limited resource availability [30]. Moreover , the correlation between faults in different logic resources of the FPGA and corresponding output errors was studied [31].

Finally. we must acknowledge that the reliability impact of reduced data precision on neural networks has also been explored before [24]. In such study, the authors have evaluated the different failure in time (FIT) rates of a CNN on a GPU, using double, single, and half-precision floating point. They have found that lower precision computation in GPUs reduces radiation sensitivity, while improving performance. Such finding further motivates our work. which intends to guide the design of reliable hardware accelerators for neural networks on FPGAs (by looking at precision and parallelism).

D. MN!ST

We have chosen the well-known Modified National Ins titu te of Standard s and Technology (MNIST) as our case study data set, since its simplicity translates to tractability for implementing and testing several design variations . We should, however, acknowledge that this aspect of our work limits the extent to which our experimental data can be fully generalized, as more complex CNNs could have different behaviors. The data set itself is composed by 60000 28 x 28 pixel images of handwritten decimal digits (from 0 to 9) [32]. As such, our neural network models receive 28 x 28 matrices as inputs and produce ten outputs (one for each decimal digit), where the index of the highest value corresponds to the classification.

Ill. EXPERIMENTAL METHODOLOGY

A. Designs Under Test

In order to evaluate the tradeoffs associated with data precision reduction, we have implemented three versions of the MNIST CNN (FP32, FPl6, INT8) using the 28-nm Zynq-7000 (XC7Z020) [33]. Table I and Fig. 2 detail resource utilization and execution times. In addition to that, we should point out that the accuracy on all of the design variations has stayed the same (95%), regardless of the quantization process performed with TensorFlow Lite. Although a higher accuracy level could have been achieved, we opted for a very minimalist CNN topology (and implementation), as detailed in Fig. 1.

In order to evaluate the tradeoffs associated with degree of parallelism, we have implemented two versions of the MNIST CNN (Min PEs, Max PEs) using the 16-nm Zynq UltraScale+ (XCZU9EG) [34]. The first design (Min PEs) has only one PE per layer of the network and represents the lower end of the parallelism spectrum. Likewise, the second design (Max PEs) has all of the necessary PEs in each layer of the network and represents the top end of the parallelism spectrum. In other words. if there is a total of N operations involved in a layer 's computati on, the single PE version is going to take N iterations to finish, while the fully parallel implementation (in this case. with N PEs), will take only one clock cycle. The aforementioned simplicity/tractability of the MNIST CNN, along with the use of a high-end FPGA (as the one found on the UltraScale+), allow it to be implemented in a fully parallel fashion. Table II and Fig. 3 provide details on resource utilization and execution times. Note that neither of the designs utilize DSPs. This is because both use 8-bit precision, which is too small for DSP inference at synthesis time. Also note that the axes in Fig. 3 are in logarithmic scale.

TABLE I

ZYNQ- 7000 R ESO URCE UTILIZATION TO IMPLEMENT THE MN!ST CNN USING 32-bit FLOATING POINT, 16-bil FLOATING POINT, AND 8-bil INTEGER

MNIST Design	LUTs (Logic)	LUTs (Mem)	FFs	DSPs
32-bit Float (FP32)	6.5k	1016	336	8
16 -bit Floa t (FPl6l	3.4k	510	240	4
8- bit Integer (INT8)	1.Sk	280	224	0

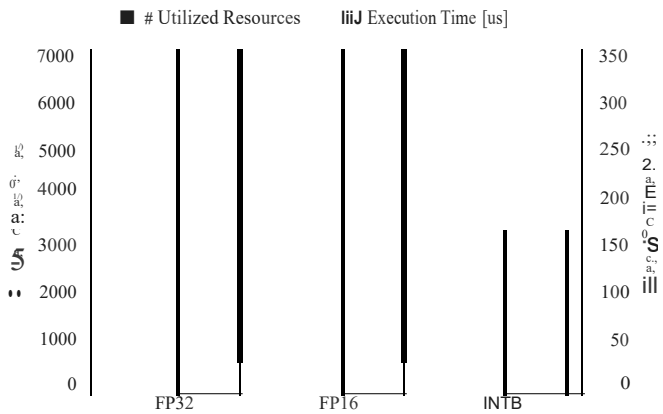

Fig. 2. Total resource utilization and execution times for three levels of data precision on the MNIST CNN. implemented on the Zynq-7000 .

We should point out, hefore moving forward, that the training phases and quantization processes for all of our neural networks were fully completed ahead of the experiments (in a fault-free environment) . This means that the neural networks ' set of weights was not affected by radiation-induced upsets in any capacity. Furthermore, we must also mention that the input stimuli for our designs under tests (DUTs) is a subset of LOO images from the **MNIST** data set (all of which are correctly classified in undisturbed executi ons).

B. *Radiation Experimental Setup*

Our radiation beam experiments were performed at the Los Alamos Neutron Science Center (LANSCE) facility of the Los Alamos National Laboratory (LANL). While LANSCE 's neutron spectrum mimics the atmospheric one, the particle flux is about 8 orders of magnitude higher than the average terrestrial flux [(13 n/ (cm^2 x h)] at sea level [35J). We ran our experiments for around 64 h, accumulating a total fluence in our DUTs of 344 x I 0 ^{9} n/ cm^2, roughly equivalent to 3 million years of natural expos ur e. The experimental setup , with the Zynq-7000 and the Zynq UltraScale+ is shown in Fig. 4. mounted at the beam line. We specifically chose Zynq devices because their heterogeneous nature considerably simplifies the setup. Using the ARM processors, we are able to provide stimuli to the FPGA and implement result checking . Since the DOR is protected with ECC, input/output corruption becomes extremely unlikely. Furthermore , given the trivial nature of the C program, we very rarely observe hangs during execution (in which case. it suffices to power cycle the device) . A server PC controls/logs the expe1iment from outside the beam room. We are also able to distinguish between transient and permanent ups ets: Whenever an output error is observed,

TABLE II

ZYNQ ULTRASCALE+ R ESO U RCE UT I LIZAT I O N TO IMPLEMENT THE MNIST CNN WITH VASTLY DIFFERENT DEGREES OF PARALLELISM

MNIST Design	LUTs (Logic)	LUTs (Mem)	FFs	DSPs
Min PEs	1.4k	280	240	0
Max PEs	212k	0	11.2k	0

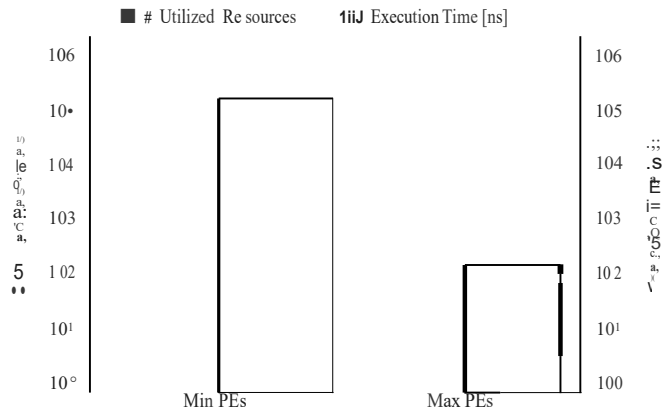

Fig. 3. Total resource utilization and execution times for two degrees of parallelism on the MNIST CNN, implemented on the Zynq UltraScale+.

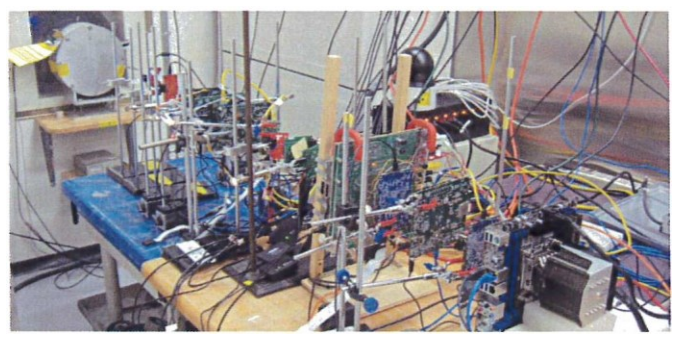

Fig. 4. Neutron beam experiment at the LANSCE facility of LANL, USA.

we run the same image again. If the second execution also presents corruptions. the upset is permanent. otherwise it is a single -event transient (SET) . However, we have not registered any of such transients in our neutron beam experiments.

C. *Ermr Criticality*

In most cases, **CNNs** are used for classification tasks, which means that not all errors need to be considered critical. In other words, even if the outputs computed by the network are different than the golden/expected ones, the classification of a given input image might still be correct. Hence, we identify two error classes in our neural networks:

1) *Tolerahle error:* The network's outputs differ from the expected ones, hut the image classification still comes out correct (i.e., the image classification is correct despite the output corruptions).
2) *Critical error:* The netw ork 's outputs differ from the expected ones, being severe enough to compromise the image classification (i.e.. an input image of the digit "7" is classified as the digit " l").

IV. EXPERIMENTAL RESULTS

A. Reduced Data Precision

Using the experimental setup and methodology desc1ibed in Section III, we have tested three versions of the MNIST CNN with varying levels of data precision (FP32. FP16. 1NT8). Fig. 5 plots the neutrons cross section of the three designs. We classify output errors as tolerable or critical, depending on whether or not the image classification was affected (as explained in Section lll-C).

As we have shown in Table I and Fig 2, lower precision hardware means lower resource utilization, which consequently means that the sensitive area of the FPGA for each design also gets reduced. Specifically, the 16-bit floating-point implementation uses about 40% less resources than the 32-bit version, while the 8-bit integer design decreases resource utilization by 64%. By analyzing Fig. 5, we can clearly see that, as we reduce precision, the probability of observing radiation-induced data corruptions diminishes. To be exact, the 16-bit floating-point version of the network has a 22% lower chance of producing errors at the output, when compared to the original 32-bit design. Similarly, when using 8-bit integers. we see a massive 72% cross section reduction from FP32.

A more nuanced result that can be drawn from our experimental data is that the rate of critical errors (as defined in Section Ill-C) is very different across the three designs. If we think about the way floating-point numbers are represented in computers (sign. exponent, and mantissa), it is easy to conclude that radiation-induced corruptions in the sign and in the exponent of a number can cause a much greater discrepancy . As the IEEE 754 standard establishes, the 32-bit floating-point representation uses 1 bit for the sign, 8 bits for the exponent. and 23 for the mantissa. while the 16-bit floating-point representation uses 1, 5, and 10, respectively. We can say that 28.13 % (9/32) of the bits in an FP32 word have a high potential of causing significant discrepancies between expected and computed outputs. while 37.50% (6/16) of the bits in an FP16 word could lead to considerable differences . In our CNN, this ultimately means that , as we reduce precision from FP32 to FP16 , the likelihood of an output error affecting the final classification of the input image increases. This result can even be intuitively generalized for a hypothetical FPS representation, which would very likely have a critical error rate higher than FP16. However, this pattern does not continue when we further reduce precision to 8-bit integers. This is because there are no exponent bits in an integer representation, so when a given bit *n* gets corrupted , the difference between expected and computed value can only be ± n2 , as opposed to a x2n discrepancy in an exponent corruption of a floating-point number. In our experiments, we have found that only 7% of the errors were critical in the FP32 design, while the error criticality rose to nearly 18% on the FP16 design and then fell back down to about 15 % on the INT8 version of the MNIST CNN. Regardless , 8-bit integer representation had the lowest absolute critical en-or rate out of all three implemen tations.

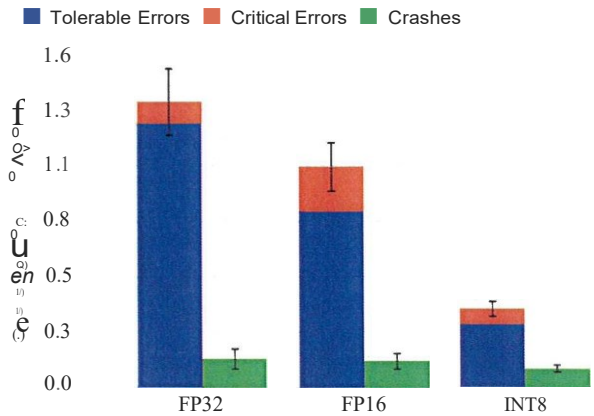

Fig. 5. Neutron cross section for the FP32. FPl6, and INT8 versions of the MNIST CNN.

Additionally, we have also observed instances where the CNN implemented on the FPGA did not provide any outputs after its expected execution time. This happens whenever the finite state machine (FSM) responsible for the control logic of the hardware gets stuck, failing to reach its *"done"* state. Such instances are shown in Fig. 5 as *Crashes.* As this type of event is much rarer than silent data corruptions (SDCs) (only accounts for 10%-20% of the total cross section), the error bars do not allow us to draw any conclusions from the experimental data. But, from an architectural perspective, we intuitively know that reducing the data precision only reduces the area occupied by the arithmetic pipeline and not the area occupied by the FSM, which means that the likelihood of observing crashes should be roughly the same across all three versions of the MNIST CNN.

The impact of reducing data precision can further be explored as we make use the notion of tolerated relative error (TRE) (24]. Fig. 6 basically shows how the neutron cross sections would be reduced if we were to allow a certain percentage of tolerance for discrepancies between expected and computed outputs. With a TRE of 0% , any bitflip in the output is considered an error. but as we increase the TRE, we staii to establish intervals of tole rance, instead of Boolean decisions. For example, with a TRE of 1%. any output corruption from 99% to 101% of the expected/golden value would still be considered correct. Interestingly, with a TRE level of only 1%, the cross section of tolerable errors on the 32-bit floating-point version of the CNN would be reduced by 43 %. This is because, as we previously discussed, 23 out of 32 bits are reserved for the mantissa , which means that most of the radiation-induced corruptions will not have a very significant impact on an FP32 word. Extending the comparison, if we were to use the same 1% tolerance interv al, the error rate on the FP16 implementation would only improve by 8%, while the 8-bit integer version would not change at all. It can then be said that , by treating neural networks as the inherently approximate computing units that they are, the perceived/effective radiation sensitivity can be significantly reduced. The caveat is that TRE only helps reducing the tolerable portion of the cross section, as the image classification

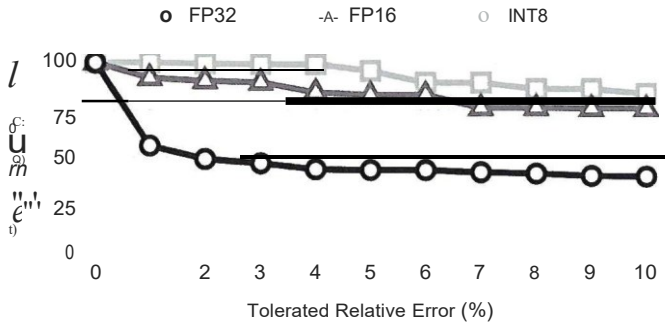

Fig. 6. Reduction in the cross sections of the FP32, FP16, and 11'17'8 versions of the MNIST CNN, when establishing incremental tolerance intervals.

TABLE Ill

MEBRFs FOR THE FP32. FP1 o, AND INT!l VERSIONS OF THE MNIST CNN

	FP32	FPl6	INT8
MEBRF [x 10^{15}]	0.61		4.08

depends on the comparison between all ten outputs, as opposed to their relationship with gold.

In order to get a realistic estimate of the failure rate of an algorithm in a radioactive environment, the mean executions between failures (MEBF) metric is commonly utilized [361, given that it takes into account both radiation sensitivity as well as performance. However, neural networks fall into a special category of applications, since they have an associated level of accuracy to them. Previous works have modeled the overall failure rate of neural networks (considering both inaccuracy and radiation as sources of faults) [25]. as they analyzed binary quantization, which is rarely lossless . Since our CNN topology is simple, and our precision reduction was not so extreme, accuracy was maintained across all DUTs at 95%. Thus, we focus our analysis strictly on the effects of radiation, using the mean executions between radiation failures (MEBRF) metric. We show the MEBRF of our designs in Table lll. conside1ing the neutron flux at sea level [35] and the entire SDC cross section measurements (tolerable+critical errors). as legislation typically does not distinguish between error categories [37). Evidently , the INT8 version of the MNIST CNN is able to complete a much higher number of failure-free executions than the floating-point implementations. To be exact. FP32 experiences over six times the failure rate of **INT8,** while **FP**16 fails over three times as much as the integer-based design. This is because reduced precision hardware not only occupies less area. but is also faster. The end result is lower radiation sensitivity and higher reliability, a true win-win situation (provided that accuracy remains stable).

B. Degree of Parallelism

Using the experimental setup and methodology described in Section Ill. we have tested two versions of the MNIST CNN with varying degrees of parallelism (Min PEs, Max PEs). As previously mentioned, the two designs in this analysis represent the two ends of the parallelism spectrum. On the one hand , we have a very sma ll, iterative design, which uses few resources (therefore, occupies less area), and takes longer

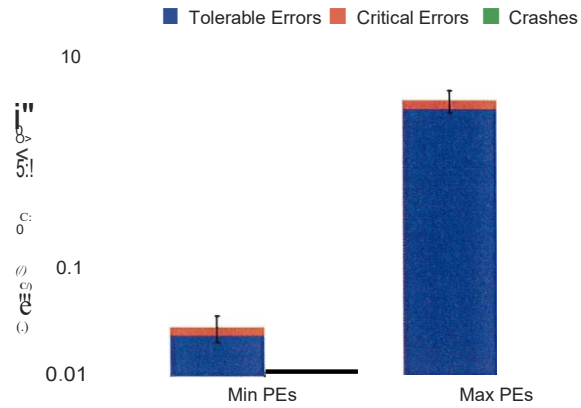

Fig. 7. Neutron cross section for the Min PE and l'vla PE versions of the MNIST CNN. Note that the y-axis is in logarithmic scale.

to complete the processing of input images (i.e ., has lower performance) . On the other side, we have a very large, fully parallel implementation, which occupies almost 100% of a state-of-the-art FPGA, but delivers extremely low latencv . Such scenario was made evident in Section in-A (Table i1 and Fig. 3), but we believe that it is worth reiterating.

Fig. 7 plots the neutron cross section of the two designs. We classify output errors as tolerable or critical, depending on whether or not the image classification was affected. Noting that the y-axis is in logarithmic scale, and the difference in radiation sensitivity between the two implementations is striking: the fully parallel design is l 33 times more likely to experience radiation-induced en'Ofs. Interestingly. it uses 130 times more resources than the version with far less PEs, confirming the direct relationship between resource utilization and cross section in SRAM-based FPGAs.

Furthermore , we should point out that the percentage of critical errors observed in our experiments was roughly the same in the two design variations (16% and 17 %, for Min PEs and Max PEs, respectively) . As there is no variation in data precision here, the impact of data corruption during computation is about the same in both cases, so this was an expected outcome. Finally, we can see that the more parallel version of the **MNIST CNN** did not experience any crashes. This is because, by being a fully pipelined. streaming architecture, it does not have/need any sort of FSM for control logic and therefore it never gets stuc k.

Again. we should emphasize that the cross section onlv measures the level of radiation sensitivity of an alg;-rit hm/circuit/ device. ln order to get an estimation of failure rate, one must also consider a performance metric (execution time) as a variable. Thus , we show the MEBRF of our DUTs in Table IV. considering the neutron flux at sea level [35]. Surprisingly, despite having a much higher cross section, the fully parallel implementation of the network is able to complete over seven times more error-free executions than the smaller design. This is because, as reported in Section III-A (Fig. 3), the more parallel design is over IOOO x faster.

Different than the analysis in Section IV-A , where the precision reduction led to smaller and faster hardware , we are seeing that the degree of parallelism in an architecture is

TABLE IV
MEBRFs FOR THE MIN PE AND MAX PE VERSIONS OF THE MNIST CNN

	Min PEs	Max PEs
MEBRF [x 10^{17}J	0.65	5.00

always a tradeoff between area and performance, which both directly impact the overall reliability of the system. From our experimental data, it seems that a faster (more parallel) design is the one that will deliver the lower failure rate, since the performance gains are outweighing the increased area and radiation sensitivity. However, as we said before, our case study CNN was specifically chosen to enable this analysis, and for most state-of-the-art neural network architectures, instantiating all of the necessary PEs in the FPGA would be unfeasible. Therefore, it follows that the optimal neural network hardware accelerator architecture (from the reliability standpoint) should really just be as parallel as possible.

C. *Technology Node*

Even though showcasing the difference between FPGAs in different technology nodes was not one of the main goals of this article, the INT8 circuit that was tested on the Zynq-7000 happens to be exactly the same as the Min PE design that was tested on the Zynq UltraScale+. Therefore. we have decided to plot and report their dynamic neutron cross sections in the same graph (Fig. 8). We can see that the older 28-nm CMOS technology is around one order of magnitude more sensitive to the radiation of the newer 16-nm FinFET technology, which, considering statistical errors, is in line with the static cross section numbers reported by Xilinx [38].

A number of previous works have dived into the exploration of differences in radiation sensitivity across technology nodes, through a mixture of charge collection simulations, and real-world beam experiments with neutrons/heavy ions (39], [40]. As such, our experimental data merely corroborates.

V. CONCLUSION

We have seen that reducing the data precision representation in CNNs, through state-of-the-art machine-learning frameworks that allow for little to no accuracy loss on their quantization processes, can significantly improve the overall reliability of safety-critical applications that rely on image processing. In summary, using a lower precision simplifies the hardware implemented on the FPGA, lowering its resource utilization, which means that it becomes less likely to get hit by impinging particles. At the same time , reduced precision hardware is usually faster, which further contributes to lowering the failure rate of neural networks. In addition to that, we have explored how the concept of error criticality affects the reliability analysis of a CNN intended for a classification task and how the adoption of

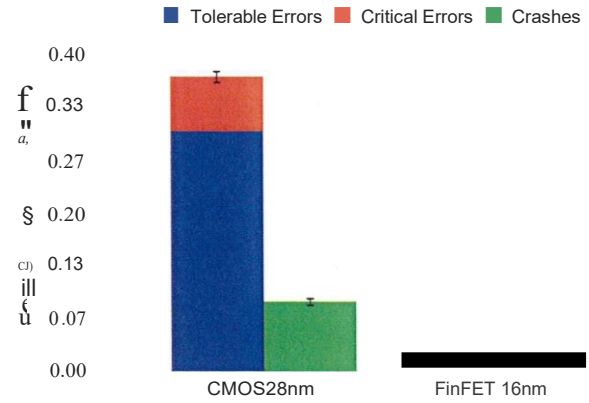

Fig. 8. N utron cross section for the MNIST CNN on the 28-nm Zynq-7000 and the I6-nm Zynq UltraScale+.

tolerance intervals can significantly impact the overall reliability of an application. ln general, we have concluded that. as long as the accuracy level remains the same after the quantization, the designer of an efficient hardware accelerator should opt for using 8-bit integer as opposed to floating-point representations.

Furthermore, we have analyzed how different degrees of parallelism (and the associated tradeoff between area and performance) affect the reliability of neural network accelerators on FPGAs. We have shown that, in general, the performance gains obtained through higher parallelism overshadow the cost paid in increased circuit area. Although the studied maximum degree of parallelism is often out of reach for most state- of-the-art CNN topologies, our results point to the fact that hardware accelerators for neural networks should be as parallel as possible. for improved reliability.

We believe that our findings regarding data precision and parallelism can offer some guidance for designing efficient and reliable hardware accelerators for neural networks in the future. As we have stated before, the machine learning field is in constant evolution, which means that the computing requirements for accelerating today's CNNs might not be the same in a few years' time. Our work can help novel accelerator architectures to be built with reliability in mind, ultimately enabling the deployment of neural networks in safety-critical applications. Finally. it is worth mentioning that we did not see any differences in error likelihood across classes in our 100-image input set. As a by-product of our experiments, we were also able to compare the radiation sensitivities of FPGAs in different technology nodes. Our experimental data corroborate with prior stu dies.

As future work , we intend to evaluate the reliability of Xilinx 's proprietary deep-processing unit (DPU) [41]. along with their recently released set of tools Yitis Al, which enable for quantization and compilation of neural network models straight out of industry standard frameworks such as TensorFlow. We also intend to analyze and improve (through adoption of efficient hardening techniques) the reliability of systolic aITay structures for matrix mu ltiplication, which seems to be the current choice of architectural core for both Xilinx 's DPU and Google's tensor-processing unit (TPU) (42] .

REFERENCES

[I j McKinsey. (Dec. 25, 2020). *Automotive Rcvol,1tio n- Perspecti1·e Towards 203V.* ronline J. Available: https ://www.mck i nsey com/ in dus tries/automotive-and-assen1b l y/our -i nsight %s/d is mptive-tre nds- that-will-transforn1-the-auto-industry/de-de#

[2] Mercedes-Benz . (Dec . 25, 2020). *Mercedes-Benz Innovation: Autonomous.* [Online]. Available: https://www.mercedes-benz.com/en/innovation/autonomous/

[Jj Volvo. (Dec. 25. 2020). *Autonomous Drivingllntellisa(c* ronlim:J. Available : htt ps://www. volvoc ars.com/e n-kw /own/own-and-enjoy/autonomous-driving

[4] Tesla. (Dec. 25, 2020). *Autopi /0 1.* [Online[. Available: https:!/www. tesla.com/autopilot

[5 j T. Samad. J. S. Bav , and D. Godbole . "Network-cenllic systems for military perations in urban terrain: The role of UAVs," *Proc. IEEE ,* vol. 95, no. I, pp. 92- 107 , Jan . 2007 .

[6] NASA. (Dec. 25, 2020). *Mars 2020 Missim, Perseverance Rmver.* [O nline [. Available: https://mars.nasa.gov/mars2020/

[7] NASA . (Dec. 25, 2020) *Mars Helicupler.* ro nl i ne j. Avai lable : https://mars.nasa.gov/technology/helicopter/

[81 H.-C. Shin *el al.,* "Deep convolutional neural networks for computer-aided detection: CNN arch ite ctures , dataset characteristics and transfer learni ng." *IEEE Trans . Med. /mag. ,* vol. 35. no . 5, pp . 1285- 1 298, May 2016.

[9] V.-E. Neagoe, A.-D. Ciotec, and A.-P. Barar, "A concurrent neural network approach to pedestrian detection in themrnJ image ry," in *Proc. 9th Int. Co11f Commun. (COMM),* Jun. 2012, pp. 133-136.

[10] M. Zhang, H. Li, G. Xia, W. Zhao, S. Ren , and C. Wang, "Research on the application of deep learning target detection of engineering vehicles in the patrol and inspection for military optical cable lines by UAV." in *Proc. /Ith Int. Symp. Co111 p111. l n1el/ . Design (ISCID),* Dec. 2018 , p p. 97-101.

[11J H. R. Kerner, K. L. Wagstaff, B. D. Bue, P. C. Gray, J. F. Bell, and H. Ben Amor , ·'Towar-d generalized change detection on planetary surfaces with convolutional autoencoders and transfer le arning," *IEEE J. Se/. fopics Appl. Eanh Observ. Remote Sens.,* vol. 12. no . 10 , pp. 39(Xl-3918, Oct. 2019 .

[11] M. Andersch, J. Lucas, M. A. LvLvarez-Mesa, and B. Juurl i nk, ··on latency in GPU throughput microarchitectures." in *Proc. IEEE Int. Svmp. Pei.form. Anal. Sys1. Softw. (ISPASS),* Mar. 2015, pp. 169-170.

[13J A. Aldahlawi. Y.-B. Kim. and K. K. Kim, "GPU architecture optimiza- tion for mobile compu ting ," in *Proc. In t. SoC Design Con/. (ISOCCi.* Oct. 2019, pp. 247-248.

[14] M. Qa, aimeh, J. Zan1tireno, P. H. Jones. K. Denolf , J. Lo, and K. Vissers, "Analyzing the energy-efficiency of vision kernels on embed- ded CPU. GPU and FPGA platfo nns," in *Proc. IEEE 27th Ann11. l n1. Symp. Field-Program. Custom Comput. Mach. (FCCM).* Apr. 2019, p. 336.

[15] T. Piovesan. H. C. Sa11o ri, J. E. Baggio, and J. R. Pinheiro. ··cutieSat electrical power supplies optimization---Comparison between conven- tional and optimal design methodology. " in *Proc. 12th IEEE /Ill . Con/. Ind. Appl. (I NDUSCON),* Nov. 2016, pp. 72-78.

f16] M. Wirthl in , "High-reliability FPGA-based syste ms: S pace, high-energy physic s, and beyond." *Proc. IEEE,* vol. 103 , no. 3, pp. 379- 389 , Mar. 2015 .

[17] Xilinx. (De c. 25. 2020). *7 Series FPGAs Configura1ion User G||ide.* [Online]. Availa ble : http s://www .xi linx.com/suppo rt/doc umentat io n/user_guides/ug470_7Series _%Con fig .pdf

[.18J B. Reagen *et al.,* "Ares: A framework for quantifying the res ilience of deep neural networks," in *Proc. 55th ACM/ESDNIEEE Design A11rom. Con{ (DAC).* New York, NY, USA , Jun. 2018. p. 17.

(191 F. Fe rnande s dos Santo s. L. Draghetti. L. Weigel. L. Carro . P. Navaux , and P. Rech , "Evaluation and mitigation of so ft-errors in neural network-based object detection in three GPU archi tect ures ," in *Proc. 47th Annu. /EEEIIFIP Int. Conj Dep<'ndahl,• Svsr. Ncrw. Workshop,·· (DSN-W),* Jun . 2017, pp. 169- 176 .

[20J F. F. dos San to s *er al.,* " Analyzing and increasing the rel iab ili ty of convolutional neural netw orks on GPU,," *IEEE Tran s. Rel.,* vo l. 68, no . 2, pp. 663--677 , Jun . 2019.

[21] F. Libano. P. Rech. L. Tambara , J. Tonfat, and F. Kastens m.idt, " On the reliability of linear regression and pattern recognition feedforward artificial neural networks in FPGAs." *IEEE Trans. N11c:l.* S c i., vol. 65, no . I. pp. 288- 295. Jan. 2018 .

[221 F. Benevenu ti, F. Libano, V. Pougel. F. L. Kas tens midt, and P. Rech , "Comparat ive analysis of inference errors in a neural network imple- mented in SRAM-based FPGA induced by neutron irradiation and fault injection methods," in *Proc. 31st Symp. lntcgr. Cirrnits Sys/. Design (SRCCI).* Aug . 2018, pp. 164- 169 .

r23 J F. Fernandes dos Santos, C. Lunard i, D. Ol ivei ra, F. Libano. and P. Rech , " Reliability evaluation of mixed-precision architectures .·· in *Proc. IEEE Int. Symp. High Perform . Comput. Archit. (HP CA).* Fe b. 2019. pp. 238- 249.

[241 F. F. dos San tos . P. Navaux , L. Carro , and P. Rech , ·'Impact of reduced precision in the reliability of deep neural networks for object detec tion;·· in *Proc. IEEE E111: Test Symp. (£TS),* May 2019, pp. 1 27- 13 2.

r2sJ F. Libano, B. Wilson, M. Wirthlin, P. Rech, and J. Bmnhaver. '·Under- standing the impact of quantiza tion , accu racy, and radiation on the reliability of convolutional neural networks on FPGAs ," *IEEE Trans . Nucl. Sci.,* vol. 67. no. 7, pp. 1478- 1484 , Jul. 2020.

[26] S. Bia nco . R. Caden,\ L. Celona , and P. Napoletano, "Benchmark analysis of representative deep neural network architectures." *IEEE Access,* vol. 6. pp. 64270-6 4277 , 20 18 .

[27] TensorFlow. (Dec . 25, 2020). *TcnsorFlow Model Optimi:ation Toolki1-Post-Trai11ing Integer Quantization .* [Online]. Available: https:// blog.tensorflow.org/20I 9/06/ tensorflow-integer -quant ization .htm %1

(18] 1. Hubara, M. Courbariaux , D. Soudry, R. El-Yaniv, and Y. Bengio , "Quantized neural networks : Training neural networks with low pre- cision weights and activations ," *J. Mach. Learn. Res.,* vol. 18 , no. I. p. 6869-6898, Jan. 2017.

[291 TensorFlow. (Dec . *15,* 2020). *TensorFlow lite.* [Onl ine]. Availa ble: https :!/www .te nsor flow.org /lite

[301 F. Libano *et al.,* "Selective hardening for neural networks in FPG As,'· *IEEE Trans. Nucl. Sci.,* vol. 66, no. I. pp. 216-221, Jan. 2019 .

[31J B. Du. S. Azimi, C. de Sia. L. Bozzoli. and L. Sterpo ne, "On the reliability of convolutional neural network implementation on SRAM- based FPG A," in *Proc. IEEE /111. Symp. Defect Fault Tolerance \!LS I Nanotechnol. Syst. (D FT),* Oct. 2019. pp. 149- 154.

r32J Y. Lecun. L. Bottou. Y. Bengio , and P. Haffner , "Gradient-based learn- ing applied to document reco gnition, " *Proc. IEEE.* vol. 86, no. II, pp. 2278-2324, Nov . 1998 .

[33] Xilinx. (Dec. 25, 2020). *Zynq-7000 SoC.* [O nline]. Available: hnps ://www.x ilinx .com/p roduct s/silicon-devices /soc/zynq-7( )(X).ht ml

(34] Xilinx. (Dec. 25, 2020). *Zynq UltraScale+ MPSoC.* [Online]. Available: https://www.xilinx.com/products/silicon-devices/soc/zynq-ultrascale- mps %oc.h tml

[35] JEDEC. (Dec. 25, 2020). *Jesd89a. JEDEC Standard.* Accessed: 2006. ronlineJ. Available: https: //www .jedec .org/sites/de faul t/tiles/ docs/ jesd89a.pdf

[36] P. Rech. L. L. Pilla. P. 0. A. Navaux. and L. Carro , "Impact of GPUs p allelism management on safe ty -cri tical and HPC applications re l iabilit y." in *Proc. 44th Annu. IEEE/IFIP Int. Con/. Depcndahlc Syst. Ncn, •.* , Jun. 20 14 , pp. 455-4 66.

137] ISO. (Dec. 25, 2020) *ISO 26262.* [On li ne]. Availa tile: h ttps: //www . iso.or g/obp/ui/#iso:std :iso :1626 2:-2:ed-2: v I :en

[38] Xilinx . (Dec. *15* , 2020). *Device Reliabilit y Rcpon .* ro nline j. Avai lable: https: //www .xi linx.com /s uppo rt/documen tation /use r_guide s/ug I 16.pdf

r39J Y.-P. Fan!! and A. S. Oates. "Neut ron-in duced charge collection sim- ulation of bulk FinFET SRAMs compared with conventional planar SRAM s," *IEEE Trans . Device Marer. Rel.,* vol. 11 , no . 4 , pp. 55 1- 554 , Dec. 2011.

r40] P. Nse ng iyumva *et al.,* ·'A co mpar is on of the SEU response of planar and FinFET d flip-flo ps at advanced technology node s," *IEEE Tran s. Nu c/.* Sci.. vol. 63, no. I, pp. 266- 272, Feh. 2016 .

[41] Xilinx. (Dec. 25. 2020). *Zynq DPU.* [Online] . Avai lable : hnps ://www. xi linx.com /s uppor t/docume ntat ion/i p_docume ntat io n/dpu/v 3 %_ I / pg338- dpu.pdf

[42] N. P. Jouppi *et al,* "In-datacenter perforn1ance ana ly sis of a tensor proces sing un it ," in *Proc. ACM/IEEE 44th Annu Int. Symp. Comput. Archir. (!S CA).* Jun. 2017, pp. 1- 12 .