

Efficient and Mathematically Robust Operations for Certified Neural Networks Inference

Fabien Geyer
Airbus Central R&T
Munich, Germany

Johannes Freitag
Airbus Central R&T
Munich, Germany

Tobias Schulz
Airbus Central R&T
Munich, Germany

Sascha Uhrig
Airbus Central R&T
Munich, Germany

Abstract—In recent years, machine learning (ML) and neural networks (NNs) have gained widespread use and attention across various domains, particularly in transportation for achieving autonomy, including the emergence of flying taxis for urban air mobility (UAM). However, concerns about certification have come up, compelling the development of standardized processes encompassing the entire ML and NN pipeline. This paper delves into the inference stage and the requisite hardware, highlighting the challenges associated with IEEE 754 floating-point arithmetic and proposing alternative number representations. By evaluating diverse summation and dot product algorithms, we aim to mitigate issues related to non-associativity. Additionally, our exploration of fixed-point arithmetic reveals its advantages over floating-point methods, demonstrating significant hardware efficiencies. Employing an empirical approach, we ascertain the optimal bit-width necessary to attain an acceptable level of accuracy, considering the inherent complexity of bit-width optimization.

I. INTRODUCTION

The last few years have seen an increasing use of machine learning (ML) and neural networks (NNs) in many domains, mainly due to their good results on challenges where no other solutions has yet been found. One important domain is transportation, where NNs are seen as the way forward for bringing autonomy, not only for cars, but also for other modes of transportation such as flying taxis for urban air mobility (UAM) or more generally aviation.

Various concerns about certification of ML and NNs have been raised by the aeronautical regulation and certification bodies [1, 2, 3, 4], especially for functions seen as safety critical. Similar concerns have also been raised in other transportation industries as shown by a recent survey [5]. To overcome these concerns, methods for certification of the complete ML and NNs pipeline are currently under development [6, 7] and will be standardized in the future [8]. The certification process covers the aspects of a ML pipeline, including the software aspects such as data verification and validation, training, inference, and tests. Of course, the embedded hardware used for inference must show appropriate characteristics such as numerical accuracy and performance, while being compliant with the related hardware standards.

In this paper, we focus on the inference part and the hardware used for it. This is a challenging task as the hardware used for training – mainly clusters of graphical processing units (GPUs) or tensor processing units (TPUs) – is generally vastly different from the one used for inference. The training

and inference hardware are usually at different optimum points when considering the different trade-offs that have to be made in terms of efficiency, flexibility, power, memory footprint, environmental conditions, and predictability. Additionally for the aeronautical industry, certification of the inference hardware is required, meaning that predictable execution time and mathematical robustness is mandatory. In practice, this means that a good trade-off between performance (e.g., inference rate), power consumption, complexity, and numerical accuracy must be found in addition to predictable execution time.

To meet these challenges, custom hardware based on field-programmable gate arrays (FPGAs) is a promising solution since it provides the flexibility to manage these trade-offs, including the execution predictability. Moreover, an FPGA design comes as a white box being a clear advantage for certification. We investigate in this paper some key aspects about mathematical robustness when using FPGAs for accelerating NNs, namely: the challenges of using IEEE 754 floating-point arithmetic and alternate number representation as a solution to those.

The main challenge of IEEE 754 floating-point arithmetic is that it is not associative, meaning that the order of operations can influence the final result of the computation. In an aeronautical certification context, this can lead to a challenge as hardware with a significantly higher mathematical performance is used for training compared to the inference device embedded into an airborne vehicle.

In this paper, we will investigate some methods to alleviate this issue by evaluating different summation and dot product algorithms for floating-point arithmetic. Moreover, we will investigate the use of fixed-point arithmetic, as this alternative number representation format does not suffer from the issue with the associativity of the operations. We will illustrate that it also enables non negligible gains in terms of hardware requirements, as operations on integers are simpler to implement than IEEE 754 floating-point arithmetic. We will also investigate how many bits are required to reach a sufficient accuracy by an empirical approach, as bit-width optimization is known to be NP-hard [9].

Finally, we will evaluate the impact of floating point and fixed point arithmetics in terms of hardware resources for two exemplary convolutional neural networks (CNNs) for image classification. Our goal is to find the optimum point in terms of numerical accuracy, bit-width, and hardware resources for

an exemplary FPGA platform.

This paper is organized as follows. First, we review the related work in Section II. Section III provides some background on the challenges of certification of ML in the aeronautical industry. We review existing methods for accurate floating-point in Section IV and fixed-point computations in Section V, and the hardware resources needed for them in Section VI. A numerical evaluation of our approach is shown in Section VII. Finally, Section VIII concludes the paper.

II. RELATED WORK

A recent survey on NNs approximations for custom hardware [10] shows that different methods have been proposed in the literature to speed-up the inference of NNs, such as alternative number representation, quantization, pruning or activation function approximation. Similarly, [11] recently surveyed tools for reduced precision computation, a growing trend for enhancing performance metrics in embedded systems and high performance computing (HPC). It highlights the lack of automated precision customization support in standard compiler frameworks and the ongoing research to improve automation, emphasizing the need for better tools, especially those based on static analysis.

Finding the appropriate number representation for NNs has already been extensively investigated in the literature, with the evaluation of bfloat16 [12, 13], posits [14, 15, 16], Microsoft floating point [17], FlexPoint [18], or adaptive floating-points [19].

In the scope of fixed-point operations, various works focused on integer-only inference and training for NNs. FxpNet is proposed in [20], a framework for training deep CNNs using low bit-width arithmetics in both forward and backward passes, adapting the bit-width of stored parameters during training. It employs integer batch normalization and fixed-point optimization methods to minimize floating-point operations, leading to power and chip area savings, with experimental results demonstrating comparable accuracy to state-of-the-art binarized and quantized approaches. [21, 22] recently introduced HAWQ and HAWQV3, an integer-only inference where the entire computational graph is performed only with integer operations. They address the hidden costs of current low-precision quantization algorithms, and present a novel mixed-precision integer-only quantization framework that enables integer-based computations and hardware-aware quantization.

Finding the optimal bitwidth given a mathematical formula has been a challenge since the age of digital signal processings (DSPs). [23] proposed an automated static method for optimizing bit widths of fixed-point feedforward designs, ensuring guaranteed accuracy. It employs semi-analytical precision analysis and adaptive simulated annealing to minimize both integer and fraction parts. In the scope of NNs, [24] proposed a novel technique involving linear programming and integer variables to optimize NNs precision without compromising output quality beyond a user-defined threshold. It is based on the method from [25], which combines forward and backward static analyses through abstract interpretation, expressed as a

set of constraints with first-order predicates and affine integer relations, simplifying verification by an SMT Solver. A similar approach was applied to code generation for NNs using error analysis in [26, 27]. While the methods from [24, 26, 27] were shown to be successful on small NNs, they scale poorly to larger networks like ResNet, as shown later in Section V-C.

III. BACKGROUND ON CERTIFICATION

Certification of hardware and software in the aeronautical industry is a rigorous process aiming at ensuring the safety, reliability, and compliance of aviation systems with stringent regulatory standards. It involves thorough testing, analysis, and documentation to verify that the onboard equipment and software meet strict requirements set by aviation authorities such as the European Union Aviation Safety Agency (EASA) or the Federal Aviation Administration (FAA). This certification process plays a critical role in guaranteeing the airworthiness of aircraft, promoting technological advancements, and maintaining the highest levels of safety for passengers and crew.

In addition to hardware and software certification, the aeronautical industry is increasingly focusing on the incorporation of onboard ML and its robustness in critical aviation systems. Ensuring the reliability and efficiency of ML algorithms and their hardware implementation is crucial for tasks such as predictive maintenance, autonomous decision-making, and enhanced flight operations. This necessitates a comprehensive evaluation of the ML models' performance under various operational scenarios, as well as a meticulous examination of the hardware used for inference.

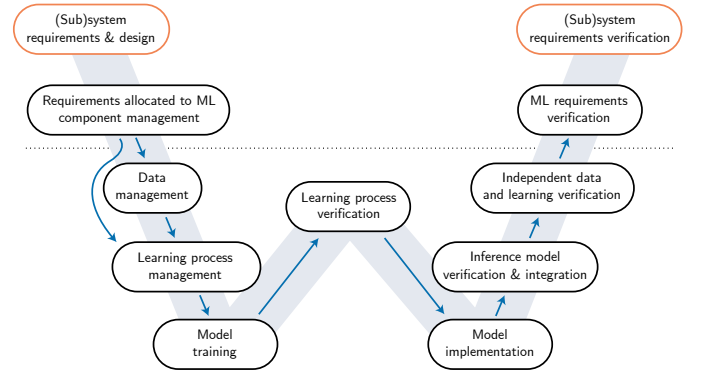


Figure 1: W-shaped development cycle for design assurance for NNs from [7]

To address this challenge, a W-shaped development process as illustrated in Figure 1 has been proposed for tailor the classical V-shaped cycle to ML applications [6, 7]. Various efforts have been started to standardize this development process in order to meet the high reliability and robustness requirements [1, 2, 3, 4].

In this paper we focus on the challenges for the development of hardware accelerators used for onboard NNs inference. These accelerators must exhibit sufficient performance to enable execution of large NNs at high enough execution rate by employing techniques such as pipelining, parallelization

and numerical approximations. However, the necessity for heightened efficiency clashes with the demand for precise numerical accuracy caused by limited computational hardware resources.

IV. FLOATING-POINT ARITHMETIC

Floating-point arithmetic is a method of representing real numbers in a way that allows a wide range of values to be expressed using a fixed number of bits. It has been standardized under IEEE 754 [28], a widely accepted standard that defines the format and rules for performing arithmetic operations with floating-point numbers. It is commonly found in off-the-shelves central processing units (CPUs), GPUs and TPUs.

Despite this standardization, issues related to the order of operations and precision limitations persist. Due to the finite precision of floating-point numbers, operations like addition, subtraction, multiplication, and division may not always yield exact results, leading to rounding errors and loss of precision, especially when performed in a different order than intended. This can potentially impact the accuracy of numerical computations, making it crucial to be aware of these limitations when working with floating-point numbers.

As parallelism is often used for accelerating computations, the order of operations is not guaranteed, leading to numerical inaccuracies. This holds significance not just within the realm of certification of NNs but also in other domains like scientific computing, where the reproducibility of results is considered crucial [29, 30]. In addition to several rounding methods, we evaluate different implementations of operations commonly found in NNs, namely summation and dot products.

A. Rounding

IEEE 754 defines several rounding modes, which determine how a floating-point number should be rounded to fit into a specific precision. The default rounding mode is *round to nearest even (RNE)*, where numbers are rounded to the nearest representable value. If the number falls exactly midway between two representable values, it chooses the one with an even least significant bit.

In this paper, we will also evaluate two additional rounding modes: *round towards zero (RTZ)* which always truncates the fractional part, effectively rounding towards zero regardless of the sign of the number; and *round to nearest away (RNA)*, where numbers are rounded to the nearest representable value, and if it is equidistant from both, it is rounded away from zero.

B. Summation algorithms

The order of operations is crucial when summing floating points to ensure accurate results and prevent rounding errors that could accumulate with each operation. Various methods have been proposed to address this, as showed in [31]. For this paper, we focus on four approaches, as they are easily implemented in hardware and do not require any sorting.

The *naive accumulation* summation algorithm, also known as the straightforward or simple summation method, involves

iteratively adding each element of a given set of numbers to an accumulator or running total.

The *pairwise summation* algorithm recursively divides the set of numbers into pairs, adding the pairs together, and then continuing the process until a single sum is obtained.

The *Kahan summation* algorithm [32] and its extensions [33], also known as compensated summation algorithms, keep track of the accumulated error during the summation process and compensating for it during each step of the process.

Finally, the *exact summation* algorithm corresponds to a fixed-point accumulator used for summing floating point introduced by [34]. This method is also detailed later in Section VI.

C. Dot product algorithms

Similar to the compensated summation algorithm, a compensated dot product algorithm – labeled *ORO* dot product in the text – has been proposed in [35]. The algorithm uses error-free transformations of the sum and product of two floating point numbers to perform accurate dot products.

V. FIXED-POINT ARITHMETIC

Quantizing the operations to 8 bits integers [36] is currently gaining more traction due to its efficiency on CPUs and hardware accelerators. Yet, this approach isn't without drawbacks: it requires a post-processing phase of the trained NN to scale the numbers, and the drop in performance can be significant in some cases as shown later in Section VII-B.

We investigate a similar approach using fixed-point arithmetic, a method widely used in DSP and gaming due to their speed compared to floating points. Our approach does not require scaling of the weights of the NN and is simpler from a computational of view than using floating points. Additionally, the order of the operations is not relevant here, compared to floating point.

A. Definition

Fixed-point arithmetic is a method of representing numbers by storing a fixed number of digits of their fractional part. Numbers are represented as integers which are split into three parts: a sign bit, a magnitude part with m bits and a fractional part with f bits. Conversion from a real number x to its fixed point representation is done via the following function:

$$\text{round}(x \cdot 2^f) \quad (1)$$

with *round* a rounding function as described in Section IV-A.

Mathematical operations can be easily performed using the underlying integers. To add or subtract two values of the same fixed-point type, it is sufficient to add or subtract the underlying integers.

To multiply two fixed-point numbers, it suffices to multiply the two underlying integers, giving a result with a fractional part of $2f$ bits. To avoid an increasing number of bits for the fractional part when performing multiple multiplications, rescaling is required. This is performed by shifting right the underlying integer and taking care of rounding.

As illustrated later in Section VII-D and Figure 8, correct rounding for the multiplication operation can dramatically improve the results.

B. Dot product algorithm

From the description of the previous section, the main loss of information appears during the multiplication of two fixed-point numbers, where a right shift and rounding operation is required. We reformulate the dot product of the \mathbf{x} and \mathbf{y} vectors as:

$$\mathbf{x} \cdot \mathbf{y} = \sum_i SR(x_i y_i) \quad (\text{naive dot product}) \quad (2)$$

$$= SR \left(\sum_i x_i y_i \right) \quad (\text{accurate dot product}) \quad (3)$$

with x_i and y_i the underlying integer at position i of the vectors \mathbf{x} and \mathbf{y} , and SR the right shift and round operation.

As only one rounding operation is required in Equation (3), it will produce more accurate results than Equation (2).

C. Mathematically bounding the error

As mentioned in Section II, the works from [26, 27] propose a method based on affine arithmetic in order to mathematically bound the absolute error of the outputs of an NN. It can easily be derived from Equation (1) that the magnitude of the maximum difference between a real value and its representation is 2^{-f} . Propagating this error through the operation of the NN is then performed using affine arithmetic.

While such method is applicable on relatively small NNs with only a few layers, it becomes impossible to use on deeper architectures. To illustrate this, we computed the bounds given by this method on a pretrained ResNet18 CNN [37] using fixed point arithmetics with different bit widths. The results are shown in Figure 2, where the error propagation of the CNN tends to get larger with increasing number of layers. Accordingly, a static analysis is not useful in practice and not further evaluated in this paper.




Figure 2: Evaluation of mathematical bounds and empirical values of the error of executing and converting the ResNet18 CNN [37] in fixed point arithmetic

VI. RESOURCE USAGE FOR ARITHMETIC

We evaluate here the resources required for different binary number formats in FPGAs. Due to their flexibility, FPGAs offer a wider range of formats than the commonly found ones in off-the-shelf CPUs or GPUs such as 8, 16, 32 or 64 bits integers or floating-points. Furthermore, the formats can be adjusted for the specific NNs that shall be executed depending on the precision that is needed for the given application. In order to perform the multiplication and adding operations needed, essential parts of any NNs processor are the multiplyaccumulate (MAC) units. These hardware blocks perform the multiplication of two values and accumulate the result of the multiplication. The number of MAC units which can be utilized in parallel at a certain clock frequency determines the maximum achievable performance of the device. The performance is typically given in operations per second (OPS) or floating-point operations per second (FLOPS). As one MAC operation consists of a multiplication and an accumulation, both computations are counted separately.

FPGAs offer different resources that can be configured and connected to implement the intended logic. The main resources, available in all different types of FPGAs, are lookup tables (LUTs) which are configured to implement the combinatorial logic and flip-flops (FFs). In addition, DSPs can be integrated in the design to speed up specific operations commonly needed for example in filters, fast-Fourier-transforms or other suitable algorithms. These DSPs differ depending on the FPGA architecture and vendor. FPGAs are available in different sizes, different numbers of LUTs, FFs, DSPs and the relation between the elements needed can be selected. There are FPGAs with a high number of LUTs compared to available DSPs and vice versa in order to select the right FPGA for the task as certain designs might be able to leverage DSPs while this is not possible in a different design. For our evaluations we selected the AMD VU9P, based on the Virtex Ultrascale Plus technology, because of the balanced relation between LUTs, FFs and DSPs [38].

In this paper, we analyze the performance achievable on FPGAs for fixed-point as well as floating-point numbers of arbitrary size of the fractional part / mantissa. For fixed-point, a MAC unit consisting of the integer Multiplier and Adder/Subtractor AMD IP cores of Vivado 2023.1 was synthesized and implemented for different numbers of fractional bits while using 10 bits for the integer part, which has shown to be the necessary bit width to achieve the desired accuracy on ResNet18 without the need to scale the numbers. Additionally, variants were implemented for different rounding modes and whether the rounding is done after the accumulation (accurate dot product) or after every multiplication (naive dot product). The results of this analysis are the resource usages (number of LUTs, FFs, DSPs) and the maximum frequencies for a single MAC unit on the aforementioned FPGA. Furthermore, the results include resource usage for fabric only (without DSPs) for a better transferability to other FPGA architectures, and with DSP usage to achieve the maximum performance

on the given FPGA. In order to transfer the resource usage and frequencies to performance estimates, it was theoretically analyzed how many of these MAC units fit into the exemplary FPGA. It was assumed that 70 % of the LUTs and FFs can be utilized to avoid potential timing closure issues and routing congestion while 100 % of the DSPs can be used. The number of MAC units fitting in the FPGA is multiplied by the achievable frequency which defines the upper bound of the performance of the complete chip and is shown in Figure 3.




Figure 3: Performance for fixed-point MACs on AMD VU9P FPGA with and without DSP usage

In the figure it is visible that in the case of no DSPs usage, for an increased number of bits the performance exponentially decreases while there is only a slight difference when comparing the different rounding and accumulation styles. However, with the usage of DSPs, the performance decreases in steps. In this case the DSPs are the limiting factor for the full utilization of the FPGA and the number of DSPs used by the MAC units of different bit sizes is increasing stepwise i.e. for fractional bit sizes of 9 to 17 the same number of DSPs per MAC unit is needed. Furthermore, there is no difference for the different rounding styles because the rounding only adds LUTs which are still available on the FPGA. Thus, even though the amount of resources consumed are slightly different the resulting performance is identical.

For floating-point, a MAC unit was designed with the AMD floating-point IP core consisting of a multiplier and an adder implemented and synthesized for different bit sizes of the mantissa. However, using a floating-point adder as an accumulator, later referred to as *naive* method, leads to a low frequency of the FPGA as the addition has to happen within one cycle. A second MAC unit was designed with the Floating-point IP core accumulator instead of the adder, later referred to as *exact* method. This accumulator is implemented as a fixed-point accumulator internally which leads to a very high precision but also a very high resource consumption [34]. However, since it can be pipelined, a high frequency is achievable. The theoretically achievable maximum performance on the given FPGA are shown in Figure 4. As expected, the performance decreases for a higher number of bits for the mantissa. Contrary to the fixed-point analysis, if DSPs are used the bit size impacts the performance not stepwise because the

LUTs are the dominant resource for the floating-point MAC units.

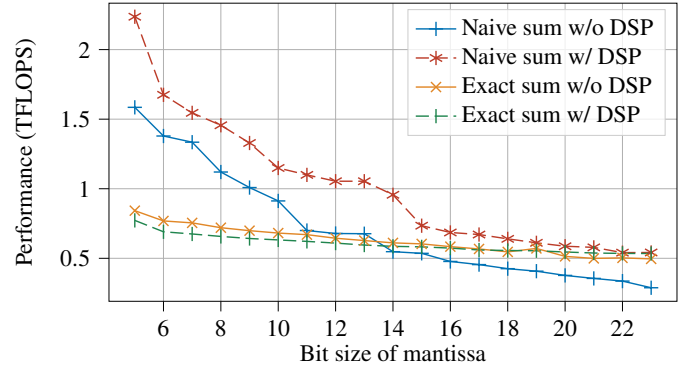


Figure 4: Performance for floating-point MACs on AMD VU9P FPGA with and without DSP usage

VII. NUMERICAL EVALUATION

We numerically evaluate in this section the different approaches presented earlier and assess their impacts in terms of numerical accuracy and hardware resources.

A. Methodology

To evaluate the impact of the bitwidth and the different summation and dot product algorithms previously listed, we implemented our own Open Neural Network Exchange (ONNX) runtime using Go and C. Floating point arithmetic are implemented using Go's `math/big.Float` arbitrary-precision arithmetic library. We use our own implementation for fixed point arithmetic. This enables us to precisely target the mathematical operations and number representation under investigation while being compatible with existing ONNX toolchains.

As exemplary models for our numerical evaluation, we use the pretrained models from the ONNX model zoo¹ for MNIST [39] and ResNet18 [37]. For the numerical evaluations, the full test set of MNIST is used for the MNIST model, and a subset of ImageNet-1k dataset [40] is used for ResNet18.

To evaluate the accuracy of our computations, we use the ONNX runtime from Microsoft² with its CPU execution provider as reference. Our main metric is to evaluate if the top-1 classification from our model with lower precision is the same as the top-1 classification from the reference (32 bits floating-point) model. For our evaluation and use-case, we aim at achieving a metric of 100 %, i.e. match the classification from the reference.

B. Evaluation of int8 quantized models

As a first benchmark, we evaluate the performance of int8 quantized models. For the MNIST model, we use the already publicly available quantized version of the model from the ONNX model zoo. For ResNet18, we use the Intel Neural Compressor open-source tool [41] to quantize the model.

¹<https://github.com/onnx/models>

²<https://github.com/microsoft/onnxruntime>

Table I: Metric for the int8 quantized models

Model	Same top-1 as reference
ResNet18	85.80 %
MNIST	54.69 %

Results are presented in Table I. It is clear from the values of our evaluation metric that the int8 quantized models are not sufficient since the required 100 % metric is not reached. These results justify why better precision and an evaluation of alternate approaches for accelerating inference are required for our use-cases.

C. Floating-point arithmetic

We evaluate in this section the performance with floating point arithmetic and the methods described in Section IV.

1) *Impact of the summation function:* Results illustrating the impact of the summation function (with naive dot product) on ResNet18 are presented in Figure 5. There is a clear benefit at using more accurate summation functions, as it dramatically improves the accuracy of the results over the naive summation for low bit widths.

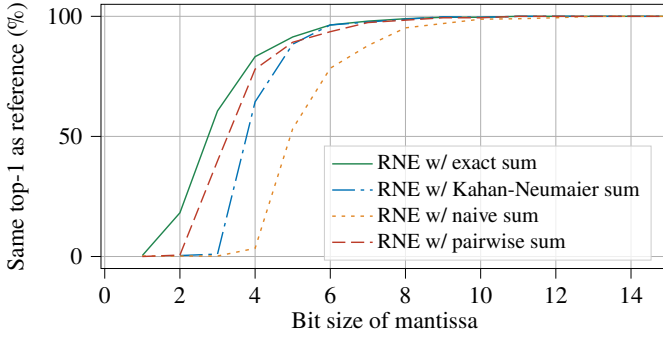


Figure 5: Impact of the summation algorithm on ResNet18 with floating-point arithmetic

Overall, the exact summation provides the best results. Yet, 11 bit for the mantissa are required for all three non-naive summation functions to achieve a metric of 100 %.

2) *Impact of the dot product function:* Results illustrating the impact of the dot product function on MNIST are presented in Figure 6. While the ORO dot product [35] enables us to gain a few percent on our metric, its impact is minimal: for both dot products, the same number of bits are required to reach 100 %. For ResNet18, the same conclusion can be made.

Overall, these results illustrate that adding the overhead of this dot product function in hardware is not worth it, as there is no gain in terms of bits required for computations.

D. Fixed-point arithmetic

We evaluate in this section the performance with fixed point arithmetic and the methods described in Section V.

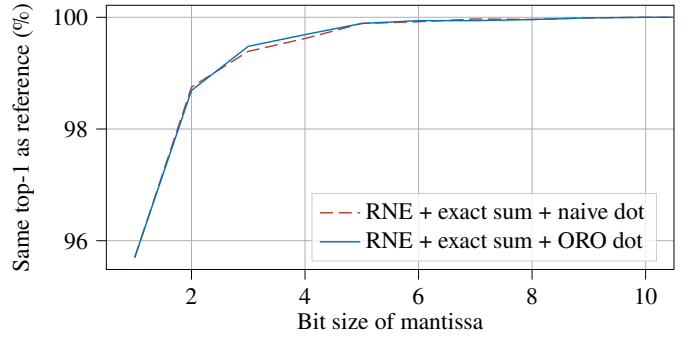


Figure 6: Impact of the dot product algorithm on MNIST with floating-point arithmetic

1) *Impact of rounding:* The impact of the rounding for the MNIST model is presented in Figure 7. Correctly rounding with RNE or RNA during the multiplication dramatically improves the results for low bit widths compared to RTZ, saving 3 bit on average.

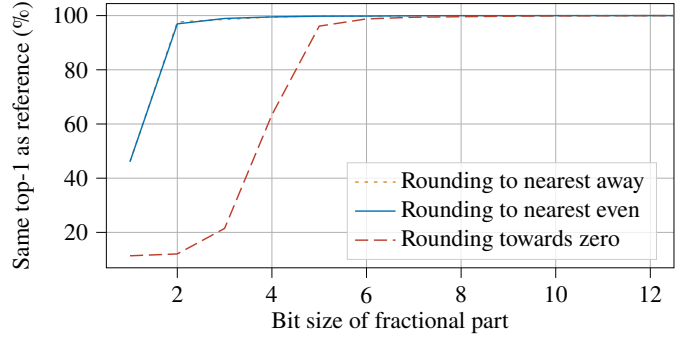


Figure 7: Impact of rounding mode on MNIST with fixed-point arithmetic with naive dot product. The RNE and RNA curves overlap.

2) *Impact of the dot product function:* The impact of the accurate dot product from Equation (3) on ResNet18 is presented in Figure 8. Unsurprisingly, better results are achieved using Equation (3) compared to the naive dot product.

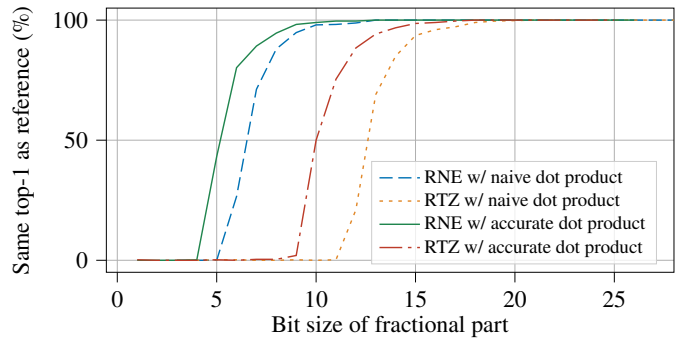


Figure 8: Impact of the rounding mode and dot product functions on ResNet18 with fixed-point arithmetic

E. Summary and hardware impact

Based on the previous benchmarks and hardware resources presented in Section VI, we summarize here our results and define the optimum points in terms of number representation, algorithms and hardware resources required.

Figure 9 presents the hardware performance which can be achieved given the different computing parameters evaluated. From these results, it is clear that using fixed point arithmetic provides the best performance.

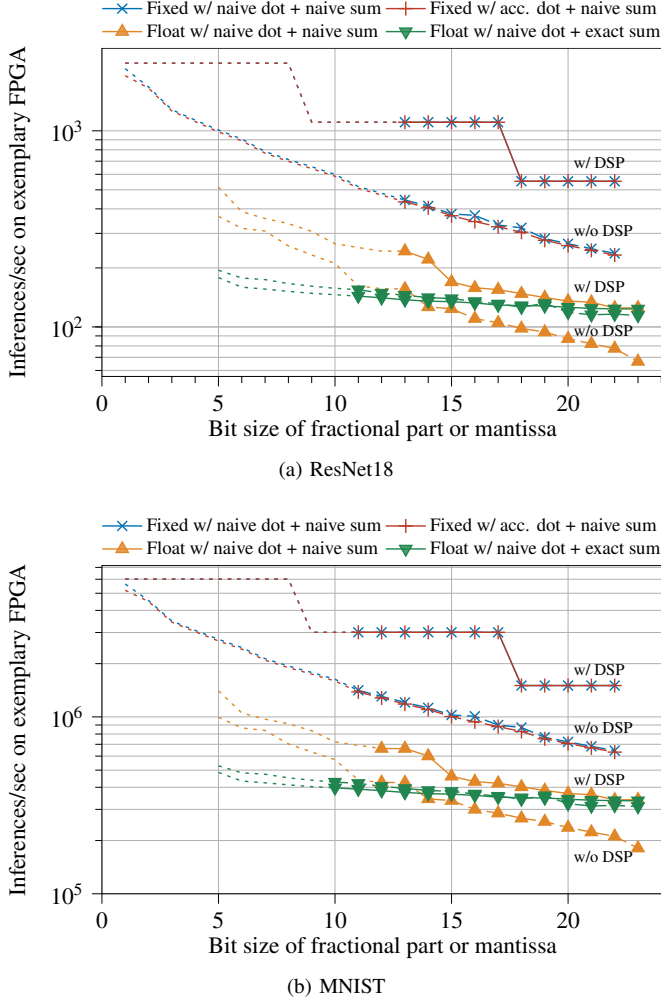


Figure 9: Hardware performance which can be achieved depending on the number representation, dot product and summation algorithms, with RNE. Dashed lines on the left side of the plots represent combinations where the accuracy of the computations is not sufficient to reach a metric of 100 %.

Tables II and III represent the optimum points for each combination of parameters. The "PBits" column represents the minimum number of bits required for either the mantissa for floating point arithmetic, or the fractional part for fixed point. The provided numbers in the "Estimated inferences/s" columns are inferred by dividing the available TOPS/TFLOPS at the given size of the fractional part or mantissa by the total

MAC operations needed for a single inference on MNIST or ResNet18, respectively.

The dot product according to ORO [35] was not implemented in hardware and marked as "noORO" in the table. Our hardware analysis is based on the AMD IP cores which implement only RNE, hence, we did not analyse RTZ and RNA but we expect very similar results compared to RNE, see "RNEonly". An estimation was done on the implementation of the KN and Pairwise algorithms in Verilog. As these algorithms are more hardware consuming than the implementation of the "Exact" sum, we did not further analyze them.

Table II: Summary of the results for MNIST

	Dot Prod.	Sum	Rnd	PBits	Estimated inferences/s	
					w/o DSP	w/ DSP
Fixed point	Accurate	Naive	RNA	11	1 451 925	3 007 621
	Accurate	Naive	RTZ	11	1 437 854	3 007 621
	Naive	Naive	RNA	11	1 426 421	3 007 621
	Naive	Naive	RNE	11	1 413 230	3 007 621
	Accurate	Naive	RNE	11	1 383 330	3 007 621
	Naive	Naive	RTZ	12	1 352 550	3 007 621
Floating point	Naive	Naive	RNE	12	426 014	662 086
	Naive	Exact	RNE	10	428 019	397 050
	ORO [35]	Exact	RNE	10	noORO	noORO
	ORO [35]	KN [33]	RNE	10	noORO	noORO
	ORO [35]	KN [33]	RNA	10	noORO	noORO
	ORO [35]	KN [33]	RTZ	11	noORO	noORO
	Naive	Exact	RNA	10	RNEonly	RNEonly
	Naive	Exact	RTZ	11	RNEonly	RNEonly
	Naive	KN [33]	RNE	10	noKN	noKN
	Naive	Pairwise	RNE	10	noPairwise	noPairwise
	Naive	Pairwise	RNA	11	noPairwise	noPairwise
	Naive	Pairwise	RTZ	11	noPairwise	noPairwise
	Naive	Naive	RNA	11	RNEonly	RNEonly
	Naive	Naive	RTZ	13	RNEonly	RNEonly

Table III: Summary of the results for ResNet18

	Dot Prod.	Sum	Rnd	PBits	Estimated inferences/s	
					w/o DSP	w/ DSP
Fixed point	Accurate	Naive	RNA	13	444	1106
	Naive	Naive	RNE	13	444	1106
	Accurate	Naive	RNE	13	434	1106
	Accurate	Naive	RTZ	18	313	553
	Naive	Naive	RTZ	21	258	553
Floating p.	Naive	Naive	RNE	13	156	243
	Naive	Exact	RNE	11	155	144
	Naive	KN [33]	RNE	11	noKN	noKN
	Naive	Pairwise	RNE	11	noPairwise	noPairwise

VIII. CONCLUSION

We reviewed in this paper various methods for achieving efficient and mathematically robust inference for neural networks (NNs) in the context of certification of hardware and software for machine learning (ML) for aeronautical applications. This is a challenging task, as special care is required on the mathematical operations to sufficiently accelerate a model in hardware while still preserving the same predictions as the

model originally trained on graphical processing units (GPUs) or tensor processing units (TPUs).

From a mathematical perspective, we numerically evaluated the various choices which are available, namely: use of floating vs. fixed point arithmetic, reduced precision arithmetic, and more accurate summation and dot product. From a hardware perspective, we assessed the impact of those choices on the resources required for an exemplary field-programmable gate array (FPGA). Overall, this enabled us to find the good balance in terms of hardware performance and mathematical precision.

Our results show that fixed-point arithmetic with sufficient bits for the fractional part yields the target accuracy for the NN and achieves the best performance.

REFERENCES

- [1] "Artificial Intelligence in Aeronautical Systems: Statement of Concerns," SAE International, Standard SAE ARP AIR6988, Apr. 2021.
- [2] "Artificial Intelligence in Aeronautical Safety-Related Systems Statement of concerns," European Organisation for Civil Aviation Equipment, Standard EUROCAE ER-022, May 2021.
- [3] "Artificial Intelligence Roadmap 2.0," European Union Aviation Safety Agency, Whitepaper, May 2023.
- [4] "Autonomy Verification & Validation Roadmap and Vision 2045," National Aeronautics and Space Administration, Technical Memorandum 20230003734, Jan. 2023.
- [5] J. Perez-Cerrolaza, J. Abella, M. Borg, C. Donzella, J. Cerquides, F. J. Cazorla, C. Englund, M. Tauber, G. Nikolakopoulos, and J. L. Flores, "Artificial intelligence for safety-critical systems in industrial and transportation domains: A survey," *ACM Comput. Surv.*, Oct. 2023.
- [6] EASA and Daedalean, "Concepts of design assurance for neural networks (CoDANN)," Tech. Rep., Mar. 2020.
- [7] —, "Concepts of design assurance for neural networks (CoDANN) ii," Tech. Rep., May 2021.
- [8] "Process Standard for Development and Certification/Approval of Aeronautical Safety-Related Products Implementing AI," SAE International, Work-in-progress standard SAE ARP 6983, Jun. 2023.
- [9] G. Constantinides and G. Woeginger, "The complexity of multiple wordlength assignment," *Applied Mathematics Letters*, vol. 15, no. 2, pp. 137–140, 2002.
- [10] E. Wang, J. J. Davis, R. Zhao, H.-C. Ng, X. Niu, W. Luk, P. Y. K. Cheung, and G. A. Constantinides, "Deep neural network approximation for custom hardware: Where we've been, where we're going," *ACM Comput. Surv.*, vol. 52, no. 2, May 2019.
- [11] S. Cherubin and G. Agosta, "Tools for reduced precision computation: A survey," *ACM Comput. Surv.*, vol. 53, no. 2, Apr. 2020.
- [12] D. Kalamkar, D. Mudigere, N. Mellempudi, D. Das, K. Banerjee, S. Avancha, D. T. Vooturi, N. Jammalamadaka, J. Huang, H. Yuen, J. Yang, J. Park, A. Heinecke, E. Georganas, S. Srinivasan, A. Kundu, M. Smelyanskiy, B. Kaul, and P. Dubey, "A study of BFLOAT16 for deep learning training," 2019.
- [13] N. Burgess, J. Milanovic, N. Stephens, K. Monachopoulos, and D. Mansell, "Bfloat16 processing for neural networks," in *2019 IEEE 26th Symposium on Computer Arithmetic (ARITH)*, 2019, pp. 88–91.
- [14] Z. Carmichael, H. F. Langroudi, C. Khazanov, J. Lillie, J. L. Gustafson, and D. Kudithipudi, "Deep Positron: A deep neural network using the posit number system," in *2019 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2019, pp. 1421–1426.
- [15] M. Cococcioni, F. Rossi, E. Ruffaldi, S. Saponara, and B. Dupont de Dinechin, "Novel arithmetics in deep neural networks signal processing for autonomous driving: Challenges and opportunities," *IEEE Signal Processing Magazine*, vol. 38, no. 1, pp. 97–110, 2021.
- [16] J. Lu, C. Fang, M. Xu, J. Lin, and Z. Wang, "Evaluations on deep neural networks training using posit number system," *IEEE Transactions on Computers*, vol. 70, no. 2, pp. 174–187, 2021.
- [17] B. Rouhani, D. Lo, R. Zhao, M. Liu, J. Fowers, K. Ovtcharov, A. Vinogradsky, S. Massengill, L. Yang, R. Bittner, A. Forin, H. Zhu, T. Na, P. Patel, S. Che, L. C. Koppaka, X. Song, S. Som, K. Das, S. Tiwary, S. Reinhardt, S. Lanka, E. Chung, and D. Burger, "Pushing the limits of narrow precision inferencing at cloud scale with microsoft floating point," in *Proceedings of the 34th International Conference on Neural Information Processing Systems*, ser. NIPS'20, 2020.
- [18] U. Köster, T. J. Webb, X. Wang, M. Nassar, A. K. Bansal, W. H. Constable, O. H. Elibol, S. Gray, S. Hall, L. Hornof, A. Khosrowshahi, C. Kloss, R. J. Pai, and N. Rao, "Flexpoint: An adaptive numerical format for efficient training of deep neural networks," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, ser. NIPS'17, 2017, pp. 1740–1750.
- [19] F. Liu, W. Zhao, Z. He, Y. Wang, Z. Wang, C. Dai, X. Liang, and L. Jiang, "Improving neural network efficiency via post-training quantization with adaptive floating-point," in *Proc. of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [20] X. Chen, X. Hu, H. Zhou, and N. Xu, "FxpNet: Training a deep convolutional neural network in fixed-point representation," in *2017 International Joint Conference on Neural Networks (IJCNN)*, 2017, pp. 2494–2501.
- [21] Z. Dong, Z. Yao, A. Gholami, M. Mahoney, and K. Keutzer, "HAWQ: Hessian aware quantization of neural networks with mixed-precision," in *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.
- [22] Z. Yao, Z. Dong, Z. Zheng, A. Gholami, J. Yu, E. Tan, L. Wang, Q. Huang, Y. Wang, M. Mahoney, and K. Keutzer, "Hawq-v3: Dyadic neural network quantization," in *Proceedings of the 38th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 139, 18–24 Jul 2021, pp. 11 875–11 886.
- [23] D.-U. Lee, A. Gaffar, R. Cheung, O. Mencer, W. Luk, and G. Constantinides, "Accuracy-guaranteed bit-width optimization," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 25, no. 10, pp. 1990–2000, 2006.
- [24] A. Ioualalen and M. Martel, "Neural network precision tuning," in *Quantitative Evaluation of Systems*, 2019, pp. 129–143.
- [25] M. Martel, "Floating-point format inference in mixed-precision," in *Lecture Notes in Computer Science*, 2017, pp. 230–246.
- [26] H. Benmagnia, M. Martel, and Y. Seladj, "Fixed-point code synthesis for neural networks," in *Artificial Intelligence, Soft Computing and Applications*, Jan. 2022.
- [27] —, "Code generation for neural networks based on fixed-point arithmetic," *ACM Trans. Embed. Comput. Syst.*, Sep. 2022.
- [28] "IEEE Standard for Floating-Point Arithmetic," *IEEE Std 754-2019 (Revision of IEEE 754-2008)*, pp. 1–84, 2019.
- [29] S. F. Jalal Apostol, D. Apostol, and R. Marsh, "Improving numerical reproducibility of scientific software in parallel systems," in *IEEE International Conference on Electro Information Technology*, 2020.
- [30] R. W. Robey, J. M. Robey, and R. Aulwes, "In search of numerical consistency in parallel programming," *Parallel Computing*, vol. 37, no. 4, pp. 217–229, 2011.
- [31] N. J. Higham, "The accuracy of floating point summation," *SIAM Journal on Scientific Computing*, vol. 14, no. 4, pp. 783–799, 1993.
- [32] W. Kahan, "Further remarks on reducing truncation errors," *Commun. ACM*, vol. 8, no. 1, p. 40, Jan 1965.
- [33] A. Neumaier, "Rundungsfehleranalyse einiger Verfahren zur Summation endlicher Summen," *ZAMM - Journal of Applied Mathematics and Mechanics / Zeitschrift für Angewandte Mathematik und Mechanik*, vol. 54, no. 1, pp. 39–51, 1974.
- [34] F. de Dinechin, B. Pasca, O. Cret, and R. Tudoran, "An FPGA-specific approach to floating-point accumulation and sum-of-products," in *International Conference on Field-Programmable Technology*, 2008.
- [35] T. Ogita, S. M. Rump, and S. Oishi, "Accurate sum and dot product," *SIAM Journal on Scientific Computing*, vol. 26, no. 6, 2005.
- [36] B. Jacob, S. Kligys, B. Chen, M. Zhu, M. Tang, A. Howard, H. Adam, and D. Kalenichenko, "Quantization and training of neural networks for efficient integer-arithmetic-only inference," in *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [37] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [38] AMD, "UltraScale+ FPGAs Product Selection Guide (XMP103)," 2023.
- [39] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [40] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2009, pp. 248–255.
- [41] F. Tian, H. Chang, H. Shen, and S. Chen, "Intel neural compressor," <https://github.com/intel/neural-compressor>, 2022.