

Tina: Acceleration of Non-NN Signal Processing Algorithms Using NN Accelerators

Boerkamp, Christiaan ; van der Vlugt, Steven; Al-Ars, Zaid

DOI

[10.1109/MLSP58920.2024.10734727](https://doi.org/10.1109/MLSP58920.2024.10734727)

Publication date

2024

Document Version

Final published version

Published in

Proceedings of the 2024 IEEE 34th International Workshop on Machine Learning for Signal Processing (MLSP)

Citation (APA)

Boerkamp, C., van der Vlugt, S., & Al-Ars, Z. (2024). Tina: Acceleration of Non-NN Signal Processing Algorithms Using NN Accelerators. In *Proceedings of the 2024 IEEE 34th International Workshop on Machine Learning for Signal Processing (MLSP)* IEEE. <https://doi.org/10.1109/MLSP58920.2024.10734727>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

TINA: ACCELERATION OF NON-NN SIGNAL PROCESSING ALGORITHMS USING NN ACCELERATORS

Christiaan Boerkamp* Steven van der Vlugt[†] Zaid Al-Ars*

*Delft University of Technology, Mekelweg 4, 2628 CD, Delft, Netherlands

[†]ASTRON, Oude Hoogeveensedijk 4, 7991 PD Dwingeloo, Netherlands

ABSTRACT

This paper introduces TINA, a novel framework for implementing non Neural Network (NN) signal processing algorithms on NN accelerators such as GPUs, TPUs or FPGAs. The key to this approach is the concept of mapping mathematical and logic functions as a series of convolutional and fully connected layers. By mapping functions into such a small sub stack of NN layers, it becomes possible to execute non-NN algorithms on NN hardware (HW) accelerators efficiently, as well as to ensure the portability of TINA implementations to any platform that supports such NN accelerators. Results show that TINA is highly competitive vs alternative frameworks, specifically for complex functions with iterations. For a Polyphase Filter Bank use case TINA shows GPU speedups of up to 80x vs a CPU baseline with NumPy compared to 8x speedup achieved by alternative frameworks. The framework is open source and publicly available at <https://github.com/ChristiaanBoe/TINA>.

Index Terms—Non-NN algorithms, signal processing algorithms, neural networks, HW accelerators

1. INTRODUCTION

In recent years, there has been a rapid increase in the number of specialized NN accelerators, either as standalone HW (e.g., TPUs) or as accelerated components in existing HW (e.g., GPUs or FPGAs). This creates a lucrative opportunity for using such accelerators to boost the performance of non-NN algorithms, in the same way that GPUs are used to accelerate non-graphics algorithms.

Existing research performed on this topic is limited. One example is JAX, a machine learning framework for transforming mathematical functions and mapping them on NN HW [1]. However, JAX is mainly focused on making ML programming more intuitive, rather than providing a general API to run non-NN algorithms on NN HW. Other efforts are either limited to a specific application (e.g., brain simulation [2]), or limited to a specific HW platform (e.g., GPUs [3] or FPGAs [4]).

Table 1. Signal processing functions implemented in this paper and the TINA building blocks used to implement them

Function	Building blocks	Section
Elementwise matrix mult.	Depthwise conv.	Section 3.1
Matrix-matrix mult.	Depthwise conv.	Section 3.2
Elementwise matrix add	Depthwise conv.	Section 3.3
Summation	Fully conn. layer	Section 3.4
DFT	Pointwise conv.	Section 4.1
Inverse DFT	Pointwise conv.	Section 4.2
FIR filter	Standard conv.	Section 4.3
Unfolding algorithm	Standard conv.	Section 4.4

This paper introduces a novel framework, referred to as TINA¹, to reduce the complexity of implementing non-NN applications on NN accelerators, thereby ushering in the era of GPNPUs (general-purpose neural processing units). In this paper, we introduce the idea of representing signal processing functions as a series of convolutions and fully connected layers. TINA makes use of the strengths of NN HW designed to accelerate NN layers, such as matrix operations. By reinterpreting traditional mathematical and logic functions as NN layers, we facilitate a form of computation that is intrinsically compatible with NN accelerators. As its input programming language, TINA uses Python, one of the most widely used programming languages in data science [5].

TINA provides a number of unique advantages, such as ensuring the portability of TINA implementations to any HW platform that supports NN accelerators. Furthermore, TINA allows non-NN algorithms to be optimized using methods previously exclusive to NNs (e.g., pruning libraries or automatically optimized quantization using quantize aware training libraries such as Brevitas [6]).

TINA consists of two main components: 1. basic building blocks, and 2. APIs that use the building blocks to implement various functions. Table 1 provides an overview of the functions discussed in this paper and the building blocks used to implement them.

¹TINA stands for "This Is Not AI"

2. BUILDING BLOCKS OF TINA

The building blocks of TINA can be divided into four NN layers: standard convolution, depthwise convolution, pointwise convolution and fully connected layers. These building blocks will be used later in the paper to implement signal processing functions. We use these specific four NN layers since they are the most widely used layers in ML applications [7, 8], thereby making them well-supported by accelerators, and easily portable to alternative hardware platforms.

2.1. Standard convolution

Convolutions are critical in deep learning for feature extraction from various types of input data, such as images and audio. This process involves sliding a kernel across the input data, conducting elementwise multiplication with the input values covered by the kernel, and summing these products to generate output feature maps, as depicted in Eq. (1).

$$O(h, w, c_{out}) = b(c_{out}) + \sum_{c_{in}} \sum_m \sum_n I(h+m, w+n, c_{in}) \cdot K(m, n, c_{in}, c_{out}) \quad (1)$$

where O is the output, I is the input, K is the kernel, c_{in} and c_{out} are indices representing the input and output channels, respectively, while b is the bias. TINA leverages PyTorch to enable precise customization of convolutional layers. This customization includes controlling the dimensions of both the input matrix, represented as (T, C_{in}, H, W) , and the output matrix, represented as (T, C_{out}, H, W) , where T is the batch size, C denotes the number of channels, and H and W are the height and width of the input/output, respectively.

PyTorch provides a variety of parameters to allow for extra controls over the convolution. For example, we can use **groups** to define blocks of connections between input and output channels. Other notable parameters include **stride**, determining the kernel movement steps, and **padding**, adjusting the input border size.

2.2. Depthwise convolution

Depthwise convolution represents a mechanism within convolutional neural networks for dissecting spatial relationships in input data. Distinct from standard convolutional layers, which convolve across both input channels and spatial dimensions, depthwise convolution applies the corresponding channel of the kernel to each input channel independently, resulting in an output with the same number of channels as the input. This operation is represented in Eq. (2).

$$O(h, w, c_{out}) = b(c_{out}) + \sum_m \sum_n I(h+m, w+n, c_{out}) \cdot K(m, n, c_{out}) \quad (2)$$

Depthwise convolution causes the kernel to slide across the input data, executing elementwise multiplications between the kernel and input values, with the sums producing the output. This procedure is independently executed for each channel, thereby generating distinct feature maps for every channel.

2.3. Pointwise convolution

Pointwise convolutions, also known as 1×1 convolutions, are a special case of standard convolutions where a kernel of dimensions 1×1 is used to process individual elements across the channels of input data, which allows for the mixing of channel information. This operation is represented in Eq. (3).

$$O(h, w, c_{out}) = b(c_{out}) + \sum_{c_{in}} I(h, w, c_{in}) \cdot K(c_{in}, c_{out}) \quad (3)$$

2.4. Fully connected layer

A fully connected layer, also known as a dense layer or linear layer, is a commonly-used layer in many neural network architectures. Its core role is to transform input data by applying a linear operation, typically followed by a non-linear activation function to introduce non-linearity into the model. The operation performed by a fully connected layer involves computing the output from the input through matrix multiplication with the kernel matrix, and adding a bias term b , as shown in Eq. (4).

$$O(c_{out}) = b(c_{out}) + \sum_{c_{in}} I(c_{in}) \cdot K(c_{in}, c_{out}) \quad (4)$$

3. ARITHMETIC FUNCTIONS MAPPING

In this section, we show how we can compute various arithmetic functions using the TINA building blocks discussed in Section 2.

3.1. Elementwise multiplication

An elementwise matrix multiplication can be implemented using a depthwise convolution. As shown in Eq. (5), a depthwise convolution applies the corresponding channel of the kernel to each input channel independently, causing the kernel to slide across the input matrix, executing elementwise multiplications between the kernel and input values, with the sums producing the output.

$$O(h, w, c_{out}) = b(c_{out}) + \sum_m \sum_n I(h+m, w+n, c_{out}) \cdot K(m, n, c_{out}) \quad (5)$$

In order to have this equation represent an elementwise matrix multiplication, we first set the bias $b(c_{out}) = 0$. In

addition, we convert the width and height of both the input matrix and the kernel into 1×1 matrices and reshape the elements of these two matrices into vectors along the channel axis where $C_{out} = H \times W$, resulting in Eq. (6).

$$\begin{aligned} O(0, 0, c_{out}) &= \sum_{m=0}^0 \sum_{n=0}^0 I(0+m, 0+n, c_{out}) \cdot K(m, n, c_{out}) \\ &= I(0, 0, c_{out}) \cdot K(0, 0, c_{out}) \\ \Rightarrow O(c_{out}) &= I(c_{out}) \cdot K(c_{out}) \end{aligned} \quad (6)$$

3.2. Matrix-matrix multiplication

A matrix-matrix multiplication is a common operation in linear algebra and can be seen as a combination of a summation and a pointwise multiplication as shown in Eq. (7), where $X(m, l)$ represents the first input matrix X with dimensions M and L , $Y(l, n)$ represents the second input matrix Y with dimensions L and N , and $Z(m, n)$ represents the output matrix Z with dimensions M and N .

$$Z(m, n) = \sum_l^L X(m, l) \cdot Y(l, n) \quad (7)$$

We can represent the matrix-matrix multiplication using a pointwise convolution by taking the following additional steps. We start with setting the bias of the pointwise convolution $b(c_{out})$ to zero resulting in Eq. (8).

$$O(h, w, c_{out}) = 0 + \sum_{c_{in}}^{C_{in}} I(h, w, c_{in}) \cdot K(c_{in}, c_{out}) \quad (8)$$

In order for Eq. (8) to represent a matrix-matrix multiplication, we apply the following transformation as shown in Eq. (9), by reshaping our input and output matrices with dimensions H and W into vectors with dimension $1 \times M$ where $M = H \times W$, and making $C_{in} = L$ and $C_{out} = N$. This in turn is equal to our original definition of a matrix-matrix multiplication in Eq. (7).

$$\begin{aligned} O(0, m, n) &= \sum_l^L I(0, m, l) \cdot K(l, n) \\ \Rightarrow O(m, n) &= \sum_l^L I(m, l) \cdot K(l, n) \end{aligned} \quad (9)$$

3.3. Elementwise addition

Elementwise addition builds further upon the elementwise multiplication introduced in Section 3.1. By opting to set the weights of the kernel to an all-ones matrix and the bias $b(c_{out})$ of the convolution as one of the input matrices, we are able to transform Eq. (6) to an elementwise addition as shown in Eq. (10). Note that I , bias and O in the equation

represent vectors of size $C_{out} = H \times W$.

$$\begin{aligned} O(c_{out}) &= b(c_{out}) + I(c_{out}) \cdot K(c_{out}) \\ &= b(c_{out}) + I(c_{out}) \cdot 1(c_{out}) \\ \Rightarrow O(c_{out}) &= b(c_{out}) + I(c_{out}) \end{aligned} \quad (10)$$

3.4. Summation

The implementation of summation within TINA using the building blocks described in Section 2 can be achieved through the utilization of a fully connected layer as described in Section 2.4. By opting to set the number of output channels to 1, the weights of the kernel to an all-ones vector and the bias to a zero vector, Eq. (4) can be transformed to a summation as shows in Eq. (11).

$$\begin{aligned} O(0) &= 0 + \sum_{c_{in}}^{C_{in}} I(c_{in}) \cdot 1(c_{in}, 0) \\ \Rightarrow O &= \sum_{c_{in}}^{C_{in}} I(c_{in}) \end{aligned} \quad (11)$$

4. SIGNAL PROCESSING FUNCTIONS MAPPING

4.1. DFT

In this section, we show how to use TINA building blocks to implement a Discrete Fourier Transform (DFT). One way an input signal can be transformed towards the frequency domain is by matrix-matrix multiplying the input signal vector with a Discrete Fourier Matrix (DFM) [9] as seen in Eq. (12), with X being the input signal and F being the Fourier matrix.

$$Z(m, n) = \sum_l^L X(m, l) \cdot F(l, n) \quad (12)$$

Using our matrix-matrix multiplication Eq. (13) as further detailed in Section 3.2 we can draw a direct parallel to Eq. (12) by setting the kernel equal to the Fourier matrix.

$$O(m, n) = \sum_l^L I(m, l) \cdot K(l, n) \quad (13)$$

4.2. IDFT

In this section, we show how to use TINA building blocks to implement an Inverse Discrete Fourier Transform (IDFT). One way an input Fourier signal can be transformed back towards the time domain is by matrix-matrix multiplying the input signal vector with an Inverse Discrete Fourier Matrix (IDFM) [9] as seen in Eq. (14), with Z as the input signal and IF being the IDFM.

$$X(i, j) = \sum_k^K Z(i, k) \cdot IF(k, j) \quad (14)$$

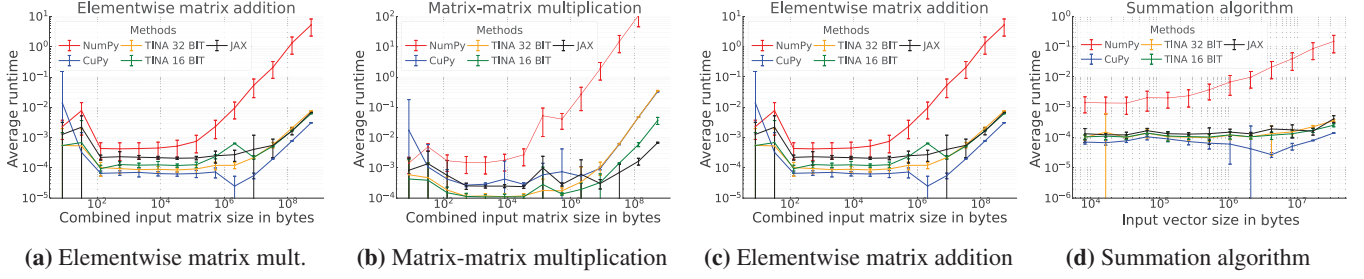


Fig. 1. Runtime of the arithmetic functions vs input size using TINA, NumPy, CuPy & JAX

Using our matrix-matrix multiplication Eq. (13) as further detailed in Section 3.2 we can draw a direct parallel to Eq. (14) by setting the kernel equal to the IDFM.

4.3. FIR filter

In this section, we show how to implement a Finite impulse response (FIR) filter using a TINA layer. Both an FIR filter and convolution involve a linear operation applied to a sequence of input values or signals as seen in Eq. (15) representing an FIR filter. $a(k)$ represents the filter coefficients defining the weights (FIR taps) applied to each delayed input sample. These coefficients indicate how the input signal is weighted at different time indices.

$$y(i) = \sum_k^K a(k) \cdot x(i - k) \quad (15)$$

We can implement this equation using the standard convolution defined in Section 2.1 by setting the dimensions C_{in} , C_{out} , H and M to 1. Then by setting the bias $b(c_{out})$ equal to zero, we get Eq. (16).

$$\begin{aligned} O(0, w, 0) &= 0 + \sum_0^0 \sum_0^0 \sum_n^N I(0 + 0, w + n, 0) \cdot K(0, n, 0, 0) \\ \Rightarrow O(w) &= \sum_n^N I(w + n) \cdot K(n) \end{aligned} \quad (16)$$

Mathematically, an FIR filter represents a discrete-time system with a finite duration impulse response, typically characterized by a set of coefficients determining the filter's behavior. Using the standard convolution, we can simply set the weights of the kernel equal to these coefficients.

4.4. Unfolding algorithm

An unfolding algorithm takes an input vector and produces an output matrix representing subsequences of the input with a successively increasing index. The unfolding algorithm can be represented using the Eq. (17).

$$Y(i, j) = X(i + j) \quad (17)$$

For an input vector X of length I and for an unfolding

window of width J , the output matrix Y has dimensions of $(I - J + 1) \times J$. As an example, for an input vector of length 4, $X = [1, 2, 3, 4]$, and a window of 2, the output $Y = [[1, 2], [2, 3], [3, 4]]$, which is a matrix of dimensions 3×2 .

We can implement this equation using the standard convolution defined in Eq. (1) (Section 2.1) by setting the C_{in} , H and M to 1. Then by setting the bias $b(c_{out})$ equal to zero, we get Eq. (18).

$$\begin{aligned} O(0, w, c_{out}) &= 0 + \sum_0^0 \sum_0^0 \sum_n^N I(0 + 0, w + n, 0) \cdot K(0, n, 0, c_{out}) \\ \Rightarrow O(w, c_{out}) &= \sum_n^N I(w + n) \cdot K(n, c_{out}) \end{aligned} \quad (18)$$

In order to reproduce Eq. (17), first we make the kernel as defined in Eq. (18) a square matrix by setting the kernel dimensions $N = C_{out}$. Then, we make the kernel weights equal to an identity matrix, which reproduces the input at the diagonal (i.e., when $n = c_{out}$), else the input gets multiplied with 0. This results in Eq. (19).

$$O(w, c_{out}) = I(w + c_{out}) \quad (19)$$

5. EXPERIMENTAL RESULTS

In this section, we discuss our experimental results. Section 5.1 compares the performance of the basic functions implemented in Section 3 and Section 4 using TINA (using PyTorch running on GPU) vs other solutions, namely NumPy [10] (runs on CPU), CuPy (automatically runs Python on GPU) and JAX (optimized for NN HW). Two results are shown for TINA, with 32 bits (running on GPU CUDA cores) and 16 bits (running on FPU Tensor cores). Then, Section 5.2 shows the experimental results of a practical use case for implementing a polyphase filter bank in TINA vs other solutions.

The experiments were executed on Google Colab's Tesla T4 GPU (2560 CUDA cores and 320 Tensor cores), alongside an Intel Xeon CPU with 2 vCPUs and 13GB of RAM. By executing our code 100 times and averaging the results, we can show statistically representative performance evaluation. Aside from the TINA 16 bit float implementation, all

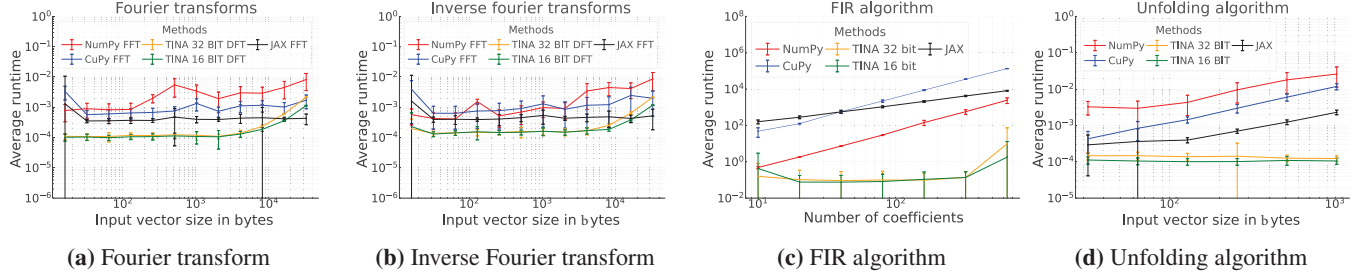


Fig. 2. Runtime of signal processing functions vs input size using TINA, NumPy, CuPy & JAX

implementations use 32 bit floating point input data generated randomly. The measurement start once the input data has been copied to the GPU memory in order to minimise the measurement of overhead.

5.1. Performance of TINA building blocks

Figs. 1a and 1b show that TINA is highly competitive on multiplication-based functions compared to other frameworks resulting in the fastest implementation over a large range of input matrix sizes. For addition-based functions shown in Figs. 1c and 1d, TINA (while competitive) is not as fast as CuPy. This could be attributed to the simplicity of matrix addition functions making it easier for CuPy to find an optimized implementation.

When comparing TINA implementations of DFT and IDFT to implementations in other frameworks as shown in Figs. 2a and 2b, we see that JAX appears to be the fastest, followed by TINA. Here, CuPy shows less competitive performance. Finally, looking at more computationally complex functions with loops, such as FIR and unfolding as shown in Figs. 2c and 2d, we can see that TINA is orders of magnitude faster than other frameworks. This is due to the fact that TINA can map functions with data independent loop iterations much more efficiently than other frameworks.

One important limitation in the TINA framework indicated by our experiment is that the main limiting factor of running operations with larger input vectors or matrices is the large amount of GPU memory used by the TINA framework as we ran into errors once we increased our input sizes beyond what is shown in Figs. 1 and 2.

5.2. Use case: polyphase filter bank

A polyphase filter bank (PFB) channelizes time domain digitized input signals to frequency channels. It is a powerful filter operation that allows for the efficient division of signals into multiple frequency bands, enabling filtering of specific frequency components and simultaneous processing of different frequency components. The PFB is found in many different application domains such as radio astronomy [11, 12], wireless communication [13], radar [14], ultrasound

imaging [15] and quantum computing [16]. A PFB receives as input a signal $x(n)$ decomposed into P branches, denoted as $x_p(n')$ which it turns into a set of subfiltered signals $y_p(n')$

$$y_p(n') = \sum_{m=0}^{M-1} h_p(m)x_p(n' - m) \quad (20)$$

where h_p are filter coefficients (otherwise known as taps) that have been divided between the P branches, and M is the number of taps. The subfiltered signals $y_p(n')$ then will be run through either a Discrete or Fast Fourier Transform resulting in the output signals. We can implement this operation by combining multiple FIR filters as detailed in Section 4.3 and have the resulting output go through a DFT filter as described in Section 4.1.

Fig. 3 shows the speedup results of accelerating PFB using CuPy, TINA 32 bit and 16 bit, as well as JAX (on GPU) as compared to a naive implementation written in NumPy (on CPU). The figure shows that TINA 32 bit achieves significant acceleration of 25x to 80x [12, 17], followed by TINA 16 bit with 20x to 30x speedup, while JAX is the next fastest alternative with 6x to 8x speedup. This shows the significant potential of TINA as a portable accelerator language for signal processing applications.

6. CONCLUSION

This paper presented TINA, a novel framework enabling the development of signal processing applications exclusively through the utilization of convolutions and fully connected layers. The paper showed how to map various functions to NN layers and also demonstrated the potential of this approach. Results show that TINA is able to achieve up to 80x GPU speedup over a naive CPU baseline with NumPy [12, 17] for a PFB algorithm compared to 8x speedup achieved by alternative frameworks.

The main drawback currently found by running TINA layers on a GPU is the large amount of memory needed to run the algorithms. Even though for this paper TINA was run on a server-grade GPU, we found the amount of GPU memory to be a limiting factor in TINA's capacity to process larger applications.

Even though these results are promising, the current pos-

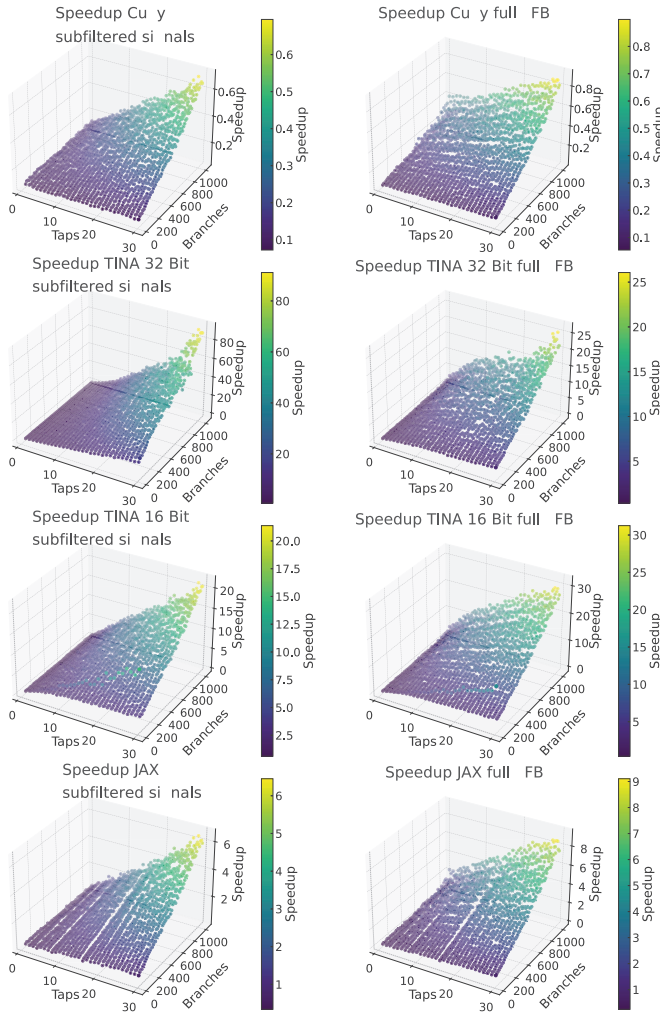


Fig. 3. GPU speedups with respect to NumPy (CPU) of PFB without (left column) and with (right column) Fourier transform in TINA, JAX & CuPy

sible application domains are rather narrow. In order to make TINA more generalisable, further research is needed into mapping more non-NN operations into TINA layers, in addition to investigating ways to make these TINA layers perform efficiently. Furthermore, we will investigate adding more TINA layers (linear, recurrent, etc.) to increase the flexibility of the framework, while keeping in mind the three TINA objectives, namely portability, accessibility, and performance.

Acknowledgements

This project has received funding from the Eureka Xecs TASTI project (grant no. 2022005), Horizon Europe research and innovation program RADIOBLOCKS project (grant no. 101093934) and the Netherlands eScience Center RECRUIT project. We would also like to thank Christos

Strydis, Lennart Landsmeer and Anna Mészáros for help with reviewing the paper and the mathematics.

7. REFERENCES

- [1] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang, "JAX: composable transformations of Python+NumPy programs," 2018.
- [2] Lennart P.L. Landsmeer, Max C.W. Engelen, Rene Miedema, and Christos Strydis, "Tricking ai chips into simulating the human brain: A detailed performance analysis," *Neurocomputing*, vol. 598, pp. 127953, 2024.
- [3] Ryosuke Okuta, Yuya Unno, Daisuke Nishino, Shohei Hido, and Crissman Loomis, "Cupy: A numpy-compatible library for nvidia gpu calculations," in *Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS)*, 2017.
- [4] J. Hoozemans, R. de Jong, S. van der Lugt, J. van Straten, U.K. Elango, and Z. Al-Ars, "Frame-based programming, stream-based processing for medical image processing applications," *J. Signal Process Systems*, vol. 91, pp. 47–59, 2019.
- [5] Joost Hoozemans, Johan Peltenburg, Fabian Nonnemacher, Akos Hadnagy, Zaid Al-Ars, and H. Peter Hofstee, "Fpga acceleration for big data analytics: Challenges and opportunities," *IEEE Circuits and Systems Magazine*, vol. 21, no. 2, pp. 30–47, 2021.
- [6] Xilinx, "Xilinx/brevitas: Brevitas: Neural network quantization in pytorch," July 2024.
- [7] Saad Albawi, T. Mohammed, and Saad Al-Zawi, "Understanding of a convolutional neural network," *2017 International Conference on Engineering and Technology (ICET)*, pp. 1–6, 2017.
- [8] Tolga Ahmet Kalaycı and Umut Asan, "Improving classification performance of fully connected layers by fuzzy clustering in transformed feature space," *Symmetry*, vol. 14, no. 4, 2022.
- [9] K. Ramamohan Rao and P. Yip, *The transform and Data Compression Handbook*, CRC Press, 2001.
- [10] Charles R. Harris et al., "Array programming with NumPy," *Nature*, vol. 585, pp. 357–362, 2020.
- [11] van Haarlem, M.P. et al., "Lofar: The low-frequency array," *A&A*, vol. 556, pp. A2, 2013.
- [12] Danny C Price, "Chapter 1 Spectrometers and Polyphase Filterbanks in Radio Astronomy," *The WSPC Handbook of Astronomical Instrumentation: Volume 1: Radio Astronomical Instrumentation.*, pp. 159–179, 2021.
- [13] F.J. Harris, C. Dick, and M. Rice, "Digital receivers and transmitters using polyphase filter banks for wireless communications," *IEEE Transactions on Microwave Theory and Techniques*, vol. 51, no. 4, pp. 1395–1412, 2003.
- [14] Daniel Zhou, "A review of polyphase filter banks and their application," p. 36, 09 2006.
- [15] Taewan Kim, Choong Lee, Jung-jun Kim, and Tai-Kyong Song, "A new dynamic decimation filter using polyphase macs for medical ultrasound imaging," *Proceedings of the IEEE Ultrasonics Symposium*, 11 2008.
- [16] Johannes Pfau, Shalina Percy Delicia Figuli, Steffen Bähr, and Jürgen Becker, "Reconfigurable fpga-based channelization using polyphase filter banks for quantum computing systems," in *Applied Reconfigurable Computing (ARC)*, 2018, vol. 10824, pp. 615–626.
- [17] Danny C Price, "pfb introduction.ipynb," 2020.