Analogue neuromorphic receiver signal processing for radio astronomy

Hidde Bos

A master's thesis presented for the project at:
ASTRON Dwingeloo

Radboud Universiteit

Radboud University
June 17, 2025

Acknowledgements

Firstly, I would like to express my gratitude to Dr. Johan Mentink for getting the opportunity to do my internship with him and getting me into contact with ASTRON. Also his help with professionalising parts of this thesis were much appreciated.

I'd like to thank Dr. Chris Broekema for helping me wrap my head around a lot of radio astronomy processing that I had many questions about, and for thinking with me towards any problems regarding my thesis.

Additionally, I am grateful for the help given by Dominique Kösters, who I could come to for the theory and implementations of memristors.

Lastly, I'd like to acknowledge Anne and Bas for listening to my confused ramblings followed by "of course, it's simple, really!".

Abstract

Since the introduction of memristors in 2008, there has been a rise of the in-memory computing (IMC) paradigm, which was able to reduce energy costs significantly for many use cases. This thesis analyses the use of IMC within radio astronomy and shows its possible implementation specifically for the front-end processing pipeline of the LOw Frequency ARray (LOFAR). This is done in the form of a literature review, analysing the individual front-end processing components in LOFAR. By finding power consumption data and extrapolating it to the researched pipeline, the findings conclude that the proposed architecture resulted in an energy intensity decrease of 4 ± 1 orders of magnitude. The vast majority of power consumption comes from the analogue delay lines, so future research can be done to analyse and mitigate this component's output. The limitations of accuracy within the proposed architecture are also discussed. The recommended implementation of such new hardware in radio astronomy can nonetheless lead to a significant shift in energy costs, leading to a possible transformation of conventional radio astronomy processing.

Contents

1 Introduction

In the world of today, environmental sustainability has proven to be a major issue. From the recent wildfires in Los Angeles and Australia, increases in droughts to biodiversity losses among other things, there are many challenges that are paired with environmental impacts [1][2][3][4]. In response to the environmental challenges, planetary boundaries have been defined to not go beyond the tipping points, Sustainable Development Goals (SDGs) have been developed to promote sustainable practices by 2030 and the Paris Agreement has emerged to combat climate change globally [5][6][7]. As a result, countries have implemented policies and allocated budgets toward mitigating the negative environmental effects. These initiatives have not only led to effects on the country level, but it has also resulted in awareness and effective changes on personal and institutional levels such as at ASTRON.

ASTRON is the Netherlands Institute for Radio Astronomy, which has supervision of a number of telescopes, among which is the LOw Frequency ARray (LOFAR). A top level overview of this radio telescope is shown in figure 1 [8]. A radio telescope like LOFAR is set up with a high number of receivers, spread out over a large distance. These start with signal processing steps done locally at the receiver end, after which the signal is transported to a central processing facility in Groningen, where further processing is done, in part by use of the central correlator and beamformer [9]. After all the aforementioned processing of the LOFAR telescope, the data can be analysed and processed further by radio astronomers at the end of the process. The total data throughput that is transported over large distances in LOFAR at 150Gbit/s is almost on the level of data throughput of the Amsterdam internet traffic [10].

ASTRON has calculated the greenhouse gas (GHG) emissions of their LOFAR telescope to determine its carbon footprint. The calculated impact - excluding science processing - was found to be 2,681 tCO_2-eq/year, which is equal to the carbon footprint of almost 500 EU citizens [11][12]. It was furthermore found that the annual total electricity consumption of LOFAR is equal to 3,603 MWh. Compared to EU households that consume roughly 1540 kWh per capita annually, LOFAR consumes as much electricity as 2,340 EU citizens [13]. Meanwhile, the Russian invasion of Ukraine has lead to a large increase in electricity costs, which meant that ASTRON had difficulties to keep LOFAR running without emergency help [14]. The significant environmental impact from LOFAR combined with the high electricity costs from the energy crisis are relevant issues, and they should be addressed appropriately. One such way may be with the use of emerging technologies, such as In-Memory Computing.

In-Memory Computing (IMC) has been an emergent paradigm that significantly reduces the power consumption of processing by functionally combining the processor and the memory units in hardware. This paradigm has risen in relevance in part due to the introduction of so-called memristors, which are components that can be used to combine memory and processing. IMC has since then proven useful for many use cases in theory and research, like in condensed matter physics and particle physics [15]. Other than that, there has been a focus on common processing

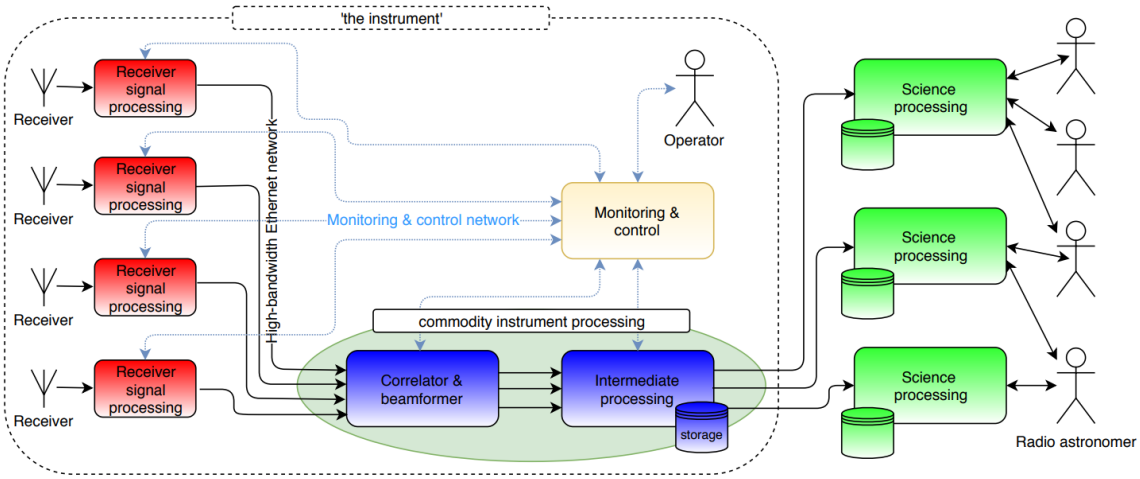

Figure 1: The top level overview of the full LOFAR pipeline. Image taken from [8].

algorithms like the Fourier transform [16]. Research has shown that these IMC implementations are feasible solutions with lower energy intensity and improved latency. For this reason, it is interesting for ASTRON to consider the application of IMC for their telescopes. The reduction of power consumption from electricity usage can lower the associated carbon footprint as well as the monetary costs from electricity. Potentially the capital costs are also decreased, depending on the costs of each hardware unit, their longevity and the amount of hardware needed. Consequently, this thesis aims to analyse the feasibility of implementing IMC in the receiver-end processing pipeline in LOFAR.

The thesis is divided into 5 sections following this introduction.

Firstly, Section 2 goes into the details of radio astronomy, specifically that of LOFAR. This includes the relevant processing steps that will be analysed further and their function in the LOFAR system.

Secondly, Section 3 aims to explain the physical mechanisms behind IMC. This includes analysing the current state-of-the-art technology and their specifications.

Thirdly, Section 4 analyses the implementation of analogue and IMC components into the LOFAR pipeline. This includes information about the power consumption.

Finally, Sections 5 and 6 finalise this thesis by discussing the results and giving the conclusion of the findings.

2 Radio astronomy at LOFAR

2.1 Introduction to radioastronomy

Observing light from outer space is the fundament of astronomy. The whole electromagnetic spectrum consists of a wide range of light wavelengths λ, spanning from the smallest gamma rays ($\lambda < 10^{-11}m$) to radio waves ($\lambda > 10^{-3}m$), which can go deep into the kilometer range. Since varying wavelength amplitudes require different measuring techniques, the area of radio astronomy has to be covered as a separate research topic. Among other things, radio astronomy is interesting because it can observe pulsars, star and planet formations and ionised atmospheres of stars and interstellar plasma, which is significantly blocked by interstellar dust clouds for other wavelengths [17]. Furthermore, since radio wave photons have very small energies, they can be treated as classical waves, which can be amplified coherently and manipulated in receiver processing domains [17]. However, complications also come along with such low energies. Namely, the light is distorted by the outer atmosphere, the long wavelengths need the telescope to span a large area for resolution and sensitivity to stay adequate and there is also radio pollution, which is radio frequency noise that causes interference problems [18]. There are manners of which to counteract these challenges, allowing for sufficient resolution of the radio signal for telescopes like the LOw Frequency ARray (LOFAR), which spans multiple countries to reach a sufficient area.

The processing of radio waves in LOFAR is done in multiple steps. Firstly, there are some minor processing steps done in the analogue domain after which it quickly is converted to the digital domain, as is shown in figure 2. After this, the digital signal is put into a PolyPhase Filter (PPF) consisting of a bank of Finite Impulse Response (FIR) filters and a Fast Fourier Transform (FFT), which then sends the outgoing signal to the Beamformer (BF) as shown in figure 3. Finally, this signal is sent to a central point in Groningen, where all the data of the LOFAR stations across Europe are combined and analysed further. Since the purpose of this project is to analyse the possibility of doing the receiver signal processing in the analogue domain completely, the focus for this section shall mostly lie in the processes that are shown in figures 2 and 3. These processing components are:

- **Analogue domain:**
- Low Noise Amplifier (LNA)
- Beamformer
- Filter
- Amplifier
- Down converter
- Analogue to Digital Converter (ADC)
- **Digital domain:**
- Finite Impulse Response (FIR)

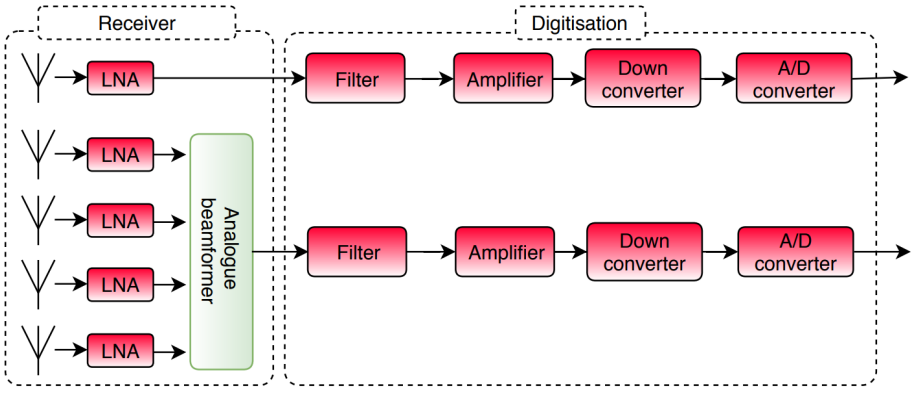

Figure 2: The analogue segment of the receiver signal processing pipeline in LOFAR. Figure taken from [8].

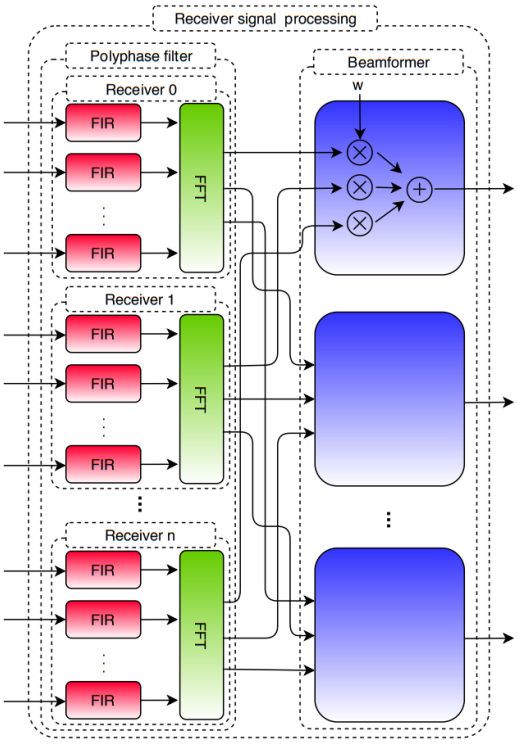

Figure 3: The digital segment of the receiver signal processing in LOFAR, which sends its signal to a central point in Groningen. Figure taken from [8].

- Fast Fourier Transform (FFT)
- Beamformer

These components will be analysed one by one in this section, starting with the minor components that are already analogue. Since there are two beamformers — one analogue and one digital — they will be combined into one subsection. However, it is important to note that only so-called high band antennas have beamformers in the analogue domain [19]. The digital components will each get a separate subsection, starting with the Polyphase Filterbank (PPF) that consists of Finite Impulse Response (FIR) filters and Fast Fourier Transform (FFT) components, which together perform frequency splitting of the signal. Lastly, the beamformers are considered.

2.2 The minor analogue processing components

The Low Noise Amplifier (LNA) is the first minor processing component. Since the incoming radio signal that needs to be analysed at a radio telescope is very weak, the signal has to be amplified to be researched. Such an amplifier would add new noise into the system. Due to the relatively high noisiness of the signal and the importance of minimising the noise for the analysis, it is imperative that the amplifier adds as little new noise as possible, therefore keeping the Signal to Noise Ratio (SNR) as high as possible. Hence, a LNA is used. Similarly, later on in the pipeline, there is another amplifier to raise the power of the signal.

A filter is used in many signal processing systems to isolate certain frequencies of the signal. Telescopes receive a large array of different signals that combine into a complex wave. These filters can be used to exclude all frequencies that are irrelevant for the analysis. For example, the High Band Antennas (HBA) at LOFAR use bandpass filters to let through frequencies between 210-250 MHz or other frequency bands [9]. This induces a ripple effect, as shown in the low-pass filter example of figure 4. In the passband, below the cut-off frequency of the low-pass filter, the signal is ideally let through with the same amplitude as before the filter. Similarly, the stopband after the transition band is ideally attenuated completely. However, this image is not fully realistic, leading to the fluctuating amplitude of the passband amplitudes and the stopband amplitudes, which is called the ripple effect. For LOFAR, the ripple effect is mitigated again once the signal reaches the central processor (CEP) in Groningen.

f_p: Passband edge frequency (cycles/sample)

δ_1: Peak Passband Ripple, in dB: $-20log_{10}(1-\delta_1)$

f_s: Stopband edge frequency (cycles/sample)

δ_2: Peak Stopband Ripple, in dB (rejection): $20log_{10}(\delta_2)$

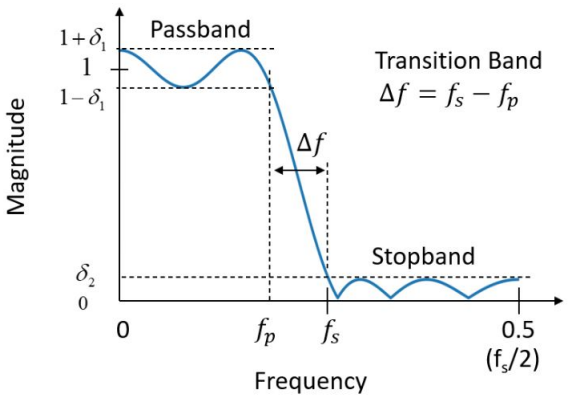

Figure 4: Ripple in passband of a filter, which is a fluctuating amplitude in both the passband and stopband of a filter. Figure taken from [20].

The down converter adjusts the radio signal's frequency such that the radio wave has a lower frequency. This down conversion simplifies the future processing steps by taking the radio signal towards a baseband frequency with a lower minimum sampling rate, without a loss of information.

The Analogue to Digital Converters (ADC) digitises the analogue signal. This is done by reading the voltage value of the wave, where a cut-off voltage is defined. If the voltage of the wave is above this cut-off voltage, the digital voltage value of 1 is given. This could, for example, be 5V. Conversely, if the voltage of the wave is below the cut-off voltage, there is a digital voltage value of 0 as the output, corresponding to 0V. This is done to reduce the SNR and to make use of the robustness of digital processing.

2.3 Finite Impulse Response

Before the digital processing steps are made in the PolyPhase Filter (PPF), the signal is initially split into sub-bands of 156 kHz or 195 kHz frequency bandwidths, depending on whether a 160 MHz or a 200 MHz sampling clock is chosen [9]. Each of these sub-bands are fed into a Finite Impulse Response (FIR) filter by splitting them into 512 equally spaced signals, as shown in figure 5. The incoming wave samples are distributed over all 512 FIR filters in a round robin way. This means that the first FIR filter gets the first sample, the second FIR filter gets the second sample, all the way until the 512th

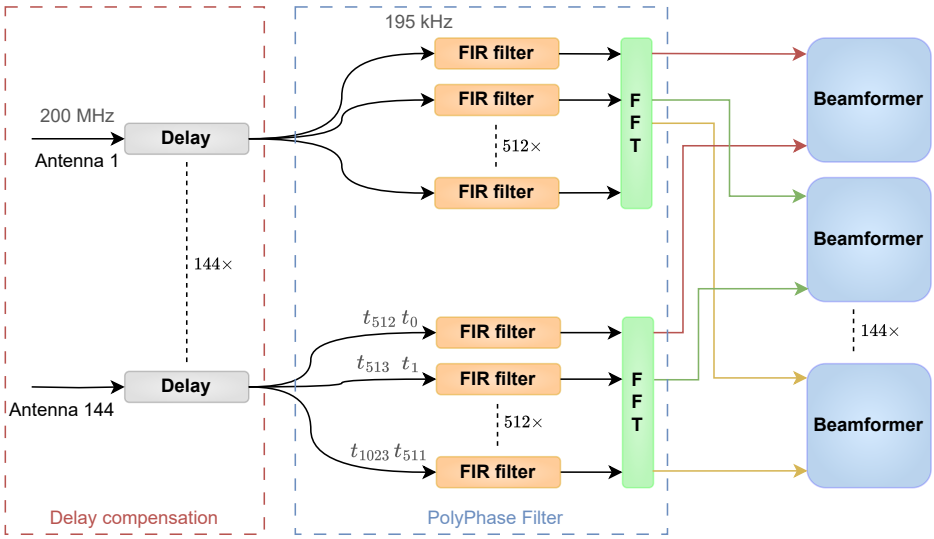

Figure 5: The digital receiver signal processing pipeline with the subdivision of 512 sub-bands per station. Each station then feeds each sub-band into a separate FIR filter which then combines as inputs of one FFT filter. A Dutch LOFAR station has 144 antennas and thus 144 beamformers at the final step of receiver signal processing. In this figure, a 200 MHz clock is assumed, leading to sub-band frequency widths of 195 kHz, 160 MHz clock sampling is also used. Figure adapted from [21].

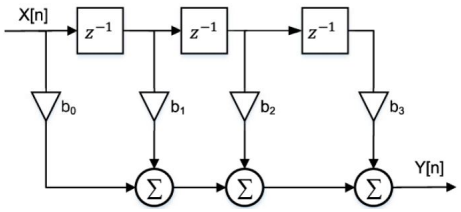

Figure 6: Schematic of a FIR filter, shown here with 4 taps. The input X is put through delays z^{-1}. Then, addition is made after multiplication with weights b, leading to output Y[n] with sample n. The calculations are given in figure 7. Figure taken from [22].

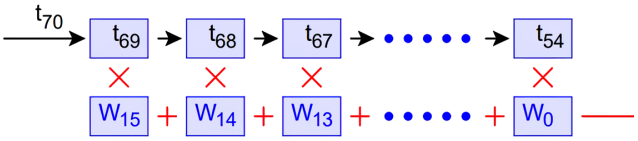

Figure 7: The 16 tap FIR filter calculations, as found in LOFAR. With the delay line, each sample in the filter gets multiplied with one of the 16 weights and then each product gets summed together. Afterwards, each sample gets shifted to the next weight, allowing the oldest sample to be fully moved out of the filter and a new sample is added at the start of the filter. Figure taken from [21].

sample into the 512th FIR filter. Then, the process repeats with the 513rd sample into the first FIR filter again. This does mean that the FIR filters each have a much lower time resolution individually.

At the core of the FIR filter is the impulse response. This means that if an impulse is given, a transfer function turns this input into a response which can be continuous or discrete. For FIR filters in LOFAR, these responses are discrete. The transfer function consists of a finite amount of weights which the impulse gets multiplied by, with a short time interval between each output, the combined total of which form the impulse response. When an FIR is of the Nth order, there are N+1 inputs and associated weights, the former of which is commonly called a tap. Figure 6 shows an example of an FIR filter of the 3rd order, where an input is put through the delay line so that input values across multiple time intervals are multiplied with the weights as shown in figure 7. The FIR filters in LOFAR have 16 taps, meaning that the signal wave's values are read and subsequently multiplied with the weights. Figure 7 shows how the oldest part of the signal t_{54} - i.e. the part of the signal that has been delayed by all 15 unit delays - is multiplied by the last weight W_0 and the newest part of the signal t_{69} gets multiplied by the first weight W_{15}. After these multiplications are added, a newer part of the signal t_{70} moves into the filter, shifting all older signal components to the next weight and pushing the oldest component t_{54} out of the filter. Altogether, this filter can perform the function of a bandpass filter.

The significance of an FIR filter compared to a regular RC bandpass filter comes from the linear phase characteristic that is achieved with an FIR filter. When a signal with a single frequency and another signal with double its frequency are delayed by the same time delay, the phase shift is double for the second signal compared to the first signal. Hence, the phase shift is frequency dependent. Then, if the delayed signals are added up, the combined signal has a different shape than it would have without the difference in phase delay. Conversely, linear phase means that for every increase in frequency, the phase shift increases linearly[23]. In such a case, a time delay will induce the same phase shift for all frequencies. Since it is important in radio astronomy to retain the wave-shape of the signal, it is vital to have linear phase in the FIR filter. This is due to the fact that the Fourier Transform can retrieve the frequencies of the signal, which states something about the source, e.g. the chemical composition.

To achieve linear phase, the FIR filter's coefficients need to be carefully chosen. If there is an impulse response h[n] of an FIR filter of order N, the transfer function is

$$H(z) = \sum_{n=0}^{N} h[n] z^{-n} \tag{1}$$

This must then also have a frequency response, which is the impulse response in the frequency domain or simply the Fourier transform of the impulse response. This is given as

$$H(e^{j\omega}) = \sum_{n=0}^{N} h[n] \cdot e^{-jn\omega} \tag{2}$$

Since a linear phase is required, it is necessary that given an even and real function of ω $\overline{H}(\omega)$, the

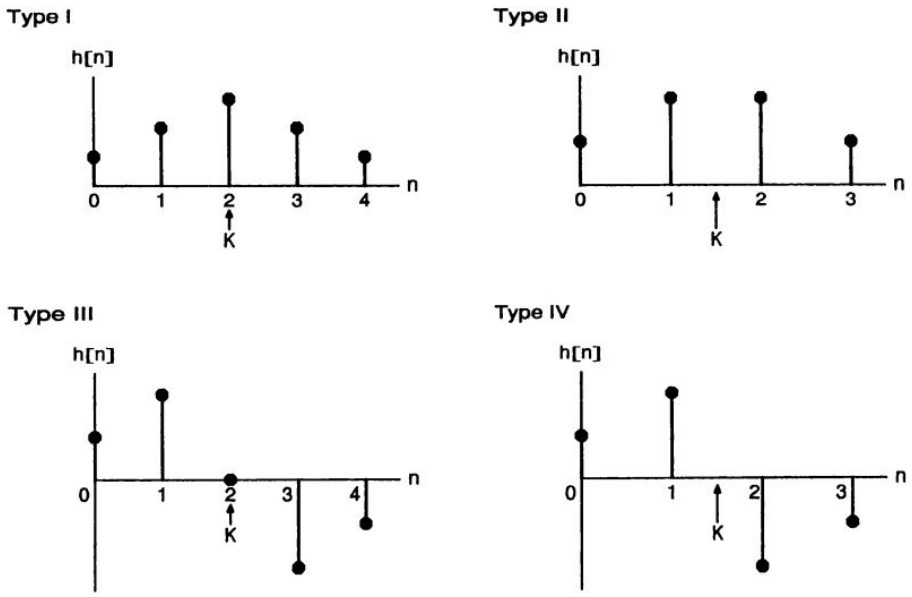

Figure 8: The four types of linear phase. Type I is for even symmetry and even N, type II is for even symmetry and odd N, type III is for odd symmetry and even N and type IV is for odd symmetry and odd N, where N is the number of samples. Figure taken from [24].

frequency response can be written as

$$H(e^{j\omega}) = \overline{H}(\omega)e^{j\phi(\omega)} \tag{3}$$

Where $\phi(\omega)$ is a linear function of ω. From this, there are four possible ways to reach linear phase with an FIR filter[24].

1. Even symmetry with odd length (N is even)
2. Even symmetry with even length (N is odd)
3. Odd symmetry with odd length (N is even)
4. Odd symmetry with even length (N is odd)

These four types are shown in figure 8. In each case, symmetry is necessary and sufficient for linear phase in an FIR filter. As such, the bandpass filter can carefully be constructed in such a way that it has linear phase, given that the weights are chosen to be symmetric.

2.4 The Fourier Transform

As the FIR filters pass a signal within a specific frequency range, the combination of all FIR filters within a receiver give the inputs for the Fourier Transform (FT), as shown in figure 5. Since the combination of all FIR filter outputs form a wave that is a linear combination of all frequencies that have been taken in by the antenna, all frequencies can be extracted from this wave [25]. This is what the FT does, it takes the wave and compartmentalises it into separate frequency bins. The corresponding frequencies are directly related to the energies of the light, which can be linked to emission spectra of certain materials. As such, knowledge about the sources and their chemical compositions can be gained from the output of the FT in LOFAR.

There are 4 different ways to do a Fourier Transform, depending on whether the time is continuous or discrete and depending on whether the time duration is finite or infinite. Depending on what problem the FT needs to solve, any one of these types are valid forms to switch a wave from a time domain to a frequency domain. In quantum mechanics, this is the regular Fourier Transform that is shown in table 1. This is commonly used to switch between the position and momentum representations of wave functions. Furthermore, this integral is also relevant in radio astronomy. The measurement equation, which has proved to be an important mathematical basis for calibration methods and techniques, is itself nothing more than a multidimensional FT[27]. Within signal processing however, the only type of FT that makes sense for a sampled,

Finite time	Infinite time	
Discrete FT (DFT) $X(k) = \sum_{n=0}^{N-1} x(n)e^{-j\omega_k n}$ k = 0,1,...,N-1	Discrete time FT (DTFT) $X(\omega) = \sum_{n=-\infty}^{+\infty} x(n)e^{-j\omega n}$ $\omega \in (-\pi, +\pi)$	Discrete time n
Fourier Series (FS) $X(k) \int_0^P x(t)e^{-j\omega_k t}dt$ k = $-\infty, ..., +\infty$	Fourier Transform (FT) $X(\omega) \int_{-\infty}^{+\infty} x(t)e^{-j\omega t}dt$ $\omega \in (-\infty, +\infty)$	Continuous time t
Discrete frequency k	Continuous frequency ω	

Table 1: The four different types of Fourier Transforms. Table adapted from [26].

finite length signal is the Discrete FT (DFT). This is also what is used in the LOFAR processing pipeline, except there it is the more efficient algorithm called the Fast Fourier Transform (FFT). This will be further examined in Section 2.4.1, but to understand this, the DFT has to be introduced first.

To understand the DFT, one may look first at the most simple scenario, where an input of two values gives two outputs as well. These outputs come from the summation of both inputs, each multiplied by a complex exponential $\exp(-j\omega_k n)$, as shown in table 1. It is easy to see that the outputs are likely also complex, which not only gives information about the frequencies but also the polarisation of the waves[29]. It's beyond the scope of the thesis how this complex data is analysed in further detail. The input as well as the output can be represented as vectors, where the multiplicative factors can be expressed as a matrix. This notation is most commonly used to describe the DFT, where the DFT matrix of size N is shown in figure 9. These matrix components are complex exponentials, given by $\omega_N = e^{\frac{-i2\pi}{N}}$ where i is the imaginary number and N is the size of the DFT matrix, which is equivalent to the amount of inputs and the amount of outputs. The inverse of this matrix is equal to the Inverse DFT (IDFT) matrix, which is a FT going from the frequency domain to the time domain. Together they need to have a normalisation factor of 1/N, which is represented as a $1/\sqrt{N}$ factor in figure 9, implying another $1/\sqrt{N}$ factor in the IDFT matrix. It is also valid to only put a $1/N$ factor in front of the IDFT matrix. The matrix components ω_N are also primitive Nth roots of unity, which means that they are complex numbers that equal one if they are raised to the Nth power, following Euler's identity (fig 10). This will also be an important property for the FFT.

The DFT matrix is a great tool for signal processing, but it is computationally expensive. This is because the input vector is an N×1 vector, which gets multiplied by an N×N square matrix. As such, there is a computational complexity of $O(N^2)$. This can become a problem when N is large, which can be necessary when precise Fourier Transform analysis is required. Ideally, the amount of computations is reduced to lower latency, required capacity and power consumption[25][31]. The solution is provided in the form of the FFT.

$$W = \frac{1}{\sqrt{N}} \begin{bmatrix} 1 & 1 & 1 & 1 & \cdots & 1 \\ 1 & \omega & \omega^2 & \omega^3 & \cdots & \omega^{N-1} \\ 1 & \omega^2 & \omega^4 & \omega^6 & \cdots & \omega^{2(N-1)} \\ 1 & \omega^3 & \omega^6 & \omega^9 & \cdots & \omega^{3(N-1)} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega^{N-1} & \omega^{2(N-1)} & \omega^{3(N-1)} & \cdots & \omega^{(N-1)(N-1)} \end{bmatrix}$$

Figure 9: The DFT matrix, the first row and first column are all ones. There is a factor of $1/\sqrt{N}$ in front of the matrix for normalisation purposes. Figure taken from [28].

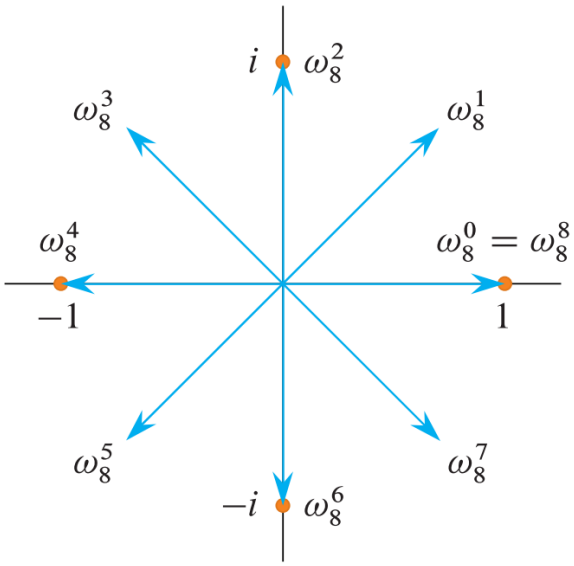

Figure 10: Roots of unity displayed on the unit circle in the complex plane. The complex exponentials ω are given for the Nth roots of unity of size 8. This means that every ω represented on this unit circle will equal 1 if it is raised to the power of 8. Figure taken from [30].

2.4.1 Fast Fourier Transform

While Gauss first discovered the FFT over two centuries ago, it went wholly unnoticed until it was rediscovered 150 years later [32]. In 1965, J.W. Cooley and J.W. Tukey did pioneering work for the FFT, based on the Danielson-Lanczos Lemma that was published during the second world war[33]. This Lemma states that a size N DFT can be rewritten as a sum of two DFTs of size N/2 [34]. This is given as:

$$F_n = \sum_{k=0}^{N-1} f_k e^{-i2\pi nk/N} \tag{4}$$

$$= \sum_{k=0}^{N/2-1} f_{2k} e^{-i2\pi nk/(N/2)} + \sum_{k=0}^{N/2-1} f_{2k+1} e^{-i2\pi nk/(N/2)} \cdot \omega^n \tag{5}$$

Where $\omega^n = e^{-i2\pi n/N}$. This Lemma can then be extended to be broken down in smaller segments by subdividing the numbers by even-ness and odd-ness. This is most simply done by looking at bit notation of the numbers using little-endian bit representation. In the case of 3 bits, the even numbers look like [xx0] and the odd numbers are [xx1]. In other words, even-ness is shown by the final index being a 0. Similarly, subdividing these even and odd numbers into more even or more odd parts is done by looking at the second index. This means [x0x] is more even than [x1x]. As such, 0 ([000]) is the most even number and 7 ([111]) is the most odd number in a 3-bit system. This rewriting ultimately does not remove complexity, since the recursive half sized summations still span the whole N sized FT. However, complexity can be reduced by this lemma in combination with the observation that the DFT has a lot of symmetry. This is taken advantage of in the Cooley-Tukey FFT algorithm to reach a complexity of O(Nlog(N))[35]. For this reason, it is important to make the FFT be of a size that is a power of 2 to allow all recursive steps, which allows for maximum efficiency.

If an N x N DFT matrix is taken with N as a power of two (if needed, add 0s until the matrix size is a power of two), reshuffle the evaluation points $f_0, f_1, ...f_{N-1}$ (equation 4) to even and odd parts. The DFT matrix then is simplified to two new matrices filled mostly with zeros as shown in equation 6. The I_{512} stands for unitary matrices of order 512 (so 512 x 512), F_{512} refers to a DFT matrix specific to either the even numbers or the odd numbers and D_{512} refers to a diagonal matrix of $1, \omega, \omega^2, ...\omega^{511}$.

$$F_{1024} f = \begin{bmatrix} I_{512} & -D_{512} \\ I_{512} & -D_{512} \end{bmatrix} \begin{bmatrix} F_{512} & 0 \\ 0 & F_{512} \end{bmatrix} \begin{bmatrix} f_{even} \\ f_{odd} \end{bmatrix} \tag{6}$$

$$F_{1024} \rightarrow F_{512} \rightarrow ... \rightarrow F_4 \rightarrow F_2 \tag{7}$$

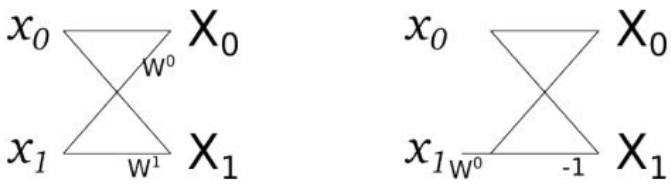

Figure 11: Improved butterfly diagram with Twiddle Factor $\omega^k = -\omega^{k-N/2}$ for $k \geq N/2$. Instead of the diagram with two twiddle factors (left), it can be simplified to a multiplication with one twiddle factor and simply a multiplication with -1 (right). Figure taken from [39]

To more easily represent the computations of the FFT, butterfly diagrams are used. The butterfly diagram for the simplest size-2 FT is shown in figure 11. The two butterfly diagrams in this figure are identical, but generally the one on the right is used as simplified notation [25]. The output signals receive a combination of the two input signals. Namely, X[0] receives x[0] + $\omega^0 \cdot$ x[1] and X[1] receives x[0] - $\omega^0 \cdot$ x[1]. These weight factors $\omega_N = exp(i2\pi/N)$ are called twiddle factors in the FFT, equal to the aforementioned root of unity components for the DFT. For N = 8, ω^0 is equal to 1, ω^4 is equal to -1, etc.

Figure 12 schematically shows the FFT after the even and odd components are separated once by use of butterfly diagrams [36]. Two DFT calculations are shown here, which require multiplications on the order of $N^2/2$. In accordance with the Danielson-Lanczos Lemma, the N/2 DFT calculations can be again subdivided into smaller segments, as shown in figure 13. The first splitting of even and odd indices only lowers the complexity to the order of $N^2/2 + N$, due to the $(N/2)^2$ computations of each smaller DFT matrix and the N final butterfly computations, as shown in figure 12. But splitting the DFTs down to even smaller variants leads to lower complexities. If the FFT size N is defined as $N = 2^r$ or $r = log_2(N)$, then it can be found that after r splitting steps that:

$$C = \frac{N^2}{2^r} + r * N = N + Nlog_2(N) \tag{8}$$

Where C is the complexity of the FFT[31][37]. For large N, it is therefore evident that the complexity is on the order of O(Nlog(N)).

In these examples above, the assumption is made that the FFT is split up from the original DFT matrix size down to many DFT matrices of size 2x2. This is called the radix-2 FFT, which is also what is used in LOFAR [31]. Similarly, radix-4 means dividing the DFT matrix into small DFT matrices of size 4. Since the last step of subdividing size 4 DFTs into size 2 DFTs isn't conserving much energy, it can be omitted. In both cases, the total amount of computations, latency and power consumption goes down significantly for large N [31][38].

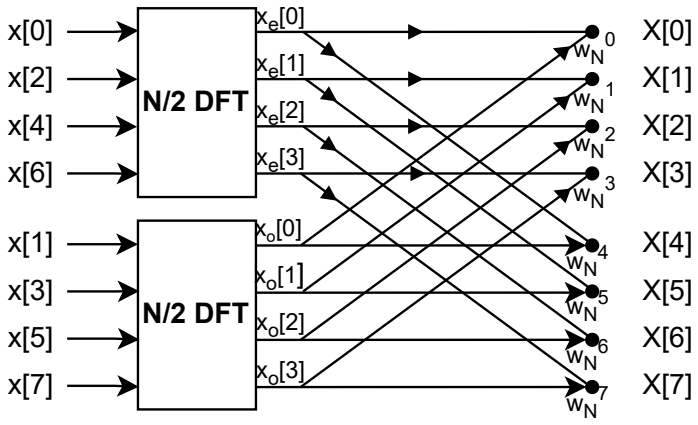

Figure 12: A diagram for N=8 which shows the first step towards an FFT from a DFT. First the even and odd parts are separated and smaller DFTs are calculated. This architecture reduces multiplicative complexity from N^2 in the DFT to $2 \cdot (N/2)^2 + N$. w_N^k is the twiddle factor for a FT of size N. Figure adapted from [36].

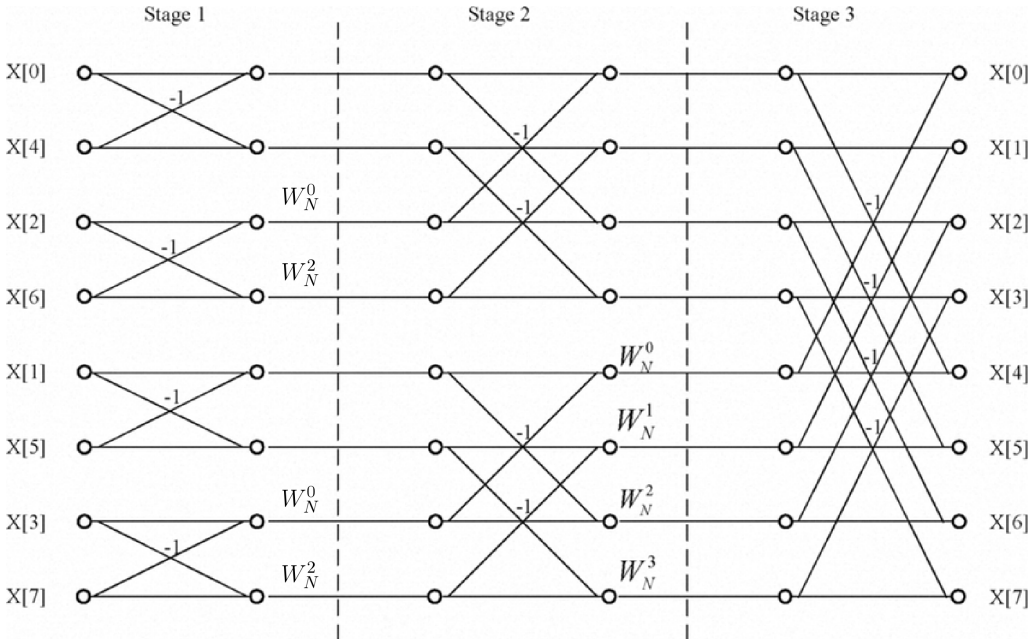

Figure 13: The full diagram for N=8 showing how the FFT simplifies DFT calculations, including recursion. Figure taken from [40].

The FFT computations from figures 12 and 13 are done using the Decimation In Time (DIT) algorithm, meaning that the signals x[n] are decimated into smaller segments. Alternatively, one could use the Decimation In Frequency (DIF) algorithm, which decimates the X[k] into smaller segments. These two algorithms create butterfly diagrams that are flipped compared to each other. This is shown in figure 14. The DIT algorithm has pre-ordered time signals and the DIF algorithm has pre-ordered frequency signals[42]. The order of the shown indices are representative of the even-ness or odd-ness of said signals, which works the same as before with the Danielson-Lanczos Lemma. Since the FFT matrices have a size that are a power of 2, the indices can perfectly be written in bit representations. As a result, the time signal indices have inverted bit representations compared to the frequency signal indices. For example, for N=8 this means that [xyz] indices in time domain lead to [zyx] indices in frequency domain.

The multistage FFT works the same as the aforementioned butterfly structures, but now there are multiple in succession. Multiplication needs to be applied to this signal before addition is applied, which then feeds into the next butterfly structure until the output is reached. As an example, in figure 14 the calculation for X[1] is given as

$$X[1] = x[2] \cdot -1 + x[0] + W_4^1(x[3] \cdot -1 + x[1])$$

where $W_4^1 = -j$ is the Twiddle Factor for N=4, k=1 and j is the imaginary number.

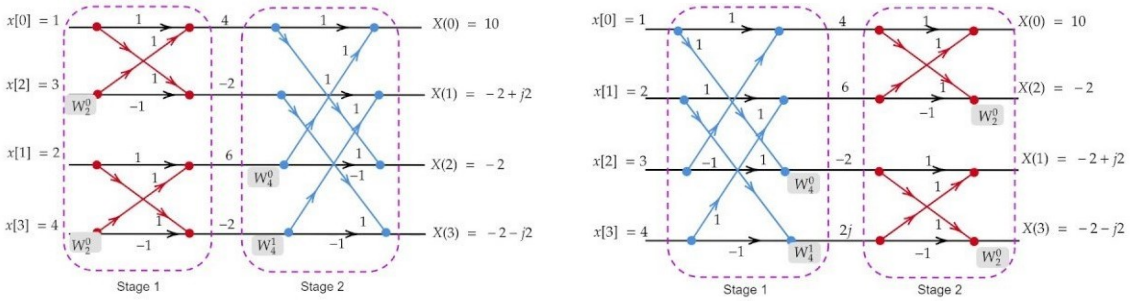

Figure 14: The difference with DIT (left) and DIF (right) butterfly diagrams with input examples. While the setups are designed differently, the same outputs are found. Figure taken from [41].

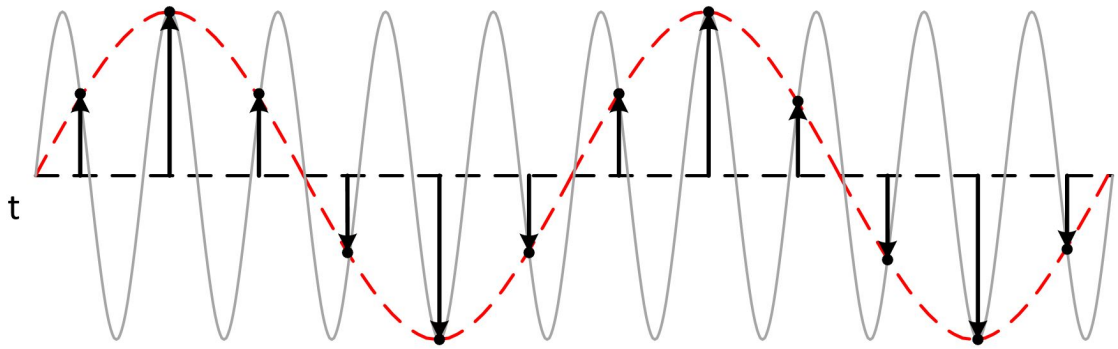

Figure 15: Example of aliasing, for the case where the sampling frequency is 1.25 times the frequency of the sine wave. Since the sampling frequency does not surpass the Nyquist frequency, the perceived wave does not closely resemble the real wave. Image taken from [43].

2.4.2 Frequency related artifacts

Next to conceptual understanding of the FT, it is also necessary to understand how it retrieves frequencies within the electrical systems that it resides in. In the figures 13 and 14, it is shown how the signals result in an output of X[i], with i any integer from 0 until N-1. These X[i] are frequency bins, with discretised frequency values [31]. X[0] is always the DC signal with a frequency of 0 Hz. After that, each frequency bin goes up in frequency values in equally spaced increments. Since the Fourier transform has to span a specific bandwidth B that should be analysed, it is evident that each frequency interval Δf between the frequency bins should equal B/N with N the size of the FT. For example, if a 156 kHz bandwidth is analysed with a FT of size N=512, then $\Delta f = 381Hz$. This means that there is an inaccuracy in the system, as the frequency resolution is limited by the allowed values in the frequency bins. However, this inaccuracy can be made insignificant if the ratio N/B is maximised. This could mean that the size of the Fourier transform has to be doubled, which also requires double the amount of FIR filters. This not only increases the amount of computations ($\sim$ linearly with N), but the sampling frequency also needs to be thought of due to the Nyquist theorem.

The Nyquist frequency f_n is very relevant for the PPF, since it is the highest frequency that can reliably be measured at a specific sample rate f_s. The Nyquist frequency is defined as $f_n = f_s/2$, meaning that if the sample rate is 2 kHz, the highest measurable frequency is 1 kHz [23]. As such, if a bandwidth up to the highest frequency of 250 MHz needs to be analysed at LOFAR, the sampling rate should be at least 500 MHz. Figure 15 shows what happens if the Nyquist rate is not reached. Here, there is just a simple sine wave of frequency 1 Hz (the gray line) and the sample rate is 1.5 Hz. If a sine wave is extrapolated from these samples, the result does not correspond to the real sine wave of 1 Hz. This is a similar effect to the shutter speed of a camera, that can not capture the rotating speed of rotors of a helicopter, making them look seemingly not move at all. If all frequencies above f_n are excluded from the bandwidth, there is only one set of frequencies that can form the measured samples. However, without taking the Nyquist theorem into account, there can be an infinite amount of waves that can describe the measured samples. This means that tuning the sampling rate and bandwidth carefully is essential for signal processing.

When doing a FT, the theory assumes that the sampled signal is repeated periodically infinitely long. In practice, this implies that the signal is sampled for a whole number of periods. It is not only hard to tune the equipment to achieve this, the incoming signals are too complex to find and predict any full periods of the signal. This leads to leakage, which means that energy seems to "leak" into all other frequency levels, due to uncertainty in the signal frequencies [31]. To counteract this effect, weighting functions called windowing functions can be used [37]. This reduces the amplitude at the boundaries of the sampling segments, to reduce the leakage at these points. Since a reduction of leakage also lowers the frequency resolution [31], various windowing functions are commonly used, each with their respective use cases.

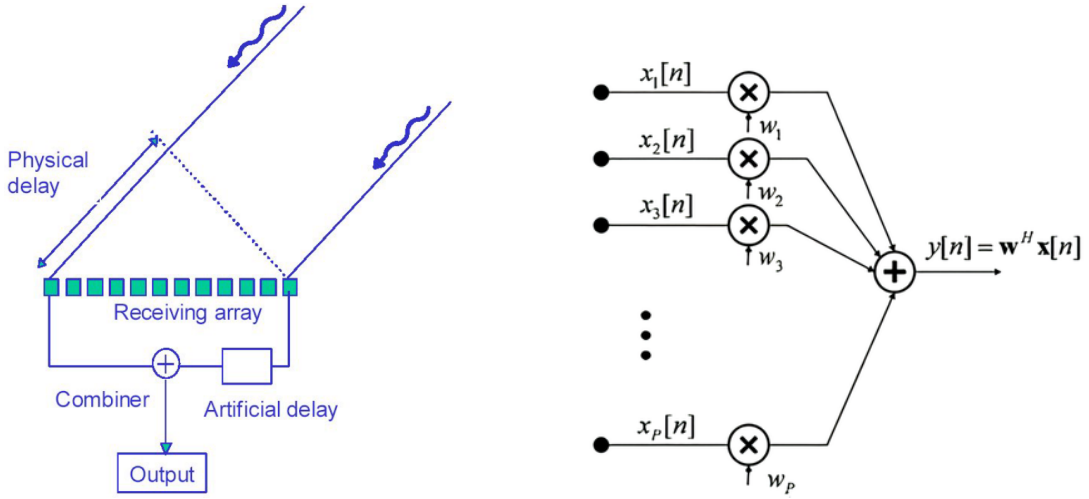

(a) (b)

Figure 16: Schematics of a beamformer. a) The delay of an incoming wave. Since it takes longer for the wave to reach one antenna than another, there is an artificial delay implemented to correct any phase shift. After this, the signals are combined to get a strong signal from the desired light source. b) The inner workings of the beamformer. Each input gets multiplied by their corresponding weights and finally combined. These weights have the purpose of steering the directionality of the antennas. Figures taken from [44] and [45], respectively.

2.5 Beamformers

Before the filtered signal can be sent towards the central processing in Groningen, one final beamforming step needs to be taken. The function of this step is the same as the analogue beamformer at the receiver immediately after the initial LNA. The shorter wavelength HBAs have such an analogue beamformer incorporated into their processing, while LBAs skip this step [19]. All antennas lead to the final digital beamforming step, however. This step is necessary, because adding them together reduces the amount of processed data by a factor of the number of antennas. The beamforming step is also necessary because the antennas take in all kinds of signals, from many angles and wavelengths. This means there is a large influx of noise, leading to a low SNR. Since the antennas in LOFAR lie flat and cannot move to point in a certain direction, conventional addition of the antennas signals would equate to the antennas pointing to zenith (pointing directly upwards). The signals coming from zenith will then add up perfectly, but signals from other angles will have a phase delay compared to one another, as shown in figure 16a. These signals will therefore have a reduced intensity. Since the sources that need to be analysed are usually at an angle, there needs to be a solution to still be able to "point" the antenna towards other angles. Figure 16a shows how the beamformer provides this solution. It does this by inputting an artificial delay into all antennas to align them towards the required angle, getting rid of any phase mismatches between antennas and thus increasing the SNR.

The delaying nature of beamformers can be described accurately by the application of weights to the input signals, as shown in figure 16b. Each antenna receives a signal x_i that can be multiplied by a complex weight, which points the antenna in the desired direction. To define these weights, it can be easier to define the offset of the incoming wave. Since LOFAR deals with far-field sources, the incoming waves can be seen as plane waves as per the Van Cittert-Zernike theorem. This significantly makes the process easier to describe [46]. If the desired signal is defined as x_0 with steering vector s, which acts on the signal $x(t)$, it is found that:

$$x(t) = x_0 * s \tag{9}$$

$$s = \begin{bmatrix} 1 \\ exp(-j2\pi dsin(\theta)/\lambda) \\ exp(-j4\pi dsin(\theta)/\lambda) \\ exp(-j6\pi dsin(\theta)/\lambda) \\ \vdots \\ exp(-j(M-1)2\pi dsin(\theta)/\lambda)] \end{bmatrix} \tag{10}$$

Where j is the imaginary number, d the distance between antennas, θ the angle of arrival and λ the wavelength of the incoming signal. This steering vector is a mathematical representation of the extra distance that light needs to travel depending on which antenna the light reaches. This is shown in figure 17, where it is expressed as a phase delay. Here it is shown how, in accordance with the vector shown above, there is an extra distance $d \cdot sin(\theta)$ travelled by light for each antenna it passes. Naturally, if the angle of arrival is 0 and thus the signal comes from zenith, all components of the steering vector are 1 just as expected. To counteract the phase change that is induced by the steering vector, one can apply the weights shown in equation 9 if the Hermitian conjugate is taken. These weight values can easily be tuned in multiple digital ways, even when the beamformers are in the analogue domain [47]. The ideal signal x_0 is maximised again following the simple characteristic that $s^H s = I$ with s^H the Hermitian conjugate of the steering vector and I the identity matrix. This holds true because the steering vector is Hermitian. Thus, the ideal signal can be maximised this way.

Real antenna arrays get more than only deep-space sources' light signals. There could be interferences from other near-field sources, including satellites and other man-made signals. As a result, there is not only a peak in signal from the angle at which the deep-space object resides, but also at angles from interfering sources. This, in combination with examples of beam patterns, is shown in figure 18. A deep-space object is shown here, which is the subject of analysis by use of the antenna. By using the beamformers, the antennas can be prepared to accurately receive the desired signal [49]. This is shown by the main lobe, which is the tall lobe in the centre of the beam pattern. The height of the lobe implies that the antenna is very sensitive to signals coming from that direction, compared to the rest of the beam pattern. There are also local maxima, shown as side-lobes. This is an inevitable feature of the antenna beam pattern. The beam pattern can, however, be altered, which is done for two reasons. Firstly, the main lobe has to be shifted along with how the deep-space object moves in the sky due to the earth's rotation. Secondly, the beam pattern can be reshaped such that the interference falls into minima of the beam pattern (spatial nulls) instead of into the side lobes. These spatial nulls mean that any signal coming from that direction is completely mitigated by the phased array. Figure 18 shows exactly how the interfering satellite's signal falls into a spatial null, yet the deep space object falls into the main lobe. Furthermore, grating lobes can also be found when the antennas are separated by a distance of more than $\lambda/2$, in which case it can be solved by carefully putting the antennas in a random grid.

There are multiple ways to retrieve the ideal signal from the total signal — which includes the interference — which is done by the use of tapering functions [17][50]. Firstly, a so-called rectangular window can be used, by tuning the weights in a way that maximises the signal. This is by putting the weights such that they act as the Hermitian conjugate of the array steering vector,

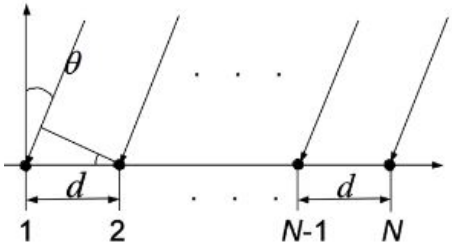

Figure 17: Schematic showing the extra distance that each wave travels depending on which antenna it reaches. For example, the wave at antenna 1 has travelled an extra $sin(\theta) \cdot d$ compared to the incoming wave at antenna 2. Figure taken from [48].

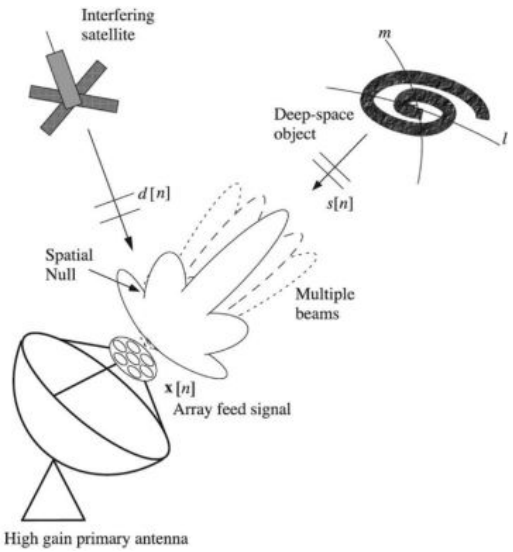

Figure 18: The signal from both the deep space object that needs to be researched and from interfering objects. The beam pattern of the antenna is also shown. This pattern can be changed — as it's an antenna characteristic — so the main lobe points to the signal of interest and interfering signals are in directions of spatial nulls. Figure taken from [49]

which was shown in equation 9. This is the simplest idea, but there exist more tapering functions like the Hamming function, which increase the SNR while losing intensity for the ideal signal. This tapering function not only affects the phase delay of each incoming signal, but also how they are amplified. Currently, LOFAR uses a simple rectangular window, with uniform amplification of all signals and only phase delays to maximise the signal.

As mentioned before, there are 2 different beamformers at the local processing steps in LOFAR. While the analogue and digital beamformers perform the same tasks, there is still a difference between them for how LOFAR is set up. Figure 19 shows the full setup of antennas within a LOFAR HBA station. Firstly, a set of 16 antennas are taken together in what is called a tile. The analogue beamformers combine these individual antennas' beams in the manner described in this section. In the same way, the digital beamformers combine the signals of all the tiles, which span the entire station.

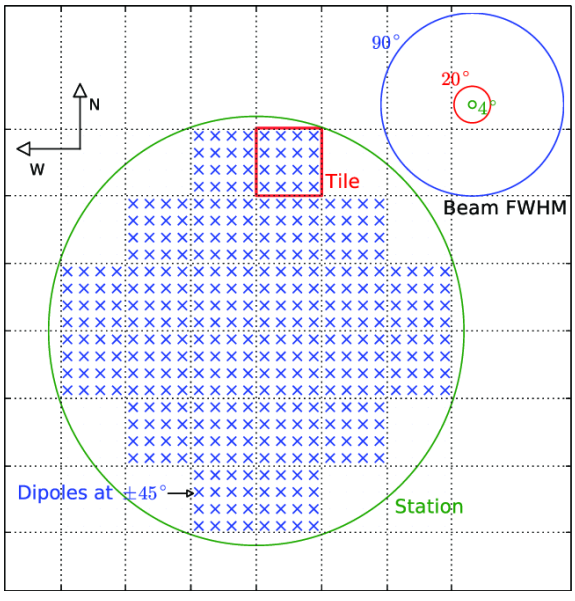

Figure 19: The tile setup for a HBA station in LOFAR. A group of 16 antennas fill up a tile, which similarly combines into a station. Figure taken from [51].

3 Neuromorphic computing

3.1 Introduction to neuromorphic computing

Computing machines have been around since ancient times. The abacus, for example, has been around as early as 4 millennia ago in Mesopotamia [52]. These early machines were then used to help humans with complex computational tasks. Typically they could execute only one type of program, such as arithmetic in the case of the abacus. If other operations were needed, a new computing machine needed to be built. In 1945, John von Neumann published a design for a digital computing machine that could run many kinds of programs — this design is known as the von Neumann architecture and it is the computing model still used globally today [53][54]. The architecture is based on the idea where a part of storage is used for a program — which is simply a set of instructions — which meant that separate processing units and memory units are being used. As a consequence, computations are done by sending electrical signals between the processing units and memory units.

Due to the inherently serial design, the von Neumann architecture is not efficient for large scale complex calculations. Despite the serial design of computers, decades of downscaling still allowed to do very complex calculations [55]. However, enormous amounts of data transport inevitably leads to a huge energy consumption. This can easily be seen by how a supercomputer from 2011 compares to the human brain. It was found to have 10 times as much memory storage as the human brain and about 4 times better calculation times (in FLOPS) [56]. Meanwhile, this supercomputer uses 1 million times more energy than the human brain[57]. The difference in energy consumption is due to the fact that, unlike computers, brains do not follow the Von Neumann architecture. Memory and processing in brains are integrated rather than separated, which means the system is highly parallel rather than serial. It is then interesting to see if such parallel systems can be used for very complex calculations, with the benefits of the extreme energy efficiency that is seen in the human brain. Such computing forms that are inspired by the architecture of the human brain is known as neuromorphic computing [58].

One of the use cases of neuromorphic hardware is analogue in-memory computing (AIMC), which is energetically favourable for two reasons. Firstly, the data transport between memory and processors is no longer necessary [59]. For large systems like the matrix vector multiplications (MVM) that are important in machine learning, retrieving the matrix values from memory is where the majority of energy consumption resides. Secondly, analogue computations only need a few resistors and capacitors, whereas digital computations may need in the order of a thousand transistors just to do one multiplication, depending on the required precision [60]. The latency of the computations also improves significantly in large systems, as all computations can be done simultaneously instead of serially.

The brain is made up of two major building blocks, neurons and synapses. A human brain

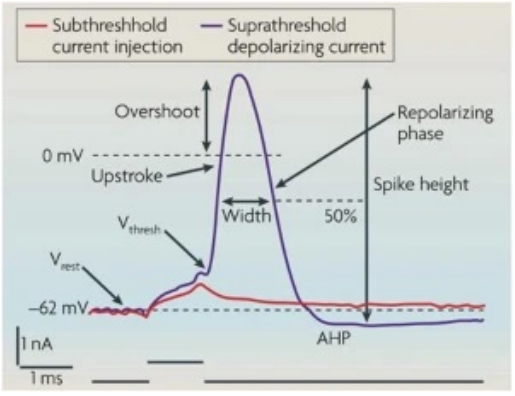

Figure 20: Action potential of a neuron. Once the voltage passes a threshold, there is a spike, called the action potential. This signal can be passed on to the next neuron in a neural network. Image taken from [61].

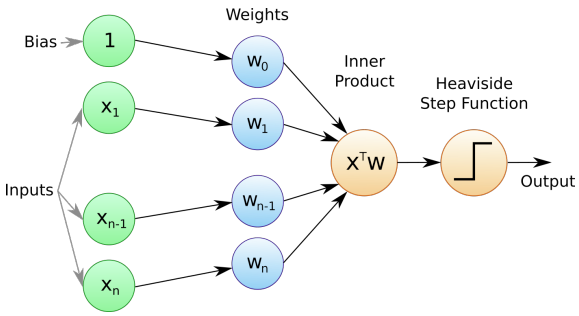

Figure 21: The most simple perceptron model. There is an n number of inputs with assigned weights which can be negative or positive. The weighted sum of the inputs have to surpass a threshold, which is set by the bias, so that the output is activated. If the threshold is not reached, there is no activation. Image taken from [62].

consists of 86 billion neurons and each neuron has 15 thousand synaptic connections to other neurons, a network of a population of neurons and synapses is called a neural network [63]. When an outside impulse is received by a human, for example by seeing a bowl of soup, a group of neurons will receive an electrical impulse. If this electrical impulse at a neuron surpasses a threshold voltage in a neuron, there will be short yet high peak in the potential, called the action potential or a spike (fig 20) [61]. The action potential in a neuron will then release neurotransmitters at the synaptic junction, which can produce a spike in the next neuron if the threshold voltage is reached. Each spiking neuron will start this process, so even looking at a bowl of soup can lead to spiking patterns of a high complexity. Still, the basic concept of the interplay between neurons and synapses — the latter of which having incorporated memory[64] — paved the way for learning to be done by computers, both in software and hardware. This lead to the introduction of Spiking Neural Networks (SNN) and In-Memory Computing (IMC) as topics of interest. This thesis only covers the latter one of these two neuromorphic computing paradigms. To understand this, artificial neural networks have to be introduced.

In the late 1950s, Frank Rosenblatt proposed an algorithm called the perceptron [65]. This is an algorithm, used for classification, that tries to imitate the biological system of neurons and synapses in a simple form — this is an example of an Artificial Neural Network (ANN). In figure 21 the basic concept is illustrated, starting with the layer of either activated or inactivated inputs, which are also called artificial neurons. Similar to the synapses in the brain, these inputs are multiplied by either inhibitory (negative) or excitatory (positive) weights. The weighted sum then is compared to a threshold from a separate neuron called the bias, which determines what the value of the output will be. This is done by taking the inner product of the weights and the transpose of the weights, which in this most simple, single-layer example is simply a element-wise product. However, a single-layer perceptron isn't sufficient for more complex computations. More layers have to be added, which will lead to systems more closely related to the neural networks that are found in brains. Furthermore, perceptrons only have one neuron, whereas multilayered ANNs have multiple neurons that are all connected to an input. These multilayered ANNs — also called Deep Neural Networks (DNN) — are then able to do more complex tasks [66]. Contrary to the element-wise products of the single-layer perceptron, the DNN works with Matrix-Vector Multiplications (MVM).

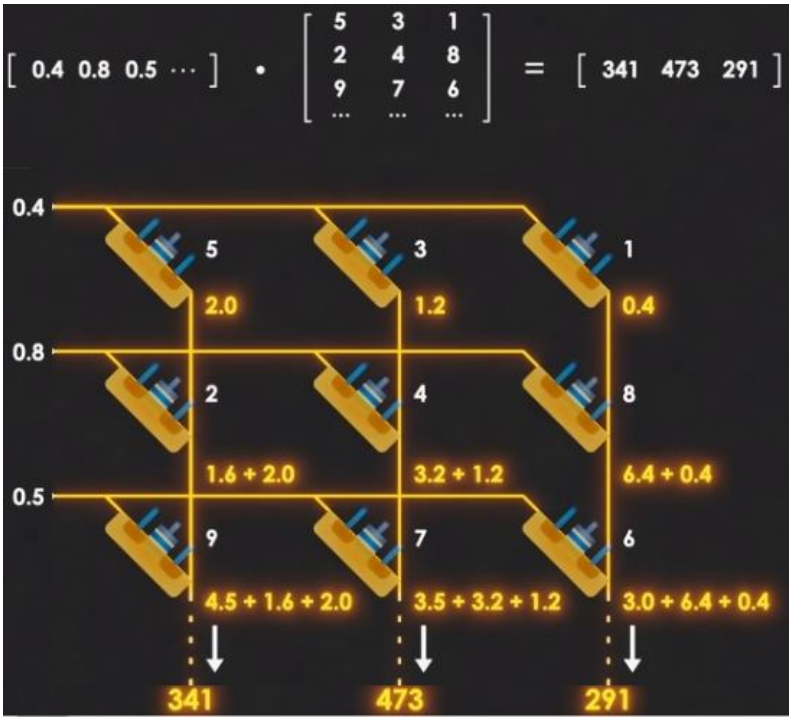

Figure 22: An MVM realised in the form of a memristive crossbar. The signal starts with a certain voltage value, which passes through a component with conductance G. Via Ohm's law $G \cdot V = I$, the incoming voltage value is translated into a new current signal. Then signals from several rows, coming with different amplitudes, are added up together via Kirchhoff's laws. On top, the calculated MVM is shown. Image taken from [67].

Activation functions in ANNs can be any mathematical function that determines the output of a neuron. A very simple one may be a step function, like the one shown in figure 21. If the incoming signal is below a predetermined value, the output is 0 — so there is no activation — otherwise it is 1. Though this can be relevant for simple classifiers, non-linear functions like the sigmoid function may be used for more complex systems [68]. An important distinction can be made here between spiking networks and non-spiking neural networks. The former either produces a spike in a neuron or it does not, with no in between, which is most like a biological brain. The latter may also allow intermediate values, increasing the potential complexity, which is what is used for the IMC use case of this thesis. Here, the complex neural systems can be replicated by embedding them into hardware.

In the case of MVMs such as the one seen in figure 22, the computations consist of additions and multiplications. Additions can be considered by Kirchhoff's law, stating that two wires that combine into one wire will add up the two currents as per: $\sum_i^N I_i = \sum_f^N I_f$ with I_i the initial currents before the junction and I_f the final currents after the junction. Multiplications can be formed with the use of resistances/conductances. The conductance G as the inverse of the resistance leads to Ohm's law as: $V = G \cdot I$. When neural networks are considered with MVMs, just like in machine learning, the weights or synapse strengths can thus be represented by conductances and the additions of each action potential from the previous neuron is represented at the wire junctions. In figure 22, a crossbar consisting of resistive devices is shown. The input vector is represented by voltages and each voltage leads to N different resistive devices. These resistances result in an effective multiplication of $V \cdot R$ to get a current output. Each current value in a column then gets summed at each junction, whose results combine into an output vector. For this project's purposes, the resistances are representative of the weights of the DFT matrix, for example. Since these weights have specific values, it is important to use components that have tunable resistances at a nanoscale. This is where the memristor comes in.

3.2 Memristor

In 1971, Leon Chua postulated the existence of a fourth fundamental electrical circuit element, besides the known resistor, capacitor and inductor. Since the four physical quantities of charge, voltage, current and flux all have a relationship with each other, Chua reasoned that — due to symmetry — the relation between charge and flux had to exist, including a circuit element [70]. This "missing" circuit element was called the memristor, as shown in figure 23. A useful model or physical example of a memristor was not found until 2008, however, when HP found experimental

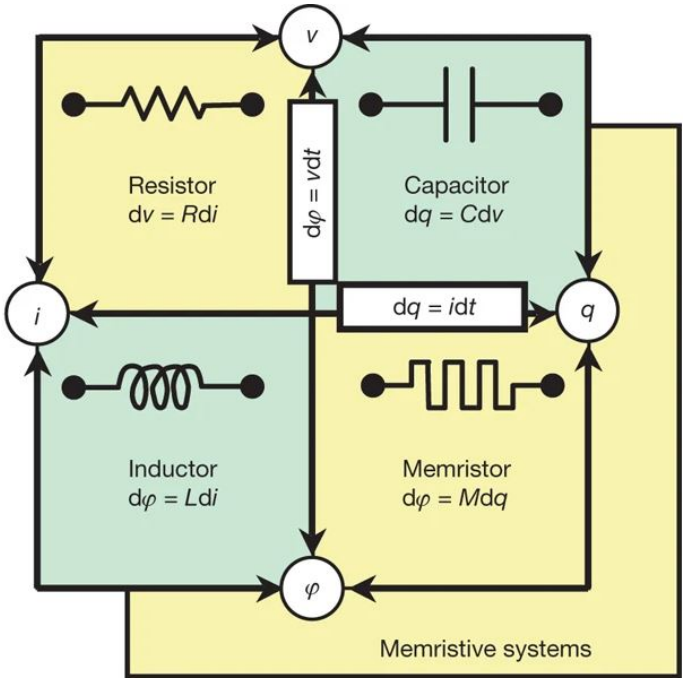

Figure 23: Memristor theory as the 4th (missing) electrical unit. In an electrical system, ϕ is the flux, v the voltage, i the current and q the charge. Each of these physical quantities are related to one another by the shown equations and the electrical devices. Figure taken from [69].

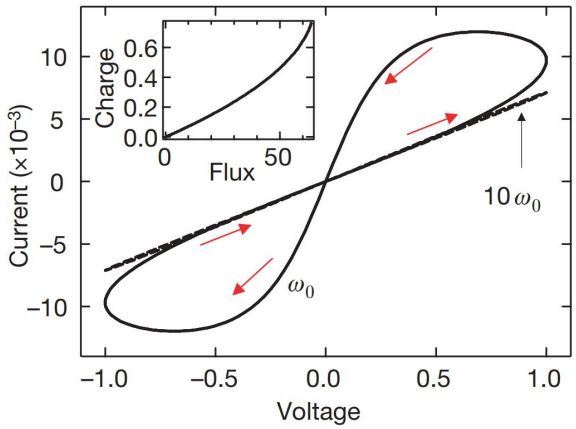

Figure 24: The hysteresis loop in the I-V characteristics for a memristor. If a sufficient voltage is applied, the slope — and thus the resistance — of the graph changes. If the voltage is removed, and thus also there is no current, and then reapplied again, the new slope is once more found. In other words, the slope is 'remembered'. This process can be reverted by reversing the voltage. If the switching frequency is too high (here shown as ω_0), the slope is unchanged, analogous to the I-V characteristics of non-memristive materials. Figure adapted from [71].

evidence using thin films of TiO_2[71]. This was done by using the memristance, a physical variable belonging to the memristor, similar to resistance, inductance and capacitance. The memristance comes from the relationship between the flux ϕ_m and the charge q, which is given by $d\phi = Mdq$, where M is the memristance. Using $d\phi = Vdt$ and $dq = Idt$, leads to $M = \frac{d\phi/dt}{dq/dt}$ and thus $M = V/I$. For linear elements, the memristor is therefore simply a resistor. When the resistance depends on the charge, however, the element becomes more interesting as this implies tunability. Looking at the I-V characteristics of a regular resistor, the dependence would be linear. For a memristor, however, a hysteresis loop should be found. This means that the charge that has gone through the material before affects the I-V characteristics for a memristor.

In figure 24, it is shown how the first memristor had a non-linear resistance. Once a voltage was applied and a current ran through the system, the resistance — indicated by the slope — changed. This would then also remain changed after the voltage is removed from the circuit. In the experiment, a sinusoidal voltage was applied to a circuit where a semiconductor thin film was put in between two metals. The semiconductor had a region doped with positive ions (p-doped) and a region with no doping at all, as shown in figure 25. The p-doping leads to a high conductivity in the semiconductor, unlike in the undoped region. Application of a voltage caused the dopants to drift, which also changed the resistivity over the circuit [72]. When the voltage is removed from the circuit, position of the dopants remains distinct from the initial position before the application of a voltage. Thus, the new resistance is "remembered", which makes this a type of Non-Volatile Memory (NVM). The dashed line shows a hysteresis collapse for high frequencies, indicating that the dopants will not drift significantly and the memristor acts as a regular resistor. Though the precise definition of a memristor is contentious [73][74][75], there are more similar memristor-like NVM devices that can act as the resistive component in IMC.

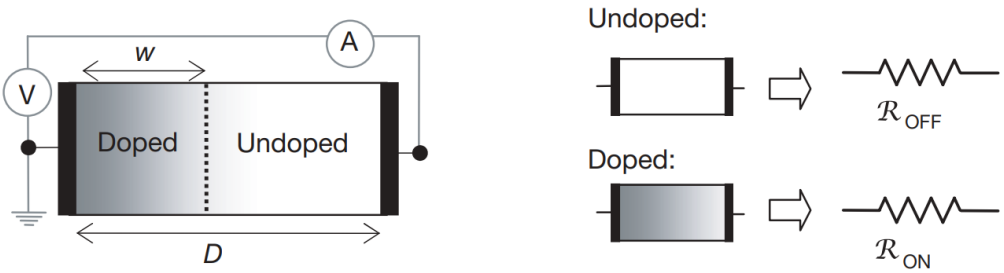

Figure 25: A schematic showing the setup of the first published physical example of a memristor. A region of p-doping and no doping were formed, which had low resistance (R_{on}) and high resistance (R_{off}), respectively. By ion migration, a low resistance filament can be formed. Image taken from [71].

3.2.1 Resistive RAM

Within the AIMC architecture, there are many interesting NVM devices which have their data stored in a variety of ways. These NVM devices include Phase Change Memory (PCM), Resistive Random Access Memory (RRAM), Electrochemical Random Access Memory (EcRAM), Magnetic Random Access Memory (MRAM), flash memory and Ferroelectric Field-Effect Transistors (FeFET) [76].
Firstly, Resistive Random Access Memory (RRAM) devices are generally based on metal-insulator-metal systems where the insulator is a dielectric material, but even semiconductors can be used. Normally the insulator would naturally lead through no current, but in RRAM systems there is formation and dissolving of conductive filaments in the insulator region. These filaments are created by positive ion (cation) or negative ion (anion) migration in the dielectric[77], for example in a system illustrated in figure 25. The conductance in the insulating material is controlled by either tuning the width of the material or changing the composition of the filament in the insulating region [76]. The dielectric material that constitutes this insulating region can be made of binary oxides, perovskites, 2D materials and more. The binary oxides, like the TiO_2 used by HP for their breakthrough RRAM device, are simple, inexpensive and compatible with CMOS processes[78]. The compatibility with CMOS processes is not shared with other possible dielectric materials like perovskites, which gives rise to considerable practical challenges. The durability of the binary oxide HfO_2 ($> 10^{10}$ cycles) is paired with sub ns resistive switching, making that the most interesting binary oxide together with TaO_2, which can reach similar numbers, even reaching an endurance of $> 10^{12}$ cycles [79]. These qualities are especially relevant for neuromorphic computing systems, where the endurance is vital for systems that ideally will not be touched for years. There are three ways of forming filaments in the dielectric material. Namely, it can be done by n-type or p-type doping on the dielectric material, by electrode diffusion or by designing bilayer structures [80][71][78].

Electrochemical dissolution and deposition of electrodes can lead to two possible filament forming processes. Firstly, the cation migration in the dielectric, which is called Electrochemical Metallisation (ECM). This can be done by applying a so called writing voltage to an active electrode such as Cu or Ag, which then oxidises. The resulting metal cations diffuse across the dielectric and bind to the electrode on the far end, which is an inert electrode like Pt and Au [78]. As these cations bind to the far end, more cations will build up until a filament is formed [82]. By applying

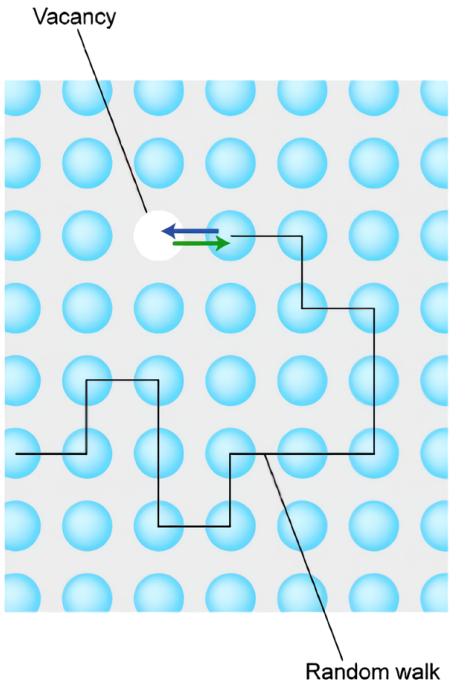

Figure 26: How the oxygen vacancies move in the material. A particle is attracted to enter the vacancy, leaving a new vacancy — which is in essence the same as the vacancy moving to a new site. Image taken from [81].

a much smaller voltage called the read voltage, the resistance can be probed. The voltage that forms the filament to create a Low Resistance State (LRS) is also called the SET voltage, whereas the voltage that switches into the High Resistance State (HRS) is called the RESET voltage [83].

A second way to form filaments is by anion migration, which is referred to as Valence Change Memory (VCM). The process here results from a redox process at one of the electrodes. The function of the electrodes here are swapped compared to the ECM process, meaning the inert metal is the active electrode and a non-noble metal is called the oxidisable electrode [84]. The redox process refers to the loss of electrons (oxidisation) in one reactant, while there is a gain of electrons (reduction) in the other reactant. This redox process happens at the oxidisable electrode's interface with the dielectric material. For example, if a high voltage is applied to the Ti (electrode) and TiO_2 interface, the following chemical reaction takes place:

$$TiO_2 + Ti \longrightarrow TiO_{2-x} + TiO_x \tag{11}$$

The TiO_x is the partially oxidised Ti electrode, which is still conductive [81]. The TiO_{2-x} in the dielectric effectively function as oxygen ion vacancies. These ions can diffuse through the dielectric in a way that is shown in figure 26. Here it is shown how a vacancy in a lattice can move around. In reality the vacancy is filled by the nearest neighbours, but this is identical to saying the vacancy moves to the donating neighbour. As the oxygen vacancy concentration is increased, diffusion by application from a high voltage will lead to the forming of a conductive filament. Both in the case of anion migration and cation migration, there exists a fundamental limit for endurance.

3.2.2 Phase Change Memory

Next to the forming of filaments, it is also possible to create high resistance and low resistance states by transitioning between crystalline and amorphous phases of chalcogenide glasses. The resistivity of the amorphous phase is very high, while the crystalline phase's resistivity is very low [87]. The resistivities of the two states can differ by as much as 3 to 5 orders of magnitude [88]. Joule heating, which is the heating that occurs when a current is run through a material, can be used to switch between the two phases of the material. As is shown in figure 27, if the material is heated up to exceed the crystallising temperature, the crystallisation can occur. Conversely, the material needs to be melted and quenched rapidly, for the material to reach the amorphous phase [89]. Melting a very high resistance state is possible due to threshold switching, where an applied voltage above the threshold voltage V_t can create conductive filaments in the amorphous phase. Through these filaments, Joule heating can be used to melt the material, given a high enough voltage [89]. These melting and quenching steps most commonly results in an amorphous dome that resembles the head of a mushroom, which is called the mushroom cell structure [85], this is shown in figure 28. The endurance has been proven to reach values of $> 10^8$ cycles with a SET speed of 1 ns [90]. However, the programming current is very high since the melting temperature of

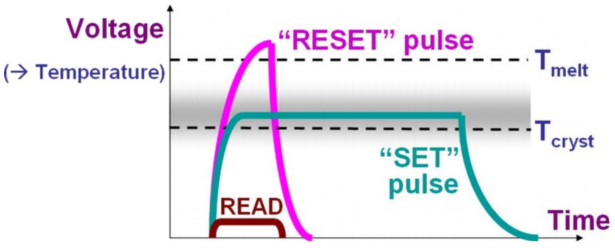

Figure 27: Comparison of the SET, RESET and READ voltage pulses in PCM. To SET a LRS, a longer pulse at high voltage is needed to reach the crystallising temperature due to Joule heating. To set an HRS, a RESET pulse has to be applied for a very short period to rapidly melt and quench so that the material can't relax into the crystal phase and stays in an amorphous phase instead. Finally, a READ voltage can be applied to probe the resistance of the material. Figure taken from [85].

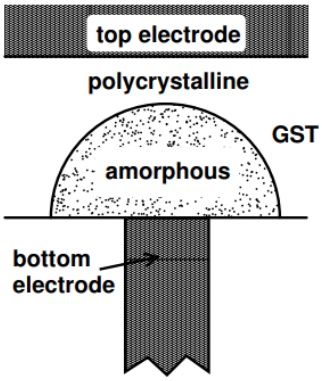

Figure 28: A schematic of the physical structure of a PCM device, shown as the characteristic mushroom cell. When putting the device in the RESET state — thus the amorphous, low resistance state — an amorphous dome appears on top of the bottom electrode. Image taken from [86].

the material needs to be reached. [90] used a crystallisation temperature of 600 K and a melting temperature of 1200 K for the melt-quench operations. Furthermore, unlike RRAM, PCM cannot yet go under 10 nm in size, making this a relatively unattractive NVM device [76][85]. Though, the HERMES core — supported by IBM — has been able to come to a 14 nm size [91].

3.2.3 Alternative Non-Volatile Memory types

While RRAM and PCM are the most important NVM devices in contemporary research, there exist other NVM devices which might become more relevant in the next decades. These are explained briefly here, with illustrations of the corresponding cell structures given in figure 29.

Firstly there is flash memory, which works on the principle where current can run through a material, past a floating gate where charge can be stored. This charge can then increase the resistivity of the gate. The amount of charge stored can be tuned. Flash memory leads to relatively high noise, however, meaning this is for the time being less feasible than PCM and RRAM[87].

Secondly, there is Magnetoresistive RAM (MRAM), where there are two magnetic layers with an insulating layer in between [92]. The magnetisation of one of the two layers are set, the other one is free to change. This can be changed by applying an external field, which then changes the magnetisation. If the two magnetisations are aligned, there is a low-resistance state for a tunnelling current to flow. If the magnetisations are not aligned, there is a high resistance state. Though MRAM has had major improvements over the past decades, scalability issues keep it from being a viable option for NVM devices. To improve on the scalability issues, Spin Transfer Torque MRAM (STT-MRAM) has been introduced [93][94]. Since STT-MRAM deals with spins, the effects are very small, which means that there is a low RHS/LHS ratio of ∼2.

Thirdly, Electrochemical RAM (EcRAM) works on a simple gate-source-drain system. The source and drain is connected by channel through which the current flows. The resistance is modulated by the gate, which is separated from the source-drain channel by a electrolyte layer that enables ion migration. By applying a voltage on the gate, an exchange of ions is induced, which changes the resistance of the device[87][94]. This has the advantages of only using low power for writing and reading processes, symmetric switching, as well as high resistance granularity[87]. However, the slow operation speed and lack of maturity make this NVM device less attractive compared to RRAM and PCM[94].

Fourthly, Ferroelectric RAM (FeRAM) is a NVM device that is similar to a dynamic RAM

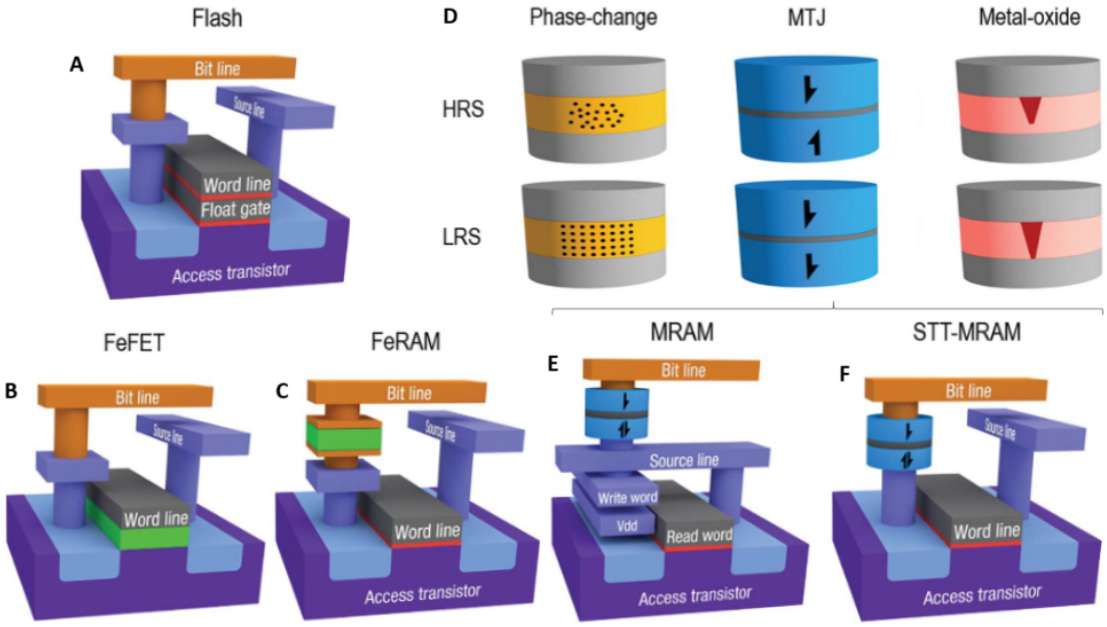

Figure 29: Cell structure of the analysed NVM devices. A-C) The cell structures of flash memory, ferroelectric FET and ferroelectric RAM. D) The nanocells with working principles of phase change, magnetic tunnel junctions (MTJ) and metal-oxides (ReRAM). The amorphous and crystalline structures of phase change memory, (mis)alignment of magnetic moments and the forming of filaments are shown for the HRS and LRS. E-F) The cell structures of the two analysed MRAM devices, which belong to the MTJ shown in D. Image adapted from [92].

Table 2: Comparison of best performances of commercialised stand-alone NVM devices. It is important to note that this is commercialised and stand-alone devices, so that better realised structures — which are not stand-alone and commercialised — are excluded in this list. This for example means that the bits per die for RRAM, which reaches 32 GB [97], is excluded. This would be a large improvement both in memory and possibly in price per bit. 'F' is the minimal lithography size and 'L' is the amount of layers in a 3D structure. Table adapted from [92].

	FeRAM	PCM	RRAM	STT-MRAM
Cell area	6-30 F^2	4/4L F^2	6-30 F^2	6-30 F^2
Bits per die	8 Mb	256 Gb	8 Mb	1 Gb
Retention	>10 years	>10 years	>10 years	>10 years
Endurance	$\sim 10^{15}$ cycles	$\sim 10^{7}$ cycles	$\sim 10^{6}$ cycles	$\sim 10^{15}$ cycles
Read time	10-100 ns	10-100 ns	~100 ns	~10 ns
Write time	10-100 ns	10-100 ns	~100 ns	~10 ns
Cell energy	~0.1 pJ	~10 pJ	~0.1 pJ	~0.1 pJ
Array energy	Low	Medium	Medium	Medium/low
2021 price	> $1000/Gb	≤ $0.30/Gb	~ $1000/Gb	$40-70/Gb

(DRAM) structure, utilising a one-transistor-one-capacitor (1T1C) structure, where the dielectric is a ferroelectric material. Writing a state can be done by applying an electric field, the two possible polarisation states of the ferroelectric material can then be utilised [95]. However, the reading process is destructive in FeRAM systems, so subsequently there needs to be a new write cycle every time. Furthermore, there are limitations regarding scalability due to a low storage capacity, while also having higher costs [92].

Finally, Ferroelectric Field Effect Transistors (FeFET) works by using an FET with a ferroelectric material used as the gate dielectric. By applying an electric field, the polarisation of the ferroelectric material can be altered such that the resistance across the FET is changed[96]. This is similar to flash memory, except with ferroelectricity in the gate. FeFET has reliability issues compared to PCM and RRAM due to parasitic effects like charge trapping, which severely limits the endurance[93]. Furthermore, a relatively short memory retention time limits the capabilities of FeFET as an NVM device[96].

There are even more NVM devices, but these are all either not researched extensively yet or have too many significant limitations to mention. The devices mentioned above, however, are all contenders for future application in IMC. PCM and RRAM are the two main types of NVM devices for the time being, especially PCM, which held 90% of the commercialised NVM market share in 2020 [92]. These other types of devices may also become interesting for ASTRON in the coming decades. Table 2 compares FeRAM, PCM, RRAM and STT-MRAM in their best performances. Here, only stand-alone commercialised devices are taken into account. Certain devices have been found to reach better performance in research, but has been excluded. One notable example is the low number of bits per die given for RRAM, while prototypes have been built that reach results of 16 or 32 GB [97].

3.3 Applicability of Non-Volatile Memory devices

The NVM devices are clear candidates for in-memory computing, which can be implemented in ASTRON's signal processing pipeline. Figure 30 shows a schematic of an IMC crossbar, with memory-based unit cells, similar to what was shown in figure 22. These unit cells have been covered, with the most interesting contemporary NVM devices being PCM or alternatively RRAM based types. In an ideal scenario, these unit cells would have exact resistances, representing the weights that belong to the various matrices. However, there are only an N number of discrete values that the resistance can have. For RRAM, this is because the filament path is limited by the amount of possible direct conductive paths that can be formed. For PCM it is the limited ways that the structures of the phases can be formed. In the best materials there exist hundreds or even thousands of programmable states, giving a modest approximation for being continuously tunable.

There is some level of imprecision by the discrete nature of programmable states. Next to this aforementioned uncertainty of tuning the resistance, there exist 3 types of noise[76]. The first type of noise is programming noise, which is the error of the programmed value compared to the target resistance value. This is due to the write-read-verify iterative loop that is part of programming the

right noises. Secondly, drift noise is the shifting of noise over time. This is either due to relaxation of the amorphous phase in PCM or due to the shifting of the ions in RRAM [98][99]. Thirdly, there is read noise that comes from the flicker noise and random telegraph noise (RTN) sources, which also change the resistance values. This happens when reading the resistances, after the components have been programmed. Besides these known noise types, there also exist noise issues specifically in analogue systems. Namely, the endurance for materials is given in digital voltage pulses. Normally, the MHz range of radio astronomy would give ample time for the materials to recover from possible drift in the time frame when a voltage is applied. Since the analogue signals spend essentially no time without an applied voltage, the endurance is unknown for these materials in such an environment.

Before these NVM devices can be applied in ASTRON's processing pipelines, research needs to be done to figure out the effects from an analogue signal into a system like LOFAR's. While direct implementation of an analogue input for IMC has been presented before [100][101], the questions remain what the effect is of implementation of a 24/7 applied signal that intends to run for decades. Finally, the read voltages need to be taken into consideration. ASTRON uses amplifiers on their incoming signals, otherwise the desired signal is too weak. However, read voltages of 0.2V may significantly increase noise levels in RRAM systems [102][103]. Either ASTRON needs their incoming voltages to stay at 0.2V at its maximum after the amplifications or other materials need to be used to allow for higher read voltages. This would make PCM devices or even Flash memory more attractive, whose writing voltages tend to be very high, though the voltage levels scale with size, so the size has to be significantly larger to raise the possible read voltages [85][104]. Currently, the threshold voltages at which Joule heating is achieved in PCM is around 1V, so the voltage has to stay at a significantly lower level to prevent any changing of the material. Having a constant current will also cause minor heating, so some cooling mechanisms like fans have to be used so that the material does not start overheating. If all the combined noises become a precision issue for ASTRON's processing purposes, there is a way to achieve arbitrary precision [105]. This will be further analysed in Section 4.4.

Having considered various NVM devices and their applicabilities, the next chapters will cover if ASTRON's front-end processing pipeline could be done by using the analogue and/or neuromorphic domains. Furthermore, the effect on energy efficiency and latency will be considered in these domains.

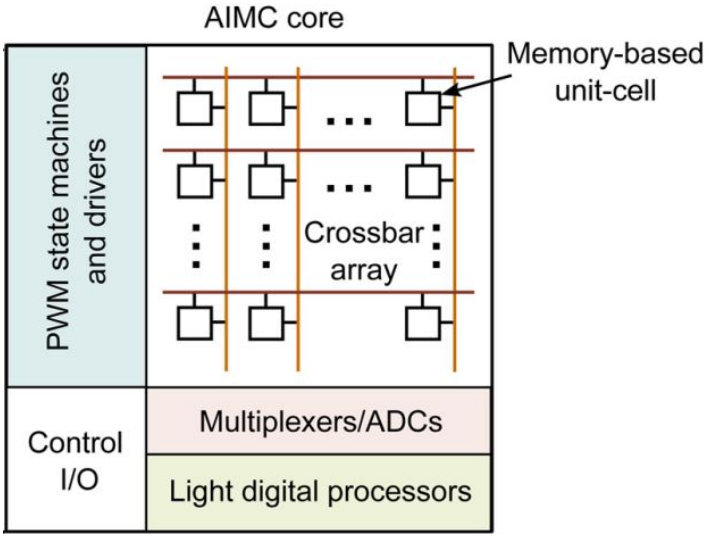

Figure 30: A schematic of an AIMC core. This includes a crossbar, with memory-based unit cells, similar to what was shown in figure 22. Figure adapted from [76].

4 Results

The shift of the PPF and beamformer from the digital domain toward a more energy efficient domain will be analysed in this section, separated into two domains for each of the components: analogue and neuromorphic. Afterwards an energy comparison will be made with the current pipeline. Finally, the precision of the outgoing signals will be discussed.

4.1 Analogue LOFAR pipeline

In this subsection, there is a focus on a shift towards the analogue domain. This includes if the LOFAR pipeline can be constructed with analogue components and if so, how.

4.1.1 Analogue FIR

As stated in Section 2.3, the FIR filter consists of a tapped delay line, which is followed by multiplicative weights whose results are then summed together. The multiplications can be done by amplifiers, capacitances, resistances and multiplicative DACs, among other things [106][107][108]. Additions are similarly easily implemented by use of an op-amp [106]. Analogue delay lines have also commonly been used, for example for music-making purposes[109]. Due to the simple nature of the FIR components, it is unsurprising that first publications of analogue FIR filters date back to the 1960s and 1970s [106][110]. As a result, the implementation of such an FIR filter in the desired analogue LOFAR system should not pose a problem.

An example of the analogue FIR system is schematically shown in figure 31a. This shows the implementation of the tapped delay line to separate the samples of the signal [106]. In this example, 12 samples were put through weights which were implemented by the use of resistances which then fed into an op-amp which summed the weighted signal taps, as shown in figure 31b. The resistances can then tune the voltages such that multiplications can be read out. The way this is done is via resistive weights using so-called Savitzky-Golay functions. The reciprocal of the absolute value of these weights can then be the values of the resistor. For example, if a minimum resistance of 10 $k\Omega$ is needed and a weight factor's maximum value is 15, then the resistor corresponding to that value can be chosen as 10 $k\Omega$. If then there is a weight value of -2, the resistance would be $|\frac{1}{-2}| \cdot 15 \cdot 10 = 75k\Omega$. The factor of 10 is due to the minimum resistance value and 15 because that is the weight factor for which this resistance reference value exists. As stated above, a sufficient number of different weight implementations exist which can be used in the analogue LOFAR system.

The power consumption of an energy efficient 32-tap FIR filter was found to be $90\mu W$ for a sampling rate of 16 MHz [111]. Within LOFAR, there are only 16 taps. If the assumption is made that the power consumption scales linearly with the number of taps, then the power consumption for a 16-point FIR in LOFAR will be $45\mu W$. The sampling rate at LOFAR is 160 MHz or 200 MHz, which is a full order of magnitude higher than what is shown here. This means that it may be more costly to implement this architecture in LOFAR, if such an increase to the sampling rate is possible at all. This architecture also makes use of high reconfigurability, which is not necessary for LOFAR. Not using this reconfigurability can further optimise the power costs of the FIR filter.

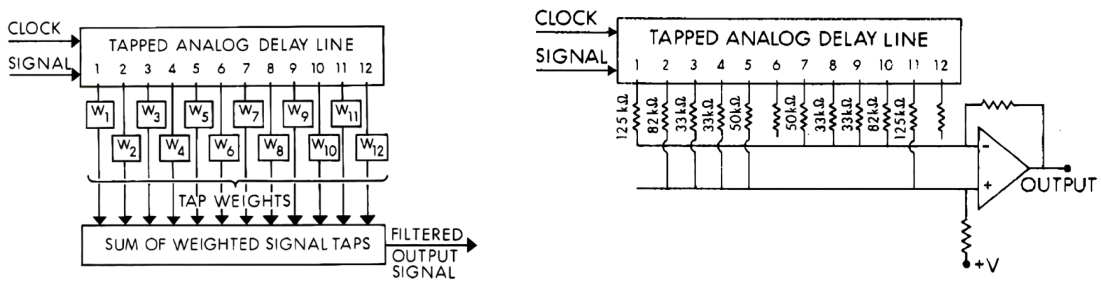

(a) (b)

Figure 31: a) Schematic of an analogue FIR filter. First, a tapped delay line is shown that then feeds into multiplicative weights before being summed together. b) A schematic of how the FIR filter is built. Resistances are used to act as weights and the summing of weighted inputs is done by an op-amp. Figures taken from [106].

4.1.2 Analogue FFT

Unlike the FIR filter, FFT filters are complex to implement in the analogue domain, though it is possible. One such example of a system that implements an analogue FFT is shown in figure 32 [112]. Here, the size-8 FFT is shown, including the peripheral components. Firstly, the incoming signal is serial so it needs to be split up into parallel components for the input to the FFT. For this purpose, sample-and-hold amplifiers (SHAs) are used that are controlled with use of a clock. The first sample is taken at time t_1, then t_2 up to t_N. Once all samples are aligned at t_{N+1}, they are put into the FFT. Within LOFAR, the incoming signal is parallel coming from the FIR filters, so the SHAs would be superfluous. Similarly, the equaliser and final SHAs are irrelevant for the analogue LOFAR system and the coefficient DACs alter the butterfly coefficients, which is unnecessary for the purposes in LOFAR as well. However, the FFT signal flow graph (SFG) is important to analyse. This consists of butterfly diagrams in a radix-2 DIT structure, which includes additions and multiplications.

Once more, a naive approach would be utilisation of amplifiers and op-amps, like with the FIR filter. However, some things have to be taken into consideration for the FFT filter. Firstly, it is critical to note that the resulting noise, energy consumption and chip size need to be limited. This is true in general, but the scale of the FFTs in LOFAR and the resulting amount of operations increase the importance here. Secondly, the FFT's twiddle factors mean that the signal is multiplied with complex numbers. To realise this in electrical engineering, it is necessary to separately multiply the the real and imaginary parts. One example of the butterfly structure is given by [112], who makes use of analogue transconductance multipliers and linear adders, which are shown in figures 43 and 44 in the Appendix (App. A). Furthermore, the butterfly diagram operations can be realised in other ways as well, such as by the use of current mirrors [113]. Furthermore, [114] uses capacitive voltage dividers and wire swapping for additive and multiplicative functions, respectively.

Figure 33 shows the set-up of the butterfly structures shown in figure 32 [115]. X_A, X_B are the incoming voltages in the butterfly diagram, Y_A, Y_B are the outgoing voltages. I and Q differential voltages represent the complex values of the signal. Transconductances G_m and linear adders Σ are shown for multiplicative and additive purposes. The transconductances change the magnitude of the signal with the coefficients C_0, C_c, C_s and converts it into a current signal. The adders include a transresistance that converts the signal back to a voltage signal. Phase shifters (PS) are used at

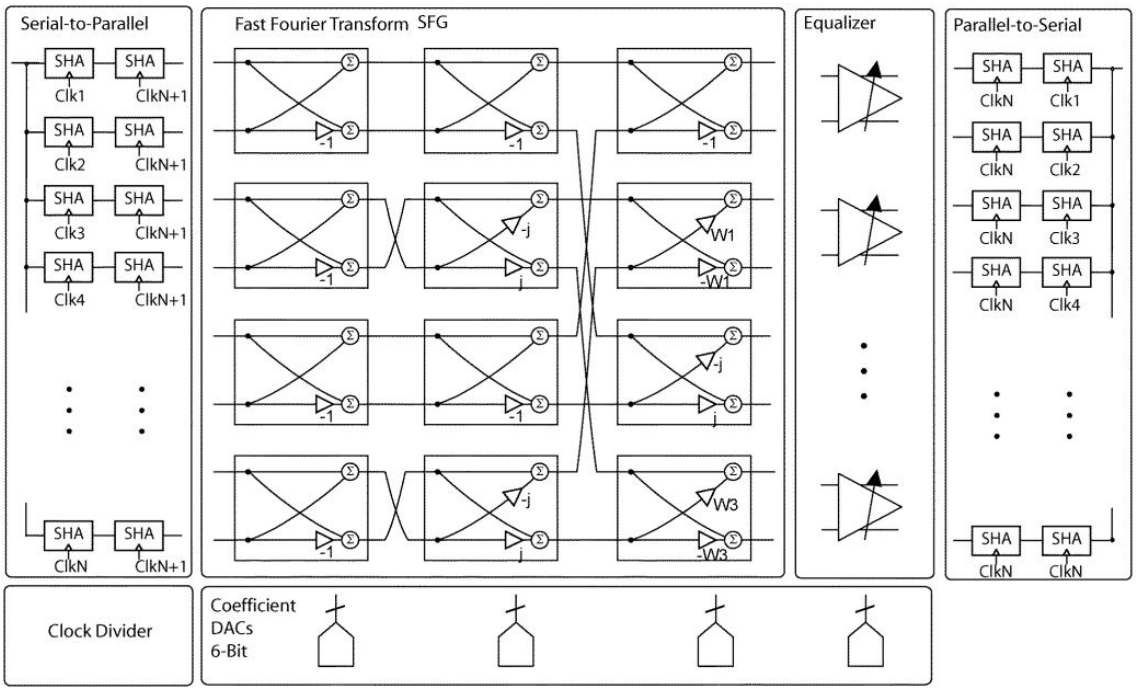

Figure 32: A circuit design of a CMOS FFT processor. In this design, the incoming signal is — unlike in the analysed LOFAR pipeline — not yet parallel. For parallelisation, SHAs are used combined with a clock divider. The same steps are used to make the outgoing signal serial again. In between, the FFT SFG is shown with butterfly structures. These butterfly structures are shown in more detail in figure 33. Figure taken from [112].

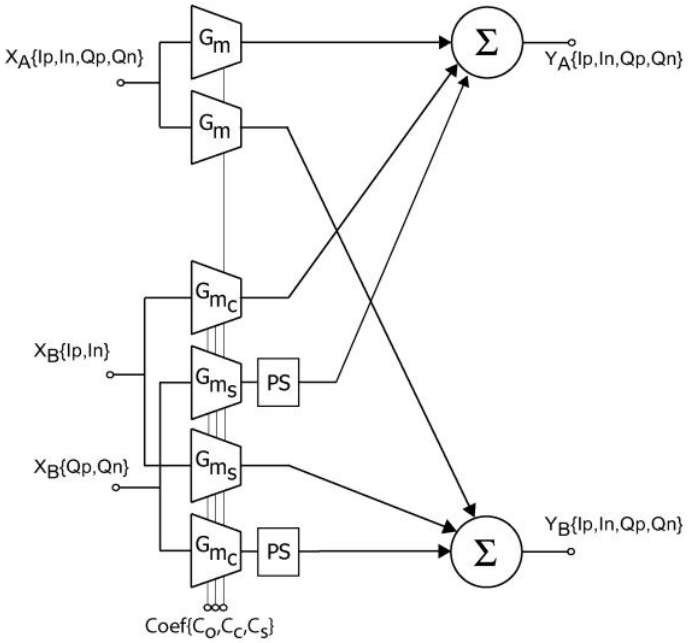

Figure 33: Schematic showing the set-up of the the butterfly structures from figure 32. Transconductances G_m and current adders Σ are used for multiplicative and additive operations, which are more thoroughly analysed in figures 43 and 44 in App. A. PS indicates an applied phase shift, tuned to 90°, 180°, 270° or 0°. Figure taken from [115].

$90°, 180°, 270°$ or 0° angles, which is implemented as an equivalent multiplication by j, -1, -j and 1 respectively. Since the twiddle factors for FFTs of N≥8 have intermediate values that are not equal to 90° phase shifts, these PS components are not sufficient to fill the purpose of the multiplicative twiddle factors. Consequently, the transconductance values G_{mc} and G_{ms} are introduced. The c and s subscripts represent the cos and sin terms to multiply with G_{m0} to represent the real and imaginary values, respectively. Then:

$$G_{mc} = G_{m0} \cdot cos(\frac{2k\pi}{N_{FFT}})$$
$$G_{ms} = G_{m0} \cdot sin(\frac{2k\pi}{N_{FFT}})$$

It must also be stated that the inputs for the FFT come from one antenna, which holds true for every FFT. Therefore, the inputs to the FFT are real valued and consequently, the FFT outputs are symmetrically valued. As a result, half of the results can be ignored as the outputs do not need to be counted twice.

[112] mentions an 11-25 mW power consumption for the 8-point analogue FFT. [114] built upon the above architecture, creating a 256 point FFT with 614.4 mW power consumption for a sample rate at 2 GS/s, which is slightly higher than the maximum LOFAR throughput of 3 Gb/s or 0.75 GS/s [116]. However, since the reference power consumption is based on complex incoming values, the power consumption might be slightly overestimated for comparison to the LOFAR system. The SNDR for the 256 point FFT is 38 dB, which is similar to the 36 dB found in the original paper by Lehne et al [112]. Within the LOFAR pipeline 512-point FFTs are needed, which would amount to a power consumption on the order of 1 W for each FFT.

4.1.3 Analogue Beamformers

The analogue beamformer is a component that already has a use case within the receiver signal processing pipeline in LOFAR. Therefore, it is more or less trivial that it is indeed possible to implement analogue beamformers at the end of this part of the pipeline as well. One example of a beamforming scheme is shown in figure 34 [117]. Here, a delay and sum architecture is used by the use of capacitors C and transistors M. The capacitors act as memory storage. This scheme was used to realise a 16 channel analogue beamformer and it was able to reach a power consumption of 16.2mW. This is significantly smaller than the number of channels needed in the LOFAR pipeline,

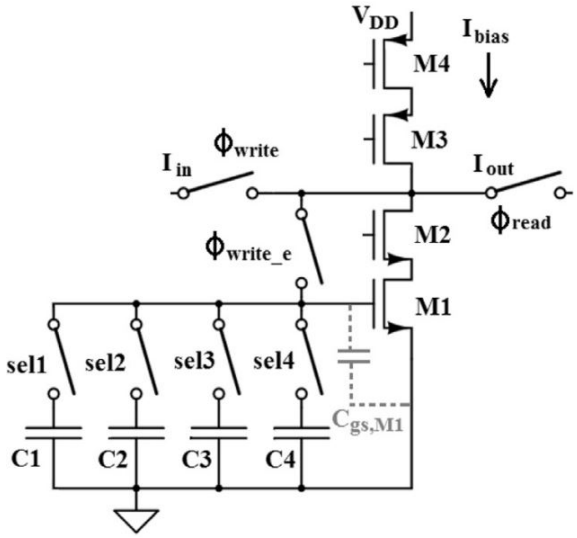

Figure 34: Scheme for an analogue beamformer where C is a capacitor and M is a core transistor. Figure taken from [117].

but research for large beamformers of around 144 channels is not readily available. If the assumption is made that power scales linearly with the number of channels, the total power consumption for one beamformer in LOFAR would be 145.8mW.

Figure 35 shows a schematic overview that indicates what the LOFAR receiver signal processing pipeline looks like with the implementation of the aforementioned analogue PPF and beamformer components. There are 144 receivers in a Dutch LOFAR station, each with 512 FIR filters, which then give their output into a size-512 FFT filter. The 512 outputs of each FFT lead to a different beamformer, those beamformers acquire one input from each receiver. The multiplicative factors, additive factors, phase shifters, etc. in this figure can be implemented in ways described in the Sections 4.1.1, 4.1.2 and 4.1.3.

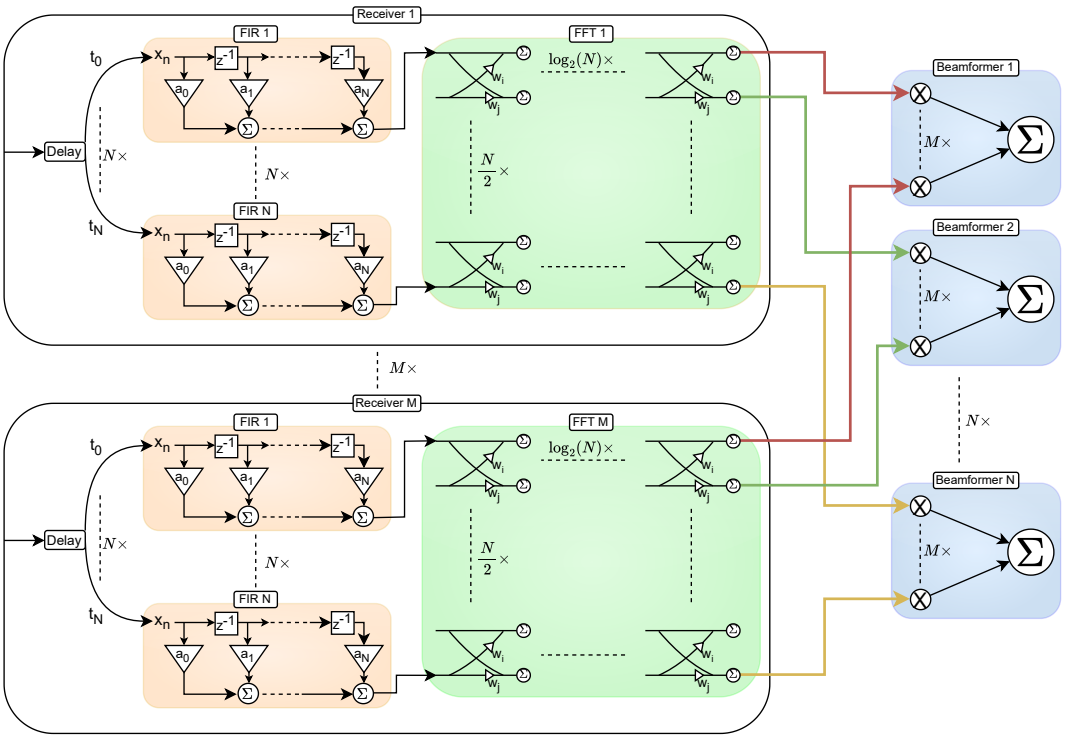

Figure 35: Pipeline schematic for the analogue LOFAR receiver signal processing pipeline. The FIR filters, FFTs and beamformers can be implemented by use of components exemplified in Sections 4.1.1, 4.1.2 and 4.1.3.

4.2 Neuromorphic LOFAR pipeline

Having seen the possible implementations of analogue components in the LOFAR receiver signal processing pipeline, further analysis will be done in this subsection on the application of neuromorphic components.

4.2.1 Neuromorphic FIR

As seen in figure 22, with the use of memristors, a crossbar can be made to replicate MVM calculations. In the case of the FIR filter, the calculations are simpler due to the simple nature of the device. There is only one multiplication for each signal coming out of the delay line. This means that there is only a vector-vector multiplication (VVM), which is relatively short as well, due to the fact that there are only 16 inputs in the FIR filters in LOFAR. As such, it is relatively trivial to make a crossbar that satisfies the conditions for the neuromorphic FIR filter. While it has never been shown before that this can be done well, the assumption here is that a FIR filter will be possible if a more complex system can also be made by use of a crossbar. One such example is the Fourier transform explained below (Section 4.2.2). Since the FIR filter has a delay line next to the VVM, it is still necessary to implement such a delay line before inputting the crossbar. The resulting schematic is shown in figure 36.

The energy consumption of the crossbar will be analysed in Section 4.2.2, but the FIR filter also has the delay line component before it. [118] created an analogue delay line of 20 taps with a power consumption of $90\mu W$ at sampling rate 40 MHz. This entails a slightly higher number of taps than what is used in LOFAR, but it is still decently comparable. If the rough assumption is made that power consumption scales linearly with the number of taps, the power consumption for a 16-tap delay line is equal to $72\mu W$.

4.2.2 Neuromorphic FT

The DFT has been used in many systems over the whole world, like in signal processing. For this purpose, the DFT has been widely analysed to figure out any ways to make it as energy efficient as possible. Besides the aforementioned FFT algorithm and the step to do this in an analogue CMOS system, the DFT also lends itself to be used for a neuromorphic use case. Since the FFT cannot be considered as a singular MVM, but rather as a multiplication of many smaller matrix-matrix multiplications, the energy gain from a faster algorithm has to be discarded for the benefit of using a memristor crossbar. This has been widely considered and later implemented[120][121].

Implementation of low power MVM crossbars has gotten a lot of attention, including in condensed matter physics and particle physics [15], and due to the Fourier transform's importance the DFT has also gotten widespread attention [16]. While the above examples analyse exact solutions to the MVM by switching between the analogue and digital domains, examples also exist analysing

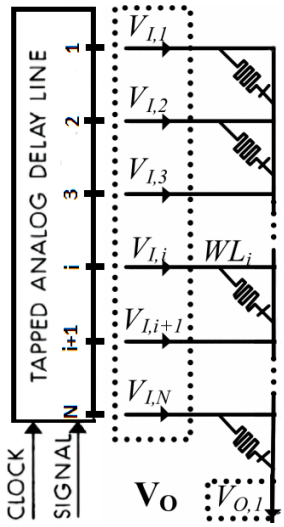

Figure 36: Schematic showing the setup of the neuromorphic FIR. First, a tapped analogue delay line is used, then the taps feed into a 16 x 1 vector crossbar. Figure adapted from [119] [106].

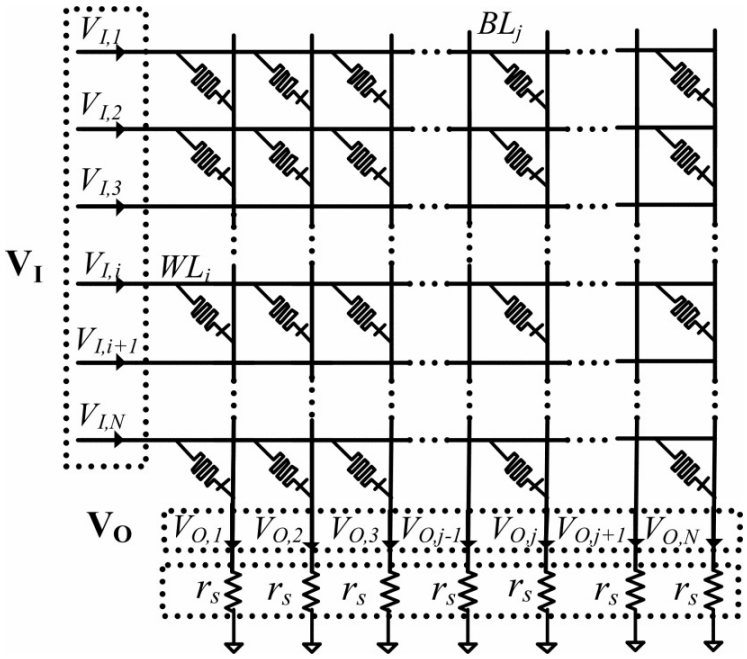

Figure 37: Schematic showing the setup of a crossbar, for example for the implementation of the DFT. The incoming voltages on the left V_i move across the word lines WL_i following Kirchhoff's laws at each junction with the bit lines BL_j. At each memristive component, the voltage is multiplied with the conductance following Ohm's law. The resulting current values across each column get added up to the output values V_0. Figure taken from [119]

a full-analogue crossbar's performance [122][119]. One such example of an analogue crossbar that is utilised for doing DFT calculations is shown in figure 37 [121]. The horizontal word lines WL feed a voltage V_i to the memristors at the junction with every vertical bit line BL. As explained in Section 3.1, through the use of Ohm's law and Kirchhoff's laws, the voltage values are multiplied with the memristors' conductance values to obtain current values, which within the BLs are summed to form the outputs of the MVM. The outputs can be converted back to voltage values at the end of the BL with the use of simple op-amp circuits [123]. For a DFT, the conductance values of the memristors are equal to the twiddle factors found in the DFT matrix of figure 9. The implementation of complex valued weights is non-trivial, as well as the addition of negative weights.

The DFT weights W and outputs X are all complex values, while the crossbar can only compute real multiplications and additions. It is therefore necessary to divide the signals into real and imaginary parts, as shown in figure 38. There are four groups of memristor crossbars, two for the real part of the incoming signal Re(x) and two for the imaginary part of the incoming signal Im(x). Each one of those groups has a crossbar representing the positive (W^+) and one representing the negative weights (W^-). Since LOFAR FT inputs are real valued, half of these crossbars can be removed. Just as before (Section 4.1.2), the outputs are symmetrical so half of them can can be ignored once more. This leaves 2 crossbars, where the real inputs x are multiplied with the real parts of the weights. Since one crossbar sums the negative values and the other sums positive values, the two crossbars still need to be added together. The values coming out of the negatively valued crossbar are positive, so subtraction of the two crossbars is necessary. This means that column k of the first crossbar gets subtracted with column k of the other crossbar before being sent on to the beamformer, for all columns k. As the weights are either positive or negative, if there are an M number of weights in a crossbar, there will be M zeroes across the two crossbars. This would mean an infinite resistance, but that is impossible. Instead, the memristor should hold the highest possible resistance so that the conductance approximates 0. Alternatively, it can be possible to put any material with very high resistance (e.g. an insulator) in the memristor's place, as these components do not need to be tuned, which can cut down on costs both in terms of resource use and capital expense. It may also be possible to make the negatively weighted crossbar one row and column smaller, since the first row and columns only consist of positive weights. This may not be a relevant way to reduce complexity or energy consumption, but in the eye of unnecessary resource use, this may be interesting.

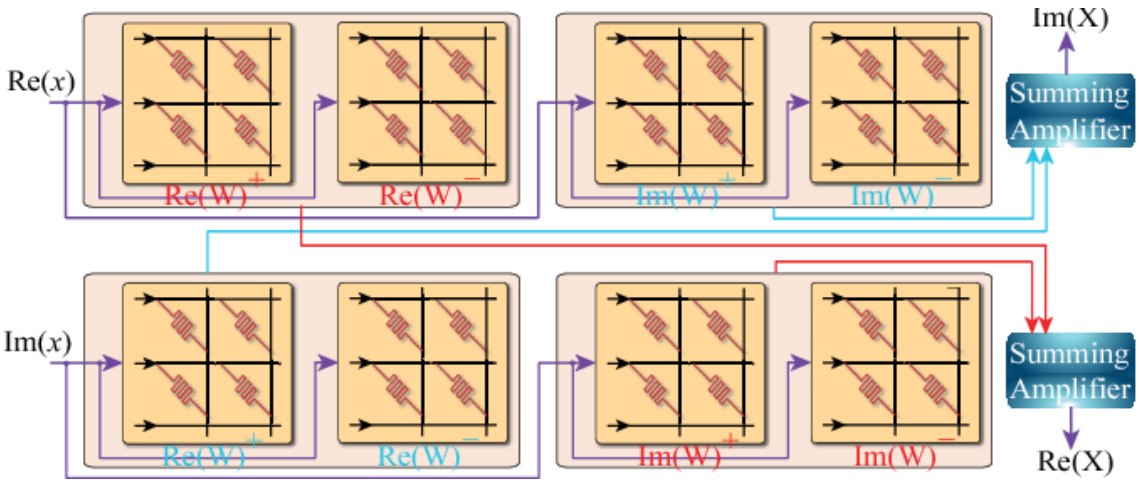

Figure 38: The setup for DFT crossbars. Since the DFT deals with complex numbers, the setup has to be changed to allow for separate real and imaginary inputs, weights and outputs. Since conductance values can only take positive real values, there are 8 necessary crossbars to compute the whole DFT. Since the incoming signal in LOFAR is real and only the real part of the output is relevant, two crossbars are enough for the DFT in LOFAR. Figure taken from [121].

This crossbar system has been implemented for 512 x 512 DFTs, with energy efficiency measurements showing a 55x improvement and latency measurements showing a 5.6x improvement compared to a CMOS FFT module [121][119]. Compared to the earlier found data for analogue FFTs, this is even a roughly 1000x improvement in power consumption. The power consumption is set at 1.12mW, the data throughput is not mentioned. However, these calculations assume the utilisation of all 8 crossbars, which are 6 more than necessary in the LOFAR pipeline. Therefore, the rough assumption can be made of a 4x further decrease in power consumption to reach 0.28mW. Measurements were also done on how the accuracy was affected by process variations in the memristive crossbar. Ideally, following Kirchhoff's laws, the output of the system should be perfectly accurate, assuming the conductance values correspond to their theoretical values perfectly. However, the process variations mean memristors have imperfect values after manufacturing, so programming the weight values into the memristor will be inaccurate [124]. Hence, inaccuracy of the result of a DFT memristor was calculated for different intensities of process variations. It has been done via a uniform variation distribution, with the maximum variation at 5% and 10%. For the maximum variation of 5%, it was found that the relative inaccuracy of the 512 x 512 DFT memristor was 0.18% and for a maximum variation of 10%, the relative inaccuracy was found to be 0.3%. However, as the data is not openly available and precise details of how it was calculated are not given, it is uncertain how accurate these numbers are precisely.

4.2.3 Neuromorphic Beamformer

Having seen the implementation of the neuromorphic DFT, a look can be shed on the relatively simple beamformer. By the separation of the real and imaginary parts of the DFT, the imaginary part of the output can be ignored. As such, the differences in phase for each of the DFT outputs are gone and thus there is no longer a need for a beamformer. This step can be completely forgone. However, for completeness' sake, the implementation of such a neuromorphic beamformer will still be shown. This may especially be interesting for possible use cases where the incoming signal for the DFT is complex, for example.

As seen in Section 2.5 (equation 9), the beamformer can also be considered as a VVM with the incoming vector being values from the DFT and the weight vector consisting of complex exponentials. Once more, a setup like figure 38 is necessary. Depending on the type of input (real or complex), some crossbars may be removed. Since this is a VVM, this amounts to a maximum of 8 column vectors. Figure 39 shows what one of these crossbars would look like. If there are other use cases, whether that is at LOFAR or elsewhere, it is easily implementable and its power consumption will be very low due to the fact that these crossbars are much smaller than the low-power DFTs. As stated above however, there is no need for these beamformers and thus beamformers produce no power consumption at all.

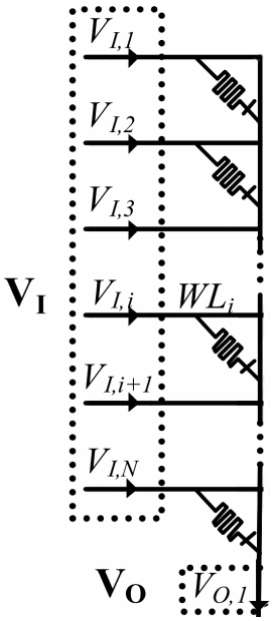

Figure 39: A schematic showing a possible way to implement a neuromorphic beamformer. While the weights of the beamformer are complex exponentials, this figure simply shows one vector-like crossbar. Figure 38 shows how it is possible to implement complex weights. For LOFAR's receiver end processing purposes, this beamforming crossbar is superfluous. Figure adapted from [119]

Figure 40 shows a schematic overview that indicates what the LOFAR receiver signal processing pipeline looks like with the implementation of the neuromorphic components. There are 144 receivers in a Dutch LOFAR station, each with 512 FIR filters, which then give their output into a size-512 DFT filter. The outputs of the DFTs do not lead to beamformers, as that is made redundant by how the system is designed. Apart from the delay lines, all components are made by use of memristive crossbars.

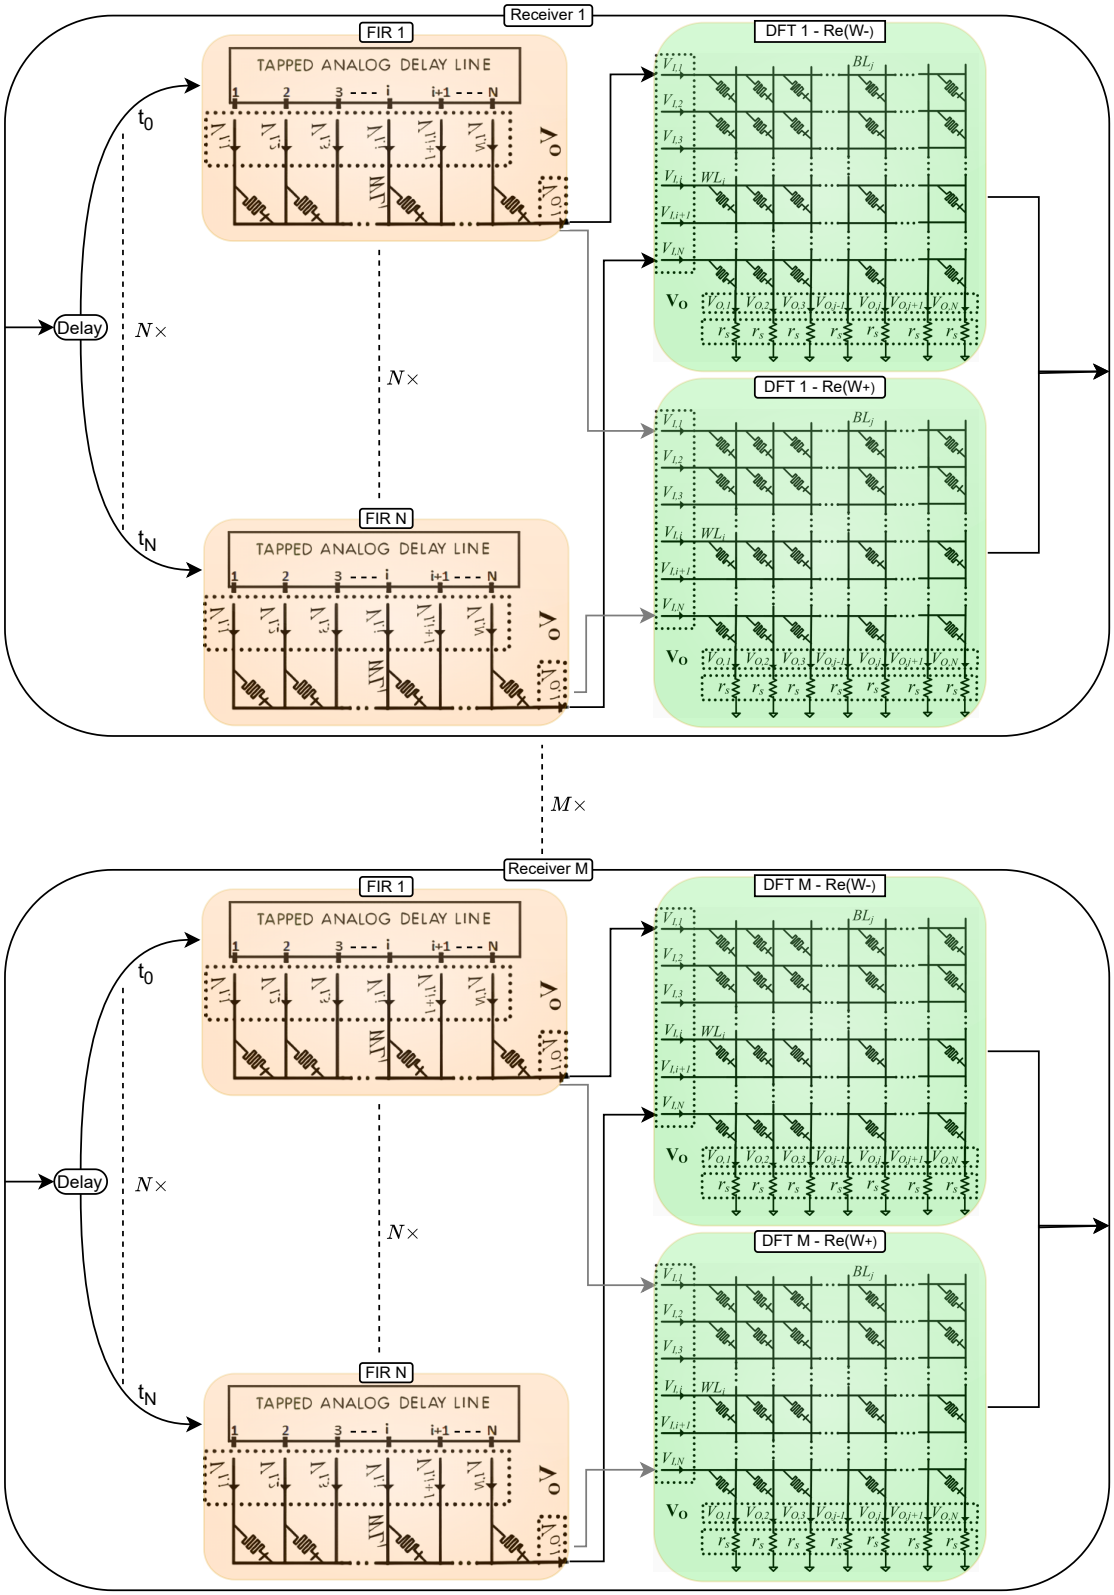

Figure 40: Schematic of the LOFAR receiver-end signal processing pipeline in the neuromorphic paradigm.

4.3 Energy calculations

Currently, each of LOFAR's uniboards use 4 FPGA units, each of which uses a maximum of 350W and an estimated 250W on average[125]. There are 6 uniboards in each Dutch LOFAR station, leading to an estimated power consumption of 6000 W.

To compare with the current FPGA setup at LOFAR, power consumption extrapolation was done from the data collected in Section 4.2.2. This was done by calculating the amount of operations done in total in a station and multiplying that with the power consumption of the neuromorphic DFT. The DFTs are the easiest, as there are 144 receivers in a Dutch station and thus 144 DFTs, with 1 DFT estimated to have a power consumption of 0.28 mW. Then simply, there is a total DFT power consumption of $144 \cdot 0.28 = 40.32$ mW.

For the FIR filters, this is a bit more complex. Namely, the shape of the crossbars are very different, there are no zeroes and the delay line also needs to be taken into account. There are 512 FIR filters per receiver, with 144 receivers and a 16 x 1 column vector size. If the total amount of operations are compared, it is found to be equal to 4.5 crossbars with size 512 x 512. The DFT uses 2 crossbars that are size 512 x 512, so all FIR filter crossbars combined have a power consumption of $2.25 \cdot 0.28 = 0.63$ mW.

As stated in Section 4.2.1, there is also an approximate 72 μW for every FIR filter delay line. With a total amount of $512 \cdot 144 = 73,728$ FIR filters, this amounts to a power consumption of 5.3 W. The total the power consumption of the neuromorphic LOFAR pipeline is still 5.3W, because the crossbars consume an insignificant amount of power compared to the delay lines. However, this is a very high number, especially considering the fact that the analogue FIR filter from Section 4.1.1 consumed $45\mu W$, which includes a 16-tap delay line, but also the weights and summations that come after. As such, the assumption is that this delay line power consumption is overestimated or outdated.

If the comparison is made with the regular analogue system, similar calculations can be made to find the total power consumption in a Dutch LOFAR station by use of the data found in Section 4.1. Firstly, the estimated power consumption of an analogue FIR filter is $45\mu W$. A Dutch LOFAR station has $144 \cdot 512 = 73,728$ FIR filters, so the total power consumption is equal to 3.3 W. This is already energetically cheaper than the neuromorphic FIR filter delay lines, which supports the suspicion that the delay line power consumption is overestimated.

The analogue FFT power consumption has been estimated to be circa 1 W. With 144 FFTs in a station, this leads to 144 W for analogue FFT power consumption.

One beamformer in LOFAR has been estimated to be 145.8 mW. If extrapolated to a full Dutch LOFAR station, this means that the total analogue beamformer power consumption is equal to 21 W.

Table 3 shows a full overview of power consumption, including the current FPGA architecture, the regular analogue setup and the neuromorphic architecture.

Since the neuromorphic power consumption is a rough estimate, another kind of calculation is also considered. This is a direct calculation done by extrapolating data from the HERMES core [91]. This was previously used in literature for calculations with a 512 x 512 crossbar in an architecture going from the digital domain to the analogue domain [15]. The found power consumption was 0.13W with a 1 GHz operation frequency. This includes all the periphery, like a DAC and ADC. As such, the power consumption of the crossbar itself is even lower. By assuming

Table 3: The energy comparison of the current LOFAR setup, the analogue alternative and the neuromorphic alternative. All data is for a full station of 144 antennas, which is accurate for all Dutch stations. The total power consumption of neuromorphic processing is shown as a range of values. The lower bound of the neuromorphic total is calculated without delay lines and the upper bound is with the delay line power consumption. The value of the neuromorphic beamformer is 0 W, because in the neuromorphic architecture, there is no need for a beamformer.

Pipeline element	Current FPGA setup	Analogue	Neuromorphic
FIR	N/A	3.3 W	5.3 W
FIR excluding delay lines	N/A	N/A	0.63 mW
FT	N/A	144 W	40.32 mW
BF	N/A	21 W	0 W
Total	6000 W	168 W	40.95 mW - 5.3 W

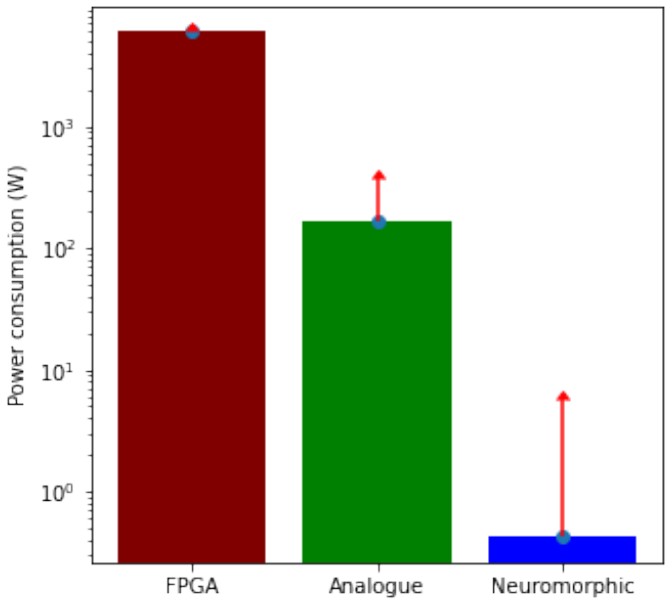

Figure 41: The energy comparison between the current FPGA setup, the analogue alternative and the proposed neuromorphic paradigm. Analogue implementation is circa an order of magnitude lower in energy intensity and a power consumption decrease 4 ± 1 orders of magnitude is possible for the neuromorphic system. There is a large uncertainty in the latter system, with the upper limit chosen at 5.3, corresponding to the neuromorphic system that includes the delay lines at the reported energy values.

power consumption goes up linearly as data throughput is increased and correcting for the ADC and DAC, the extrapolated power consumption in LOFAR would be equal to $96.53mW$ for the crossbars. This means that the crossbars in this rough second estimate would be 2.3x more energy intensive than the previously found numbers excluding delay lines. It is more uncertain, but the power consumption of the crossbars are still much lower than the values found for the analogue delay lines.

Figure 41 shows the energy comparison between the current FPGA setup, the analogue setup and the proposed neuromorphic setup. The analogue system has a power consumption that is one order of magnitude lower than the FPGA setup. In the neuromorphic setup, it can be seen that there is a reduction of 4 ± 1 orders of magnitude in energy consumption in the neuromorphic domain compared to the current FPGA setup. There is a high uncertainty attributed to the delay lines, leading to an upper limit of 5.3 W. The uncertainty shown for the analogue case is based on the fact that the main contributor to the power consumption — the FFT — has data retrieved from literature with the FT in a similar environment to the one required in LOFAR, hence it is not highly uncertain.

Next to the energy consumption of neuromorphic hardware, the size and amount of chips are also relevant factors for the implementation of neuromorphic hardware. IBM has presented their state-of-the-art IMC 256 x 256 core, fitting 64 of them onto a chip with a $4\mu m$ pitch [91][126]. If the same materials were used, with the same chip size, the number of chips needed for a Dutch LOFAR station can be extrapolated from the sizes and amounts of crossbars that are needed. Since DFTs consist of two 512 x 512 crossbars and there being 144 DFTs in a Dutch LOFAR station, that equates to 18 chips. There are also $512 \cdot 144$ FIR filters of size 16 x 1, which — in total size — equates to 18 crossbars of size 256 x 256. In other words, all FIR filters combined would take up 28% of a chip. In total, using the IBM chip as the example, LOFAR may need 19 chips in total.

4.4 Arbitrary precision

Analogue — and by extension, neuromorphic — circuits generally have difficulties regarding accuracy and precision. On top of that come the NVM devices' precision issues presented in Section 3.3, leading to an overall question whether the proposed architecture has sufficient accuracy. However, a circuit architecture with combined ADCs has been presented that allowed analogue devices to do high-precision computing [105]. This architecture is shown in figure 42a. Here, crossbars are shown

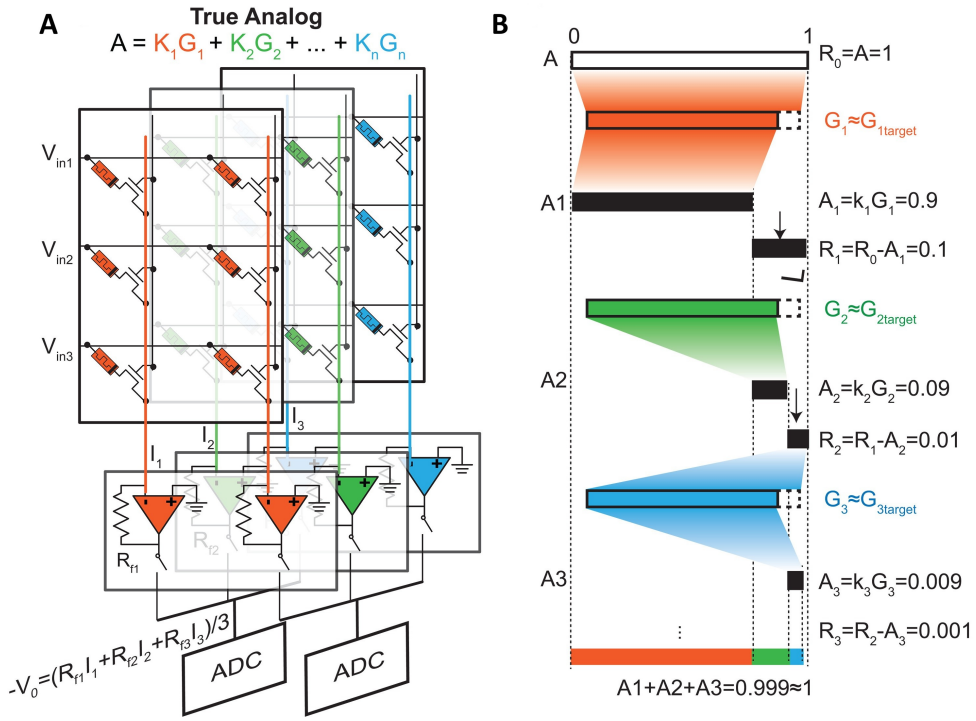

Figure 42: a) Arbitrary precision circuit for neuromorphic systems with joint ADCs. b) Example of programming a desired output A=1 with weighted parallel crossbars A_i. Figure taken from [105].

that work the same as previously explained, but now they are laid parallel to multiple identical crossbars. These parallel crossbars compensate for errors coming from the initial crossbar. One simple example can be given for when the output of a column is supposed to be 1, shown in figure 42b as $R_0 = A = 1$. The initial crossbar is then programmed to have conductance values that lead to an output of 1, but an error of 10% means that the output is actually 0.9 instead. To compensate, a second crossbar is added that once more is programmed to lead to an output of 1. To not overshoot the target value when the two crossbar outputs are summed, the resistance R_f is altered so the second crossbar is weighted by a factor of 0.1. Once again, the second crossbar may have a 10% error, leading to an output value of 0.9 and a total output of 0.99. Then this cycle can be repeated with a third crossbar and even more if necessary, up to an arbitrary precision. It does not matter if a crossbar's error leads to an output value higher than 1, since that can also be corrected by subtracting it with the next crossbar's output. Though this is a simplified example, an extensive algorithm is given following this principle for any complex crossbar, leading to arbitrary precision for neuromorphic computations.

5 Discussion

This study aimed to find whether a shift towards a neuromorphic domain would be a feasible solution to significantly reduce the energy output. The results have shown a decrease in power consumption of 4 ± 1 orders of magnitude for the implementation of neuromorphic front end processing. Next to the significant energy mitigation, there is consequently less strain on cooling mechanisms. This is a first order estimate, however, so there is a significant uncertainty in this calculation. Mainly, the power consumption from the delay lines was imprecise, while it in the most extreme case would account for more than 99% of the total system. Due to the variance in power values, more research would have to be done to more accurately define the system's power consumption. Alternatively, new architectures can be made that use less separate delay lines or other circuit components that are less energy intensive in the LOFAR pipeline. Overall the calculations were based on literature, where the specifications are different than the ones needed in LOFAR. The actual power consumption will be different as a result, which may only be known by doing experimental research. Finally, the NVM devices are still in its developmental stages. Their performances may yet improve greatly, meaning that the state-of-the-art may be much different in a decade.

While not needed for the proposed architecture, it has been shown that the beamformer can also be made using neuromorphic components. Even if it does not prove useful for LOFAR, it may be relevant for other radio-astronomical use cases or other use cases where radars are needed like in the defence industry. It is important to note, however, that such systems need tunability at a high frequency. It has been shown that the write time for NVM devices is in the ns range (table 2), so it has to be analysed how many read-write cycles are necessary to reach significant accuracy for any specific use case. Furthermore, the endurance of the NVM devices need to be taken into account. Currently, endurances go up to 10^6 read-write cycles for RRAM, which would last approximately 1 year if there is 1 read-write cycle every minute. For most systems this endurance is much too low, especially when higher frequencies of read-write cycles are needed. Though, the endurance can be significantly higher depending on which NVM device is used and it may also improve in the coming decades. It is furthermore necessary to have low read voltages. For example, in PCM a maximum read voltage is in the order of 0.1 V to prevent Joule heating. It has to be analysed whether this is feasible for the LOFAR system.

With the shift toward the analogue domain, there is a clear relevance of noise levels in the system (see also Section 3.3). Consequently, it needs to be thoroughly analysed whether the signal can be accurate enough for LOFAR's analysis purposes. For neuromorphic devices especially, the effect of a constant analogue current for extended periods of time is unknown. In part this may increase noise even more, but heating may also be relevant. For this latter point, small heating fans can be employed. If accuracy proves to be a significant issue, one possible solution is given in Section 4.4. However, due to the necessity of multiple parallel crossbars for this technique, the total power consumption scales with the number of parallel crossbars too. Furthermore, the latency will increase significantly by using this technique, while that should be one of the improvements of using neuromorphic hardware.

In principle, if the delay lines are ignored, the DFT is the major power consuming component within the front-end processing pipeline. Here it was opted for a DFT rather than an FFT despite its theoretically higher energy intensity, for the purposes of applicability of NVM devices. However, it may yet be possible to do an FFT as well using neuromorphic components, with the use of clever techniques [127]. However, the implementation of such an architecture may add unwanted energy costs in transport. If in the future neuromorphic components are implemented in the CEP in Groningen, once more the FT can be done cheaper energetically. However, one big point of power consumption of the inverse FFT (IFFT) comes from the necessary equally spaced grid [128]. The added costs that make sure there is indeed functionally an equally spaced grid may not be worth the avoided costs of a neuromorphic IFFT compared to a neuromorphic IDFT.

For steps towards the future, it may be interesting to also test the implementation of the digital neuromorphic domain. This would include energy intensive DACs and ADCs, but is more accurate. Furthermore, this paradigm would allow for cheap transport and computing over long distances, which means that processing steps in the CEP in Groningen may also be done with neuromorphic elements. Alternatively, if work on neuromorphic computing in the analogue domain yields positive results, it may be possible to do initial computing steps as such and additional neuromorphic processing in the CEP in Groningen in the digital domain.

It may also be possible to look into photonics or other alternatives that may be similarly energetically cheap. One example would be photonic beamformers [129].

6 Conclusion

In this thesis, the implementations of analogue and neuromorphic architectures for front-end processing in LOFAR were analysed, with a focus on the neuromorphic paradigm. It has been shown that the energy intensity decreases by 4 ± 1 orders of magnitude for the neuromorphic implementations, with the delay lines of the FIR being the major contributor to the power consumption. Further research still needs to be done to more accurately conclude the total power consumption, as well as for testing the noise levels that are paired with going into the analogue domain. Together with the knowledge gathered in this thesis, this will provide a new, energetically cheap way to do front-end processing in radio astronomy.

References

[1] A. H. P. W. J. N. Gavin Madakumbura, Chad Thackeray and R. Sukhdeo, "Climate change a factor in unprecedented LA fires," 1 2025. [Online]. Available: https://sustainablela.ucla.edu/2025lawildfires

[2] J. G. Canadell, C. Meyer, G. D. Cook, A. Dowdy, P. R. Briggs, J. Knauer, A. Pepler, and V. Haverd, "Multi-decadal increase of forest burned area in australia is linked to climate change," *Nature communications*, vol. 12, no. 1, p. 6921, 2021.

[3] A. Dai, "Increasing drought under global warming in observations and models," *Nature climate change*, vol. 3, no. 1, pp. 52–58, 2013.

[4] C. Bellard, C. Bertelsmeier, P. Leadley, W. Thuiller, and F. Courchamp, "Impacts of climate change on the future of biodiversity," *Ecology letters*, vol. 15, no. 4, pp. 365–377, 2012.

[5] J. Rockström, W. Steffen, K. Noone, Å. Persson, F. S. Chapin III, E. Lambin, T. M. Lenton, M. Scheffer, C. Folke, H. J. Schellnhuber *et al.*, "Planetary boundaries: exploring the safe operating space for humanity," *Ecology and society*, vol. 14, no. 2, 2009.

[6] T. Hák, S. Janoušková, and B. Moldan, "Sustainable development goals: A need for relevant indicators," *Ecological indicators*, vol. 60, pp. 565–573, 2016.

[7] United Nations, "Paris agreement," in *Report of the conference of the parties to the United Nations framework convention on climate change (21st session, 2015: Paris).*, vol. 4, no. 2017. HeinOnline, 2015, p. 2.

[8] C. Broekema, "Commodity compute- and data-transport system design in modern large-scale distributed radio telescopes," Ph.D. dissertation, Vrije Universiteit Amsterdam, 2020.

[9] M. P. van Haarlem, M. W. Wise, A. Gunst, G. Heald, J. P. McKean, J. W. Hessels, A. G. de Bruyn, R. Nijboer, J. Swinbank, R. Fallows *et al.*, "Lofar: The low-frequency array," *Astronomy & astrophysics*, vol. 556, p. A2, 2013.

[10] B. Semeijn, "Internet volume in Amsterdam doubles in four years' time — AMS-IX Amsterdam," 2023. [Online]. Available: https://www.ams-ix.net/ams/news/internet-volume-in-amsterdam-doubles-in-four-years-time

[11] G. Kruithof, C. Bassa, I. Bonati, W. van Cappellen, A. Doek, N. Ebbendorf, M. Gerbers, M. van Haarlem, R. Halfwerk, H. Holties *et al.*, "The energy consumption and carbon footprint of the lofar telescope," *Experimental Astronomy*, vol. 56, no. 2, pp. 687–714, 2023.

[12] "Global carbon budget: Per capita co2 emissions - gcb," Our World In Data, Tech. Rep., 2024. [Online]. Available: https://ourworldindata.org/grapher/co-emissions-per-capita

[13] Eurostat, "Supply, transformation and consumption of electricity," 2025. [Online]. Available: https://ec.europa.eu/eurostat/databrowser/view/nrg_cb_e/default/table?lang=en

[14] T. Können, "Stroom te duur voor radiotelescoop LOFAR," *De Ingenieur*, 9 2022. [Online]. Available: https://deingenieur.nl/artikelen/stroom-te-duur-voor-radiotelescoop-lofar

[15] D. J. Kösters, B. A. Kortman, I. Boybat, E. Ferro, S. Dolas, R. Ruiz de Austri, J. Kwisthout, H. Hilgenkamp, T. Rasing, H. Riel, A. Sebastian, S. Caron, and J. H. Mentink, "Benchmarking energy consumption and latency for neuromorphic computing in condensed matter and particle physics," *APL Machine Learning*, vol. 1, no. 1, p. 016101, 02 2023. [Online]. Available: https://doi.org/10.1063/5.0116699

[16] M. Mahdavi, "Fully parallel and reconfigurable realization of dft/idft using in-memory computing," *IEEE*, pp. 1–8, 2023.

[17] B. F. Burke, F. Graham-Smith, and P. N. Wilkinson, *An introduction to radio astronomy*. Cambridge University Press, 2019.

[18] M. Sokolowski, R. B. Wayth, and M. Lewis, "The statistics of low frequency radio interference at the murchison radio-astronomy observatory," in *2015 IEEE Global Electromagnetic Compatibility Conference (GEMCCON)*, 2015, pp. 1–6.

[19] M. de Vos, A. W. Gunst, and R. Nijboer, "The lofar telescope: System architecture and signal processing," *Proceedings of the IEEE*, vol. 97, no. 8, pp. 1431–1437, 2009.

[20] "What's the pass band ripple and stop band attenuation of a digital filter?" [Online]. Available: https://dsp.stackexchange.com/questions/38564/whats-the-pass-band-ripple-and-stop-band-attenuation-of-a-digital-filter

[21] J. W. Romein, P. C. Broekema, E. van Meijeren, K. van der Schaaf, and W. H. Zwart, "Astronomical real-time streaming signal processing on a blue gene/l supercomputer," in *Proceedings of the Eighteenth Annual ACM Symposium on Parallelism in Algorithms and Architectures*, ser. SPAA '06. New York, NY, USA: Association for Computing Machinery, 2006, p. 59–66. [Online]. Available: https://doi.org/10.1145/1148109.1148118

[22] H. Alzaq and B. Berk Ustundag, "A comparative analysis of multiplier-less 1-level discrete wavelet transform implementations on fpgas," in *2017 International Conference on Engineering and Technology (ICET)*, 2017, pp. 1–6.

[23] S. K. Mitra, *Digital Signal Processing: A Computer-Based Approach*, 2nd ed. McGraw-Hill Higher Education, 2001.

[24] T. Saramäki, S. Mitra, and J. Kaiser, "Finite impulse response filter design," *Handbook for digital signal processing*, vol. 4, pp. 155–277, 1993.

[25] J. F. James, *A Student's Guide to Fourier Transforms: With Applications in Physics and Engineering*, 3rd ed., ser. Student's Guides. Cambridge University Press, 2011.

[26] "Fourier Transforms for Continuous/Discrete Time/Frequency —." [Online]. Available: https://www.dsprelated.com/freebooks/mDFT/Fourier_Transforms_Continuous_Discrete_Time_Frequency.html

[27] O. M. Smirnov, "Revisiting the radio interferometer measurement equation: I. a full-sky jones formalism," *Astronomy amp; Astrophysics*, vol. 527, p. A106, Feb. 2011. [Online]. Available: http://dx.doi.org/10.1051/0004-6361/201016082

[28] W. contributors, "DFT matrix," 9 2024. [Online]. Available: https://en.wikipedia.org/wiki/DFT_matrix

[29] S. Yatawatta, "Polarization-based online interference mitigation in radio interferometry," in *2020 28th European Signal Processing Conference (EUSIPCO)*, 2021, pp. 1961–1965.

[30] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to algorithms*. MIT press, 2022.

[31] J. G. Proakis, *Digital signal processing: principles, algorithms, and applications, 4/E*. Pearson Education India, 2007.

[32] M. T. Heideman, D. H. Johnson, and C. S. Burrus, "Gauss and the history of the fast fourier transform," *Archive for History of Exact Sciences*, vol. 34, no. 3, pp. 265–277, 1985. [Online]. Available: http://www.jstor.org/stable/41133773

[33] G. Wolberg, "Fast fourier transforms: A review," 1988.

[34] G. Danielson and C. Lanczos, "Some improvements in practical fourier analysis and their application to x-ray scattering from liquids," *Journal of the Franklin Institute*, vol. 233, no. 4, pp. 365–380, 1942. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0016003242907671

[35] J. W. Cooley and J. W. Tukey, "An algorithm for the machine calculation of complex fourier series," *Mathematics of Computation*, vol. 19, no. 90, pp. 297–301, 1965. [Online]. Available: http://www.jstor.org/stable/2003354

[36] B. V. Veen, "The fast fourier transform algorithm," 12 2012. [Online]. Available: https://www.youtube.com/watch?v=EsJGuI7e_ZQ

[37] E. Brigham, *The fast Fourier transform and its applications*. Prentice Hall, 1988.

[38] E. R. Ramaswami Sachidanandan, N. Singh, and S. Gunda, "Design and simulation of a low-power and high-speed fast fourier transform for medical image compression," vol. 34, 03 2023, p. 18.

[39] S. Dirlik, "A comparison of FFT processor designs," 12 2013. [Online]. Available: https://essay.utwente.nl/72179/1/FFT_Comparison_Simon_Dirlik.pdf

[40] S. Roy, "FPGA Implementation of 8-Point FFT," 2 2021. [Online]. Available: https://digitalsystemdesign.in/fpga-implementation-of-8-point-fft/

[41] S. Dhok, "Mod6 Lec19: Comparing DIT-FFT and DIF-FFT Algorithms. How are they "fast"?" 6 2020. [Online]. Available: https://www.youtube.com/watch?v=lku836K6s-Q

[42] N.-K. Tsao, "The equivalence of decimation in time and decimation in frequency in fft computations," *Journal of the Franklin Institute*, vol. 324, no. 1, pp. 43–63, 1987.

[43] "Acquiring an Analog Signal: Bandwidth, Nyquist Sampling Theorem, and Aliasing," 8 2006. [Online]. Available: https://www.ni.com/en/shop/data-acquisition/measurement-fundamentals/analog-fundamentals/acquiring-an-analog-signal--bandwidth--nyquist-sampling-theorem-.html

[44] C. Broekema, "Numerical analysis of the lofar remote station beamformer," Master's thesis, Rijksuniversiteit Groningen, 2004.

[45] M. Elmer, B. D. Jeffs, K. F. Warnick, J. R. Fisher, and R. D. Norrod, "Beamformer design methods for radio astronomical phased array feeds," *IEEE Transactions on Antennas and Propagation*, vol. 60, no. 2, pp. 903–914, 2012.

[46] G. Chardon, "Theoretical analysis of beamforming steering vector formulations for acoustic source localization," *Journal of Sound and Vibration*, vol. 517, p. 116544, 2022.

[47] G. Zatorre, S. Celma, C. Aldea, N. Medrano, and E. Peralias, "Digital auto-tuning system for analog filters," in *2006 International Caribbean Conference on Devices, Circuits and Systems*, 2006, pp. 37–41.

[48] Y. Liu and H. Cui, "Antenna array signal direction of arrival estimation on digital signal processor (dsp)," *Procedia Computer Science*, vol. 55, pp. 782–791, 2015.

[49] M. Elmer, B. D. Jeffs, K. F. Warnick, J. R. Fisher, and R. D. Norrod, "Beamformer design methods for radio astronomical phased array feeds," *IEEE Transactions on Antennas and Propagation*, vol. 60, no. 2, pp. 903–914, 2012.

[50] B. D. Van Veen and K. M. Buckley, "Beamforming: A versatile approach to spatial filtering," *IEEE assp magazine*, vol. 5, no. 2, pp. 4–24, 1988.

[51] K. M. B. Asad, L. V. E. Koopmans, V. Jelić, V. N. Pandey, A. Ghosh, F. B. Abdalla, G. Bernardi, M. A. Brentjens, A. G. de Bruyn, S. Bus, B. Ciardi, E. Chapman, S. Daiboo, E. R. Fernandez, G. Harker, I. T. Iliev, H. Jensen, O. Martinez-Rubi, G. Mellema, M. Mevius, A. R. Offringa, A. H. Patil, J. Schaye, R. M. Thomas, S. van der Tol, H. K. Vedantham, S. Yatawatta, and S. Zaroubi, "Polarization leakage in epoch of reionization windows – I. Low Frequency Array observations of the 3C196 field," *Monthly Notices of the Royal Astronomical Society*, vol. 451, no. 4, pp. 3709–3727, 06 2015. [Online]. Available: https://doi.org/10.1093/mnras/stv1107

[52] G. Ifrah, *The universal history of computing: From the abacus to quantum computing*. John Wiley & Sons, Inc., 2000.

[53] J. Von Neumann, "First draft of a report on the edvac," *IEEE Annals of the History of Computing*, vol. 15, no. 4, pp. 27–75, 1993.

[54] D. E. Knuth, "Von neumann's first computer program," *ACM Computing Surveys (CSUR)*, vol. 2, no. 4, pp. 247–260, 1970.

[55] J. Shalf, "The future of computing beyond moore's law," *Philosophical Transactions of the Royal Society A*, vol. 378, no. 2166, p. 20190061, 2020.

[56] R. Kurzweil, *The age of spiritual machines: When computers exceed human intelligence*. Penguin, 2000.

[57] M. Fischetti, "Computers versus Brains," *Scientific American*, 2 2024. [Online]. Available: https://www.scientificamerican.com/article/computers-vs-brains/

[58] G. Indiveri, "Introducing neuromorphic computing and engineering," *Neuromorphic Computing and Engineering*, vol. 1, no. 1, p. 010401, 2021.

[59] N. Verma, H. Jia, H. Valavi, Y. Tang, M. Ozatay, L.-Y. Chen, B. Zhang, and P. Deaville, "In-memory computing: Advances and prospects," *IEEE solid-state circuits magazine*, vol. 11, no. 3, pp. 43–55, 2019.

[60] J. M. Rabaey, A. Chandrakasan, and B. Nikolic, *Digital integrated circuits*. Prentice hall Englewood Cliffs, 2002, vol. 2.

[61] B. P. Bean, "The action potential in mammalian central neurons," *Nature Reviews Neuroscience*, vol. 8, no. 6, pp. 451–465, 2007.

[62] E.-V. B. Network, "Introduction to artificial neural networks and the perceptron." [Online]. Available: https://evbn.org/Introduction-to-Artificial-Neural-Networks-and-the-Perceptron--QuantStart-1678014456/

[63] F. A. Azevedo, L. R. Carvalho, L. T. Grinberg, J. M. Farfel, R. E. Ferretti, R. E. Leite, W. J. Filho, R. Lent, and S. Herculano-Houzel, "Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain," *Journal of Comparative Neurology*, vol. 513, no. 5, pp. 532–541, 2009.

[64] M. B. Kennedy, "Synaptic signaling in learning and memory," *Cold Spring Harbor perspectives in biology*, vol. 8, no. 2, p. a016824, 2016.

[65] F. Rosenblatt, "The perceptron: A probabilistic model for information storage and organization in the brain," *Psychological Review*, vol. 65, no. 6, pp. 386–408, 1 1958. [Online]. Available: https://doi.org/10.1037/h0042519

[66] A. G. Ivakhnenko, "Polynomial Theory of Complex Systems," *IEEE Transactions on Systems Man and Cybernetics*, vol. SMC-1, no. 4, pp. 364–378, 10 1971. [Online]. Available: https://doi.org/10.1109/tsmc.1971.4308320

[67] Veritasium, "Future computers will be radically different (Analog computing)," 3 2022. [Online]. Available: https://www.youtube.com/watch?v=GVsUOuSjvcg

[68] A. Krenker, J. Bešter, and A. Kos, "Introduction to the artificial neural networks," *Artificial Neural Networks: Methodological Advances and Biomedical Applications. InTech*, pp. 1–18, 2011.

[69] L. Ye, Z. Gao, J. Fu, W. Ren, C. Yang, J. Wen, X. Wan, Q. Ren, S. Gu, X. Liu *et al.*, "Overview of memristor-based neural network design and applications," *Frontiers in Physics*, vol. 10, p. 839243, 2022.

[70] L. Chua, "Memristor-the missing circuit element," *IEEE Transactions on Circuit Theory*, vol. 18, no. 5, pp. 507–519, 1971.

[71] D. B. Strukov, G. S. Snider, D. R. Stewart, and R. S. Williams, "The missing memristor found," *nature*, vol. 453, no. 7191, pp. 80–83, 2008.

[72] J. Blanc and D. L. Staebler, "Electrocoloration in srtio$_3$: Vacancy drift and oxidation-reduction of transition metals," *Phys. Rev. B*, vol. 4, pp. 3548–3557, Nov 1971. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevB.4.3548

[73] L. Chua, "Resistance switching memories are memristors," *Handbook of memristor networks*, pp. 197–230, 2019.

[74] C. Mellor, "HP and Hynix to produce the memristor goods by 2013," 10 2011. [Online]. Available: https://www.theregister.com/2011/10/10/memristor_in_18_months/

[75] B. Sun, M. Xiao, G. Zhou, Z. Ren, Y. Zhou, and Y. Wu, "Non–zero-crossing current-voltage hysteresis behavior in memristive system," *Materials Today Advances*, vol. 6, p. 100056, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2590049820300035

[76] M. Le Gallo, C. Lammie, J. Büchel, F. Carta, O. Fagbohungbe, C. Mackin, H. Tsai, V. Narayanan, A. Sebastian, K. El Maghraoui, and M. J. Rasch, "Using the IBM analog in-memory hardware acceleration kit for neural network training and inference," *APL Machine Learning*, vol. 1, no. 4, p. 041102, 11 2023. [Online]. Available: https://doi.org/10.1063/5.0168089

[77] O. Kapur, D. Guo, J. Reynolds, D. Newbrook, Y. Han, R. Beanland, L. Jiang, C. K. de Groot, and R. Huang, "An ultra high-endurance memristor using back-end-of-line amorphous sic," *Scientific Reports*, vol. 14, no. 1, p. 14008, 2024.

[78] Y. Xiao, B. Jiang, Z. Zhang, S. Ke, Y. Jin, X. Wen, and C. Ye, "A review of memristor: material and structure design, device performance, applications and prospects," *Science and Technology of Advanced Materials*, vol. 24, no. 1, p. 2162323, 2023.

[79] Y. Y. Chen, L. Goux, S. Clima, B. Govoreanu, R. Degraeve, G. S. Kar, A. Fantini, G. Groeseneken, D. J. Wouters, and M. Jurczak, "Endurance/retention trade-off on HfO_2/Metal cap 1t1r bipolar rram," *IEEE Transactions on Electron Devices*, vol. 60, no. 3, pp. 1114–1121, 2013.

[80] P.-Y. Jung, D. Panda, S. Chandrasekaran, S. Rajasekaran, and T.-Y. Tseng, "Enhanced switching properties in taox memristors using diffusion limiting layer for synaptic learning," *IEEE Journal of the Electron Devices Society*, vol. 8, pp. 110–115, 2020.

[81] S. M. Regina Dittmann and R. Waser, "Nanoionic memristive phenomena in metal oxides: the valence change mechanism," *Advances in Physics*, vol. 70, no. 2, pp. 155–349, 2021. [Online]. Available: https://doi.org/10.1080/00018732.2022.2084006

[82] W. Sun, B. Gao, M. Chi, Q. Xia, J. J. Yang, H. Qian, and H. Wu, "Understanding memristive switching via in situ characterization and device modeling," *Nature communications*, vol. 10, no. 1, p. 3453, 2019.

[83] F. Zahoor, T. Z. Azni Zulkifli, and F. A. Khanday, "Resistive random access memory (rram): an overview of materials, switching mechanism, performance, multilevel cell (mlc) storage, modeling, and applications," *Nanoscale research letters*, vol. 15, pp. 1–26, 2020.

[84] M. Kaniselvan, M. Luisier, and M. Mladenović, "An atomistic modeling framework for valence change memory cells," *Solid-State Electronics*, vol. 199, p. 108506, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0038110122002775

[85] G. W. Burr, M. J. Breitwisch, M. Franceschini, D. Garetto, K. Gopalakrishnan, B. Jackson, B. Kurdi, C. Lam, L. A. Lastras, A. Padilla *et al.*, "Phase change memory technology," *Journal of Vacuum Science & Technology B*, vol. 28, no. 2, pp. 223–262, 2010.

[86] M. K. Qureshi, M. M. Franceschini, and L. A. Lastras-Montano, "Improving read performance of phase change memories via write cancellation and write pausing," in *HPCA-16 2010 The Sixteenth International Symposium on High-Performance Computer Architecture*. IEEE, 2010, pp. 1–11.

[87] G. W. Burr, A. Sebastian, T. Ando, and W. Haensch, "Ohm's law + kirchhoff's current law = better ai: Neural-network processing done in memory with analog circuits will save energy," *IEEE Spectrum*, vol. 58, no. 12, pp. 44–49, 2021.

[88] S. Raoux, G. W. Burr, M. J. Breitwisch, C. T. Rettner, Y.-C. Chen, R. M. Shelby, M. Salinga, D. Krebs, S.-H. Chen, H.-L. Lung *et al.*, "Phase-change random access memory: A scalable technology," *IBM Journal of Research and Development*, vol. 52, no. 4.5, pp. 465–479, 2008.

[89] S. Meister, S. Kim, J. J. Cha, H.-S. P. Wong, and Y. Cui, "In situ transmission electron microscopy observation of nanostructural changes in phase-change memory," *ACS Nano*, vol. 5, no. 4, pp. 2742–2748, 2011, pMID: 21425849. [Online]. Available: https://doi.org/10.1021/nn1031356

[90] Y. Zeng, J. Jin, R. Gu, X. Cheng, M. Xu, L. Gao, and X. Miao, "A fast and high endurance phase change memory based on in-doped sb2te3," *ACS Applied Nano Materials*, vol. 7, no. 12, pp. 13 983–13 990, 2024. [Online]. Available: https://doi.org/10.1021/acsanm.4c00809

[91] R. Khaddam-Aljameh, M. Stanisavljevic, J. Fornt Mas, G. Karunaratne, M. Brändli, F. Liu, A. Singh, S. M. Müller, U. Egger, A. Petropoulos, T. Antonakopoulos, K. Brew, S. Choi, I. Ok, F. L. Lie, N. Saulnier, V. Chan, I. Ahsan, V. Narayanan, S. R. Nandakumar, M. Le Gallo, P. A. Francese, A. Sebastian, and E. Eleftheriou, "Hermes-core—a 1.59-tops/mm2 pcm on 14-nm cmos in-memory compute core using 300-ps/lsb linearized cco-based adcs," *IEEE Journal of Solid-State Circuits*, vol. 57, no. 4, pp. 1027–1038, 2022.

[92] M. Lanza, A. Sebastian, W. D. Lu, M. L. Gallo, M.-F. Chang, D. Akinwande, F. M. Puglisi, H. N. Alshareef, M. Liu, and J. B. Roldan, "Memristive technologies for data storage, computation, encryption, and radio-frequency communication," *Science*, vol. 376, no. 6597, p. eabj9979, 2022. [Online]. Available: https://www.science.org/doi/abs/10.1126/science.abj9979

[93] A. Chen, "A review of emerging non-volatile memory (nvm) technologies and applications," *Solid-State Electronics*, vol. 125, pp. 25–38, 2016, extended papers selected from ESSDERC 2015. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0038110116300867

[94] H. Kwak, N. Kim, S. Jeon, S. Kim, and J. Woo, "Electrochemical random-access memory: recent advances in materials, devices, and systems towards neuromorphic computing," *Nano Convergence*, vol. 11, no. 1, p. 9, 2024.

[95] H. Kohlstedt, Y. Mustafa, A. Gerber, A. Petraru, M. Fitsilis, R. Meyer, U. Böttger, and R. Waser, "Current status and challenges of ferroelectric memory devices," *Microelectronic Engineering*, vol. 80, pp. 296–304, 2005.

[96] J. Ajayan, P. Mohankumar, D. Nirmal, L. L. Joseph, S. Bhattacharya, S. Sreejith, S. Kollem, S. Rebelli, S. Tayal, and B. Mounika, "Ferroelectric field effect transistors (fefets): Advancements, challenges and exciting prospects for next generation non-volatile memory (nvm) applications," *Materials Today Communications*, vol. 35, p. 105591, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2352492823002817

[97] Y. Development, "Emerging Non-Volatile Memory 2021," Tech. Rep., 2 2021. [Online]. Available: https://web.archive.org/web/20221225072318/https://www.i-micronews.com/products/emerging-non-volatile-memory-2021/?NAND_monitor_March2021

[98] M. Le Gallo, D. Krebs, F. Zipoli, M. Salinga, and A. Sebastian, "Collective structural relaxation in phase-change memory devices," *Advanced Electronic Materials*, vol. 4, no. 9, p. 1700627, 2018.

[99] Q. Wang, Y. Itoh, T. Tsuruoka, M. Aono, D. He, and T. Hasegawa, "Oxygen vacancy drift controlled three-terminal reram with a reduction in operating gate bias and gate leakage current," *Solid State Ionics*, vol. 328, pp. 30–34, 2018.

[100] J. Li, Y. Xue, Y. Li, H. Jia, Z. Zhou, L. Yang, S. Ren, J. Chen, Y. He, K. Xue *et al.*, "Fully analog iteration for solving matrix equations with in-memory computing," *Science Advances*, vol. 11, no. 7, p. eadr6391, 2025.

[101] J. Büchel, G. Camposampiero, A. Vasilopoulos, C. Lammie, M. Le Gallo, A. Rahimi, and A. Sebastian, "Kernel approximation using analogue in-memory computing," *Nature Machine Intelligence*, vol. 6, no. 12, pp. 1605–1615, 2024.

[102] F. Zahoor, T. Z. Azni Zulkifli, and F. A. Khanday, "Resistive random access memory (rram): an overview of materials, switching mechanism, performance, multilevel cell (mlc) storage, modeling, and applications," *Nanoscale research letters*, vol. 15, pp. 1–26, 2020.

[103] W. Shim, J.-s. Seo, and S. Yu, "Two-step write–verify scheme and impact of the read noise in multilevel rram-based inference engine," *Semiconductor Science and Technology*, vol. 35, no. 11, p. 115026, 2020.

[104] S. Raoux, F. Xiong, M. Wuttig, and E. Pop, "Phase change materials and phase change memory," *MRS bulletin*, vol. 39, no. 8, pp. 703–710, 2014.

[105] W. Song, M. Rao, Y. Li, C. Li, Y. Zhuo, F. Cai, M. Wu, W. Yin, Z. Li, Q. Wei, S. Lee, H. Zhu, L. Gong, M. Barnell, Q. Wu, P. A. Beerel, M. S.-W. Chen, N. Ge, M. Hu, Q. Xia, and J. J. Yang, "Programming memristor arrays with arbitrarily high precision for analog computing," *Science*, vol. 383, no. 6685, pp. 903–910, 2024. [Online]. Available: https://www.science.org/doi/abs/10.1126/science.adi9405

[106] K. R. Betty and G. Horlick, "Transversal filtering of analog signals with a tapped analog delay line," *Analytical chemistry*, vol. 48, no. 14, pp. 2248–2252, 12 1976. [Online]. Available: https://doi.org/10.1021/ac50008a049

[107] X. Lin, J. Liu, H. Lee, and H. Liu, "A 2.5- to 3.5-gb/s adaptive fir equalizer with continuous-time wide-bandwidth delay line in 0.25-*muhboxm*cmos," *IEEE Journal of Solid-State Circuits*, vol. 41, no. 8, pp. 1908–1918, 2006.

[108] M. A. A. Pinheiro, A. Petraglia, and M. R. Petraglia, "Performance of analogue, digital and hybrid finite impulse response filter banks affected by realisation errors and noise," *IET signal processing*, vol. 12, no. 8, pp. 949–956, 10 2018. [Online]. Available: https://doi.org/10.1049/iet-spr.2018.0105

[109] R. Marston, "Analog Delay Lines," *Folkscanomy Electronics*, pp. 67–71, 10 1986. [Online]. Available: https://www.n5dux.com/ham/files/pdf/Analog%20Delay%20Lines.pdf

[110] R. Lerner, "Band-pass filters with linear phase," *Proceedings of the IEEE*, vol. 52, no. 3, pp. 249–268, 1964.

[111] C.-W. Tseng, Z. Feng, Z. Fan, H. An, Y. Wang, H.-S. Kim, and D. Blaauw, "A low-power highly reconfigurable analog fir filter with 11-bit charge-domain dac for narrowband receivers," *IEEE Solid-State Circuits Letters*, vol. 7, pp. 74–77, 2024.

[112] M. Lehne and S. Raman, "A 0.13-μm 1-gs/s cmos discrete-time fft processor for ultra-wideband ofdm wireless receivers," *IEEE Transactions on Microwave Theory and Techniques*, vol. 59, no. 6, pp. 1639–1650, 2011.

[113] P. Reshma, V. P. Gopi, V. S. Babu, and K. A. Wahid, "Analog cmos implementation of fft using cascode current mirror," *Microelectronics Journal*, vol. 60, pp. 30–37, 2017.

[114] K. Kleine, J. H. Reed, and A. J. Michaels, "A 256-point analog discrete-time fft," in *2020 IEEE 63rd International Midwest Symposium on Circuits and Systems (MWSCAS)*, 2020, pp. 966–969.

[115] M. Lehne and S. Raman, "A discrete-time fft processor for ultrawideband ofdm wireless transceivers: Architecture and behavioral modeling," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 57, no. 11, pp. 3011–3022, 2010.

[116] J. Sabater, S. Sánchez-Expósito, P. Best, J. Garrido, L. Verdes-Montenegro, and D. Lezzi, "Calibration of lofar data on the cloud," *Astronomy and computing*, vol. 19, pp. 75–89, 2017.

[117] S. Sharma and T. Ytterdal, "Low-power low-area beamformer design using switched-current aram using external capacitors," *Microelectronics Journal*, vol. 48, pp. 76–80, 2016.

[118] J.-Y. Jeong, J.-S. An, S.-J. Jung, S.-K. Hong, and O.-K. Kwon, "A low-power analog delay line using a current-splitting method for 3-d ultrasound imaging systems," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 65, no. 7, pp. 829–833, 2018.

[119] R. Cai, A. Ren, Y. Wang, and B. Yuan, "Memristor-based discrete fourier transform for improving performance and energy efficiency," in *2016 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*, 2016, pp. 643–648.

[120] R. Velik, "Discrete fourier transform computation using neural networks," in *2008 International Conference on Computational Intelligence and Security*, vol. 1, 2008, pp. 120–123.

[121] G. Yuan, C. Ding, R. Cai, X. Ma, Z. Zhao, A. Ren, B. Yuan, and Y. Wang, "Memristor crossbar-based ultra-efficient next-generation baseband processors," in *2017 IEEE 60th International Midwest Symposium on Circuits and Systems (MWSCAS)*, 2017, pp. 1121–1124.

[122] Y. Xi, B. Gao, J. Tang, A. Chen, M.-F. Chang, X. S. Hu, J. V. D. Spiegel, H. Qian, and H. Wu, "In-memory learning with analog resistive switching memory: A review and perspective," *Proceedings of the IEEE*, vol. 109, no. 1, pp. 14–42, 2021.

[123] H. Roy and M. Sharad, "Current mode neuron for the memristor based synapse," *ArXiv*, vol. abs/1905.05307, 2019. [Online]. Available: https://api.semanticscholar.org/CorpusID:153313108

[124] Y. Zhu, G. L. Zhang, T. Wang, B. Li, Y. Shi, T.-Y. Ho, and U. Schlichtmann, "Statistical training for neuromorphic computing using memristor-based crossbars considering process variations and noise," in *2020 Design, Automation Test in Europe Conference Exhibition (DATE)*, 2020, pp. 1590–1593.

[125] G. Schoonderbeek, A. Szomoru, A. Gunst, L. Hiemstra, and J. Hargreaves, "Uniboard 2 , a generic scalable high-performance computing platform for radio astronomy," *Journal of Astronomical Instrumentation*, vol. 08, 11 2018.

[126] M. Le Gallo, R. Khaddam-Aljameh, M. Stanisavljevic, A. Vasilopoulos, B. Kersting, M. Dazzi, G. Karunaratne, M. Brändli, A. Singh, S. M. Mueller *et al.*, "A 64-core mixed-signal in-memory compute chip based on phase-change memory for deep neural network inference," *Nature Electronics*, vol. 6, no. 9, pp. 680–693, 2023.

[127] O. Leitersdorf, Y. Boneh, G. Gazit, R. Ronen, and S. Kvatinsky, "Fourierpim: High-throughput in-memory fast fourier transform and polynomial multiplication," *Memories-Materials, Devices, Circuits and Systems*, vol. 4, p. 100034, 2023.

[128] M. Wise, "Radio Astronomy Lecture 8: The Techniques of Radio Interferometry III: Imaging," 5 2017. [Online]. Available: https://www.astron.nl/astrowiki/lib/exe/fetch.php?media=ra_uva:2017:ra_uva_lecture8_keynote9.pdf

[129] P. Martinez-Carrasco, T. H. Ho, D. Wessel, and J. Capmany, "Ultrabroadband high-resolution silicon rf-photonic beamformer," *Nature Communications*, vol. 15, no. 1, p. 1433, 2024.

[130] K.-C. Kuo, *Advances in Solid State Circuit Technologies*. InTechOpen, 4 2010.

[131] G. Han and E. Sanchez-Sinencio, "Cmos transconductance multipliers: a tutorial," *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 45, no. 12, pp. 1550–1563, 1998.

7 Appendices

Here, additional information is given to expand on the results section.

7.1 Appendix A. Analogue pipeline

Within the analogue FIR, there is a necessity for multipliers and summations. [112] gives these in the form of transconductances and linear adders as shown in figures 43 and 44, respectively.

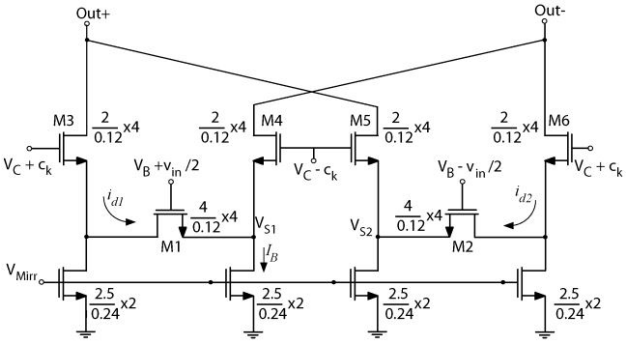

Figure 43: The coefficient multiplier circuit based on a linear transconductor. Image taken from [112].

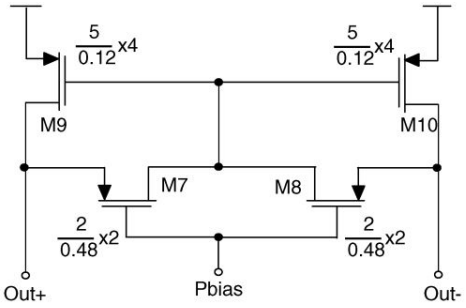

Figure 44: A linear adder used for implementation of the FFT shown in figure 32. Image taken from [112].

Figure 43 shows an coefficient multiplier circuit based on a linear transconductor, which is a device that performs voltage-to-current conversion [130]. This can fix issues with preserving the parameter that is being observed. For example, if the measured quantity from the FIR filter into the FFT is current, but the FFT is only feasible by using voltage to represent the coefficients of the matrices, then the transconductor can convert that voltage to the form of a current value. The transconductance multipliers can be made in the form of many MOS structures [131].

Figure 44 shows a linear adder, which can be used for implementation of an analogue FFT. This linear adder — proposed in [112] — works as a transresistor. This means that, similar to the multiplier, multiple differential current signals are summed to reach a single differential voltage as its output.