TUM

TECHNISCHE UNIVERSITÄT MÜNCHEN
INSTITUT FÜR INFORMATIK

Resonate-and-Fire Neurons as Frequency Selective Input Encoders for Spiking Neural Networks

Daniel Auge, Etienne Mueller

TUM-I2083

Technischer Bericht
Technische Universität München
Institut für Informatik

Resonate-and-Fire Neurons as Frequency Selective Input Encoders for Spiking Neural Networks

Daniel Auge and **Etienne Mueller**
Technical University of Munich

***Abstract*—More and more embedded applications demand to cope with complex data while still being energy-efficient. Neural networks provide the processing capabilities, but often cannot be utilised because of restricted power goals. Spiking neural networks have been shown to potentially solve this problem due to their hardware friendliness and energy efficiency. One remaining problem is the conversion of input data into event-based spikes in order to be processed. In this study, we examine using resonating neuron models to perform spectral transform and temporal spike encoding simultaneously directly on the analogue signal. Additionally, we compare the approach with the fast Fourier transform and demonstrate that both show comparable results while the neuromorphic realisation consumes significantly less energy. With the resonating neurons as input stage for large spiking neural networks, it is possible to realise energy efficient networks in neuromorphic hardware without the need of any digital logic. A limitation of the approach is the comparatively large silicon area needed to realise the circuit.**

I. Introduction

Routine processing of analogue data involves initial sampling and quantisation by analogue-to-digital converters, and subsequent transformation into the frequency domain using the fast Fourier transform (FFT). In embedded processing chains, general purpose digital signal processors, micro controllers, or application specific circuits are then employed to solve the desired processing task. To cope with highly complex tasks, neural network-based solutions are occasionally integrated directly into modern embedded systems. Their usage is often restricted by the silicon area, power consumption, processing speed, and accuracy, that determine the price and feasibility of an embedded application.

Biologically inspired spiking neural networks are being investigated as alternatives to digitally implemented networks. These are proven to perform calculations using less neurons compared to ordinary neural networks [1] and exhibit good hardware integrability. As their biological counterparts, spiking neurons intercommunicate using short energy pulses. Information is thereby encoded into the presence of spikes and their precise timing intervals.

To exploit the capabilities of spiking networks, real-world data must be converted into the spiking domain. In biological systems, parameters including light intensities, pressure levels, and chemical compositions are converted into spike patterns using specific sensors. For example, resonating hair cells reacting to a specific frequency of pressure changes in the cochlea are exploited for hearing. Following the biological model, we feed the signal directly as a current to the input of the electrical resonator circuits to obtain the spectral transform and spike encoding concurrently. The signal processing then proceeds entirely in spiking neural networks without analogue-to-digital converters or other digital circuits. In this study, we propose using resonate-and-fire (RF) neurons [2] as an input stage for spiking neural networks, although the associated computation stages are not examined.

The resonance behaviour of Hodgkin-Huxley-type neurons has been intensively studied [3]–[8]. These studies focused on understanding the behaviour of biological neural networks. Our study in contrast, aims to investigate the practical utility of these properties. Due to their simplicity, RF neurons are well suited for application-oriented approaches, although less biologically plausible. Integrated circuits implementing these neurons have been demonstrated [9], [10].

Research on spiking neural networks commonly involves using frequency-domain inputs [11]–[13], especially in the context of speech processing. Contextually, the digital speech signal is transformed using the FFT and then coded into spike events using time-to-first-spike or rate coding techniques. Wysoski, Benuskova, and Kasabov [2010] use wavelet filters to transform and encode the time-domain signal. Although further elaboration of this proposal is lacking, it necessitates buffering of the signal before application to the parallel input of the network. In contrast, our approach ensures computations are performed entirely in the domain of spiking neural networks. Independent of neural networks, an asynchronous event-based Fourier analysis has been proposed [15]. The authors developed the algorithm to process data produced by an asynchronous image sensor. Several other approaches exist for non-spiking artificial networks. Sainath, Weiss, Wilson, *et al.* [2017] utilized training convolutional filters directly on the buffered time-domain signal to detect features. Graves, Mohamed, and Hinton [2013] used recurrent neural networks for real-time processing of words spoken. For speech enhancement applications, Xu, Du, Dai, and Lee [2014] proposed deep neural networks using the result of a short time FFT as input.

Beside speech recognition, radar signal processing is another well-suited application for such neural networks. Here, the distance measurement of targets is directly linked to the frequency decomposition of the measured signal. Therefore, analysis with RF neurons as the encoding layer is also examined in this study.

We summarize the contribution of our work as follows:

1) We introduce using analogue RF neurons as the input stage for spiking neural networks signal processing. Here, the analogue input signal is directly fed into the neuron as synaptic current.

2) We examine the properties of the resonating neurons (section II) as well as practical limitations of their use (section III). Therefore, we compare our approach to the well-known FFT (section III-C).
3) We illustrate the utility of the proposal in a radar application and compare the performance to that of a common digitally processed solution (section IV).

II. Preliminaries

A. RF Neuron

As proposed by Izhikevich [2001], the RF neuron is a symmetric two-dimensional neuron model with oscillatory behaviour. The dynamics of the neuron are described by

$$\dot{x} = \begin{bmatrix} -d & -\omega_0 \\ \omega_0 & -d \end{bmatrix} x + \begin{bmatrix} I \\ 0 \end{bmatrix} \tag{1}$$

where d is a damping constant, $\omega_0 = 2\pi f_0$ is the resonant frequency of the neuron, and I is the input signal as short pulses or continuous signals. In its equivalent complex form given as

$$\dot{z} = (-d + i\omega_0)z + I \tag{2}$$

the real and imaginary parts of z represent the current- and voltage-like variables. An outgoing spike is generated once the voltage-like variable exceeds a predefined threshold, and both variables are reset. In accordance with biological neurons, the voltage-like variable is also referred as membrane voltage.

B. Properties of the Resonator

The resonating behaviour of the neuron model is first examined without spike emission or reset property. We consider the input current as a continuous sinusoidal signal with amplitude a and frequency $\omega = 2\pi f$ of the form $I = a\cos(\omega t)$. Commonly, the exact courses of z are of less interest, and the main information is provided by the envelope env of the voltage-like variable's oscillation.

1) Pulsed Excitation: If the input of the RF neuron is connected to other neurons, the input signal will, depending on the neurons, most likely consist of short current pulses. Therefore, the time between successive incoming spikes must match the natural frequency of the RF neuron to sufficiently excite it to attain the threshold, as shown by Izhikevich [2001]. This input type is, however, not considered in this study.

2) Resonant Excitation: With the input current oscillating at the resonant frequency of the neuron $\omega = \omega_0$, the transient response follows an exponential process which results in a constant oscillation expressed as

$$z = C\, e^{t\,(-d+\omega_0\,\mathrm{i})} + \frac{e^{t\omega_0 \mathrm{i}}}{2d} + \frac{e^{-t\omega_0 \mathrm{i}}}{2d - 4\omega_0 i} \tag{3}$$

$$\approx \frac{e^{t\omega_0 \mathrm{i}}}{2d} \quad \text{for } t \to \infty \tag{4}$$

hence, the envelope of the process follows the expression

$$\mathrm{env} = \frac{a}{2d}\left(1 - e^{-td}\right) \approx \frac{a}{2}t \tag{5}$$

that is approximated by a linear process for a small t and small damping constant d, independent of the frequency.

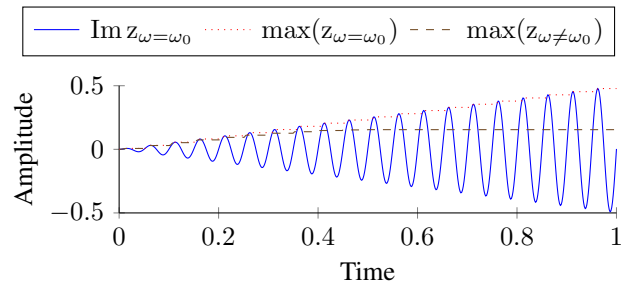

Fig. 1. Membrane voltage course of a neuron excited by a sinusoidal signal of its resonant frequency. Additionally, the oscillation's envelope is shown for a resonant and near-resonant excitation.

3) Non-resonant Excitation: With a non-resonant excitation $\omega \neq \omega_0$, the envelope reaches much lower voltages in the same time as in resonant excitation if not saturated:

$$\mathrm{env} = \frac{2a\,\omega}{\omega^2 - \omega_0^2}\sin(\frac{1}{2}(\omega - \omega_0)t), \tag{6}$$

$$\max(\mathrm{env}) \approx \frac{a}{\Delta\omega} \quad \text{for } \Delta\omega\, T \geq \pi. \tag{7}$$

The maximum voltage differences between resonant and non-resonant frequencies, and with that the frequency selectivity of the neuron is influenced by the damping constant d, the measurement time T, and the frequency difference $\Delta\omega = |\omega - \omega_0|$. The relation given in (7) is also valid for a constant excitation of the resonator with $\omega = 0$. The membrane voltage course of an exemplary neuron excited by a resonant signal is shown in Figure 1. Additionally, the maximum value of the oscillation's envelope is depicted for the resonant and near-resonant excitation. The differences between the envelopes are proportional to the duration of the measurements. The voltage difference between the resonant frequency ω_0 and the non-resonant frequency $\omega = \omega_0 + \frac{2\pi}{T}$ remains constant for all ω_0 at a constant measurement time T.

III. Approach

A. Spike Generation and Reset Behaviour

To enter the spiking domain, the neuron must generate outgoing current pulses, depending on its level of excitation. The requirements for spike generation include:

1) a neuron only responds to signals with near-resonant frequencies over a broad range of signal amplitudes/ signal-to-noise ratios (SNR),
2) determining the amplitude of the signal is possible,
3) determining the phase of the signal is also possible.

These requirements indicate two possible rules for generating spikes:

- reaching the threshold voltage, accompanied by resetting z,
- attaining the threshold voltage, followed by increasing the threshold without resetting z. The threshold adaption is linear, exponential or characteristic of an arbitrary process.

The use of the first rule requires a trade-off between the frequency selectivity of the neuron and the ability to detect and determine the amplitude of the incoming signal. In general, resetting z introduces a coarse frequency resolution, since the

differences of the voltage-like variable Im z between resonant and non-resonant frequencies increase as the development of the oscillation persists (see Figure 1). A low threshold and a resetting behaviour produces a broad range of frequencies for spike generation by a neuron. In contrast, with a high threshold, a resonant excitation with a low amplitude is unable to sufficiently excite the resonating neuron to reach the threshold. Signals with non-resonant frequencies are also incapable of sufficiently exciting the neuron, even with a large amplitude. A variable threshold enables detection of signals of small amplitude and attainment of voltage levels necessary for the development of large voltage differences by resonant and non-resonant excitations. Larger differences are achieved by not resetting z after the spike generation.

In addition to the number of spikes generated, the timing of and between the spikes contains information. The time between two spikes generated by the same neuron for resonant excitation is directly linked to the amplitude of the incoming signal. The phase of the signal also slightly influences the timing of outgoing spikes, since the outgoing spike pattern shows a phase-locked behaviour. While the envelope of the oscillating voltage-like variable is not influenced by phase changes, the oscillation shifts, moving the high points that are likely to surpass the threshold. Consequently, spike times are constantly shifted by $\Delta t = \frac{\phi}{\omega_0}$. This effect is small compared to timing shifts from amplitude variations scaling the spike time linearly with the SNR (see (5)). A spike generated after ten oscillations of a given sinusoidal signal requires a ratio between signal and amplitude noise of above 50 dB to achieve a smaller time shift than by a phase shift of ten degrees. Hence, the effect is only usable for signals that are unaffected by noise. But even without noise, the shift is only fractions of the duration of one oscillation.

The damping constant d determines the settling time of the resonant oscillation as well as its steady-state amplitude. The constant should be small enough to ensure that the steady-state is not reached during the measurement time T. Therefore, the growth of the envelope of the oscillating membrane voltage is approximately linear (see (5)). As the measurement time increases, the voltage difference between the resonant and non-resonant excitation increases (see Figure 1).

B. Frequency, Amplitude and Phase Resolution

Depending on the rule used for resetting and spike generation, the neurons differ in their selectivity of the input signal's frequency components near their resonant frequency. To discuss the frequency, amplitude, and phase resolution, a comparable metric covering the activity of a neuron must be defined. Since information is contained in the number and precise timing of spikes emitted, the metric must consider both. In general, a neuron emitting many spikes is more active than one producing only a few. Additionally, small latencies between a stimulus onset and the first spike emission as well as between successive spikes contribute to a higher activity. Depending on the spike

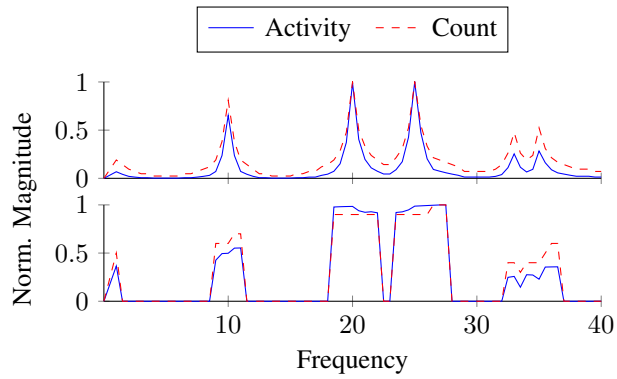

Fig. 2. Comparison of normalized spike count and activity for different spike generation and reset rules. Top: variable threshold without reset, bottom: fixed threshold with reset. Input is a sum of sines with frequencies 0, 10, 20, 25, 33, and 35 Hz and amplitudes 0.4, 0.8, 1, 1, 0.5, and 0.5.

times, we define the activity α_j of a neuron j as

$$\alpha_j = \sum_{n=1}^{N_j} n\,(T - t_n) \tag{8}$$

where N_j denotes the maximum number of spikes emitted by neuron j during the measurement time T. The variables t_n and n represent the time and index of the specific spike events of the neuron in ascending order, respectively.

Figure 2 displays a comparison of spectral analyses using RF neurons with different threshold and reset rules. Additionally, the absolute number of spikes as well as the activity metric (8) are depicted. For a variable threshold without reset, the introduction of the activity metric causes a sharper frequency selectivity and smaller activity values of near-resonant neurons. Generally, the spikes emitted by resonant neurons are early and therefore significantly influence the activity metric. Instead of sharp peaks at resonant frequencies, neurons with a fixed threshold and membrane voltage reset after spike emissions show a constant number of spikes and activation value for a broader frequency range. Neurons with a near-resonant frequency exhibit a reaction identical to the resonant neuron up to a specific frequency difference. In traversing this area, the spike rate instantly drops to zero. Moreover, the number of spikes and the activation metric show no improved differentiability. Consequently, the spike timings are similar for resonant and near-resonant neurons.

An important trade-off during the parametrisation of the resonating neuron exists between the amplitude range of the input and the frequency resolution. While ensuring a large dynamic range of the input at which oscillations are detected, the frequency resolution becomes coarse. While low thresholds enable the detection of low amplitude oscillations, high amplitudes eventually saturate and near-resonant oscillations also attain high levels of excitation. Low thresholds on the contrary render the detection of small amplitudes impossible, but enhance discrimination between resonant and near-resonant oscillations at higher amplitudes. The following relations are

used as a bases for deciding on the threshold:

$$\text{threshold} > \frac{a_{\max} T}{2N} \quad \text{and} \tag{9}$$

$$\text{threshold} < \frac{a_{\min} T}{2} \tag{10}$$

where $a_{\min}$ and $a_{\max}$ denote the detectable amplitude range and N is the maximum number of spikes during a measurement of duration T. As described previously, (9) ensures that the spike rate of the neuron does not saturate at high signal amplitudes. Equation (10) in contrast, enables the resonant oscillation of a signal with low amplitude to reach the threshold once during the measurement. For near-resonant oscillations without an exact frequency covered by a neuron to be detected, this threshold must be even lower. To specify the probability of spikes generated for non-resonant signals at neighbouring natural frequencies, an additional rule is formulated and expressed as

$$\text{threshold} > \frac{a_{\max}}{\Delta\omega}. \tag{11}$$

This prevents spike generation for an oscillation with frequency ω at a neuron with natural frequency $\omega+\Delta\omega$. For more neurons and a larger footprint of the circuit, multiple neurons with the same natural frequency but different thresholds can be realised to expand the dynamic range of detection. The setting of the threshold alters only the spike behaviour of the neuron, with the overall membrane voltage resonance remaining unchanged. The dependence on near-resonant frequencies with $\frac{1}{\Delta\omega}$ as in (7) can not be influenced.

For neurons with a fixed threshold, the amplitude range is primarily influenced by the minimum detectable signal amplitude (10) and the frequency selectivity (11). Comparing both relations produces the expression

$$\frac{a_{\max}}{a_{\min}} = \pi \Delta f\, T. \tag{12}$$

This means, if no spike must be generated at neighbouring frequencies with difference $1/T$ having a signal with amplitude $a_{\max}$, an amplitude range of 10 dB is achieved. If the selectivity is less strict, e.g. $\frac{2}{T}$ or $\frac{4}{T}$, the range reaches 16 dB or 22 dB, respectively.

Neurons with a variable threshold are primarily restricted by the maximum number of spikes generated during a measurement (see (9)). By comparing the relation to (10), the relation between the minimum and maximum signal amplitude is

$$\frac{a_{\max}}{a_{\min}} = N. \tag{13}$$

Thus, if 100 spikes are generated during the measurement, an amplitude range of 40 dB is achieved. Still, a very wide range also causes inaccurate frequency selectivity.

C. Comparision to FFT

An overview over the properties and limitations of the FFT and the neuronal-based approach is given in Table I. The main differences are highlighted in the following.

1) Energy Consumption: One of the largest benefits of spiking neural networks is their high energy efficiency [19]. The main consumers of energy in our resonate-and-fire neurons are the bias currents which enable the oscillation as well as the generation of the outgoing current pulses. Based on the proposed realisation by [10], 512 neurons require a total energy of 0.2 nJ to achieve time constants in order to analyse signals up to 12.5 MHz as needed later in section IV. [20] present a circuit design consuming 0.47 pJ per spike generated including the energy needed to drive axons with a total length of 1 mm. During the radar application example, 100 spikes were generated with a total energy consumption of ≈ 50 pJ. This total consumption of 0.25 nJ is much lower than current digital FFT blocks. Not considered are the savings introduced by omitting the ADC and further digital blocks if the whole processing is realised in neuromorphic hardware.

2) Frequency Resolution: In the FFT and neuronal approaches, the ability to distinguish similar frequencies scales with increasing measurement time. The frequency resolution of the FFT is influenced by the sampling frequency and the number of samples measured. For a given sampling frequency, the resolution is proportional to the number of samples. Additionally, due to the Nyquist criterion, the sampling frequency determines the highest measurable frequency of the signal without aliasing. However, the accuracy of the frequency decomposition decreases for oscillations with frequencies of non-integer multiples of the reciprocal of the measurement time due to spectral leakage. Resonating neurons show similar leakage of oscillations into neighbouring near-resonant neurons. The neurons develop higher amplitude differences between resonant and non-resonant frequencies with longer measurements (see (6)). Both effects, leakage at FFT and resonators, have different causes but similarly scale with $\frac{1}{\Delta\omega}$.

3) Amplitude Resolution: In comparison to FFTs with sample sizes of 12 bit and more, the amplitude range of the proposed neuromorphic approach is poor. Due to the event-based encoding of the weight of the spectral components in spikes however, the resolution within the measurable range is high. Since the information is represented by the timing of the spikes, arbitrarily small amplitude differences can be encoded.

4) Footprint Scaling: The FFT is scaled by the number of samples or the bit width of each sample point. The former is associated with a longer measurement duration while keeping the sampling time constant or by having a higher sampling rate at a constant measurement time, whereas the latter influences the amplitude resolution. Both contribute to the footprint size of the produced circuit because more or less data must be cached and processed. Depending on the application, it is also sufficient to compute the spectral transform only for a few specific frequencies of the spectrum. The most popular algorithm for this selective spectral transform is the Goertzel algorithm, that is more efficient than a block-wise solving of the transform, if $M \leq \frac{5}{6}\log_2(N)$. Here, M and N denote the number of the computed frequency components and the total number of samples, respectively. The resonating neuron approach scales in a similar way, and with a longer measurement duration, larger amplitude differences develop between resonant

and non-resonant frequencies. Consequently, narrower spaced frequencies can be differentiated. To exploit that, a larger number of neurons can be used to analyse the signal while maintaining the influence of near-resonance frequencies of the now closer lying natural frequencies. Equivalent to the higher word length in the digital domain, the use of more neurons to broaden the detectable amplitude range directly scales the footprint size of the network. In contrast to the FFT, the number of analysed frequencies scales linearly with the number of neurons needed. Thus, if only distinct frequencies must be evaluated, a significant chip area is saved.

The common silicon size of a 1024-point FFT implementation using 10 bit samples is 0.7 mm^2 for a 40 nm process. This area includes the necessary memory and logic blocks without considering ADCs and other arithmetic blocks. Depending on their natural frequency, the footprint of an RF neuron is primarily dominated by the size of the capacitors which determine the neuron's time constants. The additional logic only accounts for one third of the occupied area. A rough estimation of the area of neurons up to 12.5 MHz is 40 $\mu\mathrm{m}^2$ per neuron based on the process-dependent capacity density of 3.2 $\mathrm{fF}/\mu\mathrm{m}^2$. The resulting area for all 512 neurons is 20 mm^2, again using a 40 nm process.

Obviously, the larger size of the neuronal approach is less attractive since this directly determines the economic efficiency of the realisation. Next to more suitable fabrication processes which can shrink the size of the circuit significantly, some functional considerations, however relativise the difference. If the entire processing chain is realised using spiking neurons, as already mentioned, the ADC and further digital blocks can be eliminated. Furthermore, due to the linear complexity, the frequencies to be evaluated can be chosen arbitrarily. For example, neurons at frequencies of little relevance can be omitted or the spacing of the neurons natural frequencies can be adjusted freely.

IV. Evaluation: Radar Application

One application that directly relies on the accuracy of spectral analysis is the radar. Here, an electromagnetic wave is transmitted by the radar sensor, reflected by a target and the response received by the sensor. The waveform transmitted in frequency modulated continuous wave (FMCW) radar applications is a sinusoidal signal with linearly increasing frequency over time. By mixing the received signal with the currently transmitted signal, a sinusoidal intermediate frequency (IF) signal with a frequency difference between the two is generated for analysis. Since the transmitted signal's frequency increases linearly, the frequency difference between the transmitted and received signals is linearly dependent on the time of flight and consequently the distance of the reflecting target. The amplitude of the signal is influenced by the radar reflectance and the distance between the sensor and target. Therefore, only the frequency of the IF signal provides unambiguous information of the radial distance. Due to the high steepness of the frequency ramp, the Doppler-shift introduced by moving objects is negligible.

In common implementations, spectral analysis is performed using the FFT. This example demonstrates that it can also be achieved with the presented neuronal resonators. In an example FMCW radar with a base frequency of 77 GHz, HF and IF bandwidths of 1 GHz and 12.5 MHz, respectively, and a measurement time of 40 μs per ramp, a theoretical maximum range of 76 m and a range resolution of 0.15 m are achievable. The use of a sampling frequency of 25 MHz results in a 1024-point FFT. Accordingly, 512 RF neurons with evenly spaced natural frequencies between 0 and 12.5 MHz are to be analysed. To generate spikes, the simple fixed threshold and reset method was used. The result of the analysis of a simulated radar signal reflected at three targets is shown in Figure 3. On the signal, phase noise with an SNR of 19 dB is applied.

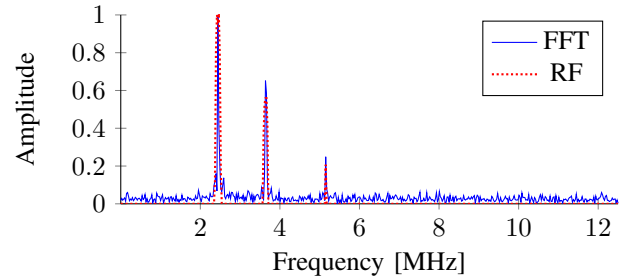

Fig. 3. Comparison of the activity metric of fixed threshold RF neuron and FFT for the simulated radar data. The signal is noisy with a signal-to-noise ratio of 19 dB. The RF neuron is able to detect the same frequencies while suppressing noise. The differences in their amplitude is caused by the choice of the firing threshold.

Depending on the chosen threshold, a clear decomposition of the radar signal is achieved using resonating neurons. With a predefined maximum amplitude range, the threshold is adjustable such that all relevant signals are recognised while simultaneously suppressing noise. Since the same measurement times are used, FFT and RF neurons show identical resolutions. Accordingly, due to the linear dependence between the IF signal frequency and distance, the range between the sensor and target is determined with the same accuracy.

To achieve a wider amplitude range for the RF neurons, adaption of the firing threshold is necessary. Therefore, the peaks become broader while still centred at the correct frequency or range bin. Consequently, objects which are close to each other become more difficult to distinguish. To cope with that necessitates the addition of more neurons with various thresholds or the use of a different spike generation rule.

As shown in section III-C, while being very low power, the silicon area needed to realise the neuronal input and spectral conversion stage is much larger compared to a classical FFT. In this particular use case one can consider different options to reduce the number of analysed frequencies. Depending on the application, low frequencies correspond to small distances and can be omitted if they are not relevant. Additionally, at longer distances, coarser distance measurements might be acceptable, resulting in less neurons at the corresponding frequencies.

V. Conclusion

This study presented using RF neurons to encode an analogue signal into spikes while performing a spectral transform simultaneously. The proposed method was compared to the

TABLE I
COMPARISON THE EFFECTS OF DIFFERENT PROPERTIES ON FFT AND RF NEURONS. N IS THE NUMBER OF SAMPLES, M IS THE NUMBER OF EVALUATED FREQUENCIES. THE SILICON AREA AND ENERGY CONSUMPTION ESTIMATIONS ARE BASED ON A 1024 POINT FFT AND A 40 NM PROCESS [20], [21].

PROPERTY	FFT	RF NEURON
sampling	Nyquist: $f_s \geq 2f_{\max}$	-
spectral leakage	spectral leakage at $f \neq n\frac{1}{\mathrm{T}}$	leakage of all signal components with $\frac{1}{\Delta\omega}$
frequency resolution	$\frac{1}{\mathrm{T}}$	$\approx \frac{1}{\mathrm{T}}$ (at resonator)
amplitude range	sample size	depending on spike generation
spectrum dimensionality	complex (magnitude and phase)	real (only magnitude usable)
reversibility	inverse FFT	-
complexity	NlogN (block solving) or NM (Goertzel)	M
silicon area	$0.7\,\mathrm{mm}^2$	$20\,\mathrm{mm}^2$
energy consumption	15 nJ	0.7 nJ

FFT for resolution, trade-offs, and scaling of the silicon implementation. We demonstrated that data from both methods were consistent. The encoding of the frequency components into spike events enabled further analysis of the signal in spiking neural networks, leveraging the advantages of neuromorphic hardware. The limitations of the proposed approach are a large silicon area needed to implement the circuits as well as a low measurable amplitude range. Future research will address these to provide an economically more attractive solution.

BIBLIOGRAPHY

[1] W. Maass, "Networks of spiking neurons: The third generation of neural network models," *Neural networks*, vol. 10, no. 9, pp. 1659–1671, 1997.

[2] E. M. Izhikevich, "Resonate-and-fire neurons," *Neural networks*, vol. 14, no. 6-7, pp. 883–894, 2001.

[3] B. Hutcheon and Y. Yarom, "Resonance, oscillation and the intrinsic frequency preferences of neurons," *Trends in Neurosciences*, vol. 23, no. 5, pp. 216–222, 2000.

[4] M. J. E. Richardson, N. Brunel, and V. Hakim, "From Subthreshold to Firing-Rate Resonance," *Journal of Neurophysiology*, vol. 89, no. 5, pp. 2538–2554, 2003.

[5] I. Lampl and Y. Yarom, "Subthreshold oscillations and resonant behavior: Two manifestations of the same mechanism," *Neuroscience*, vol. 78, no. 2, pp. 325–341, 1997.

[6] H. G. Rotstein, "Frequency Preference Response to Oscillatory Inputs in Two-dimensional Neural Models: A Geometric Approach to Subthreshold Amplitude and Phase Resonance," *The Journal of Mathematical Neuroscience*, vol. 4, no. 1, p. 11, 2014.

[7] M. St-Hilaire and A. Longtin, "Comparison of Coding Capabilities of Type I and Type II Neurons," *Journal of Computational Neuroscience*, vol. 16, no. 3, pp. 299–313, 2004.

[8] E. M. Izhikevich, "Which model to use for cortical spiking neurons?" *IEEE transactions on neural networks*, vol. 15, no. 5, pp. 1063–1070, 2004.

[9] T. Asai, Y. Kanazawa, and Y. Amemiya, "A subthreshold MOS neuron circuit based on the Volterra system," *IEEE transactions on neural networks*, vol. 14, no. 5, pp. 1308–1312, 2003.

[10] K. Nakada, T. Asai, and H. Hayashi, "A silicon Resonate-and-fire neuron based on the volterra system," in *Int. Symp. on Nonlinear Theory and Its Applications*, 2005, pp. 82–85.

[11] A. Tavanaei and A. S. Maida, "A spiking network that learns to extract spike signatures from speech signals," *Neurocomputing*, vol. 240, pp. 191–199, 2017.

[12] G. Bellec, D. Salaj, A. Subramoney, R. Legenstein, and W. Maass, "Long short-term memory and learning-to-learn in networks of spiking neurons," in *Advances in Neural Information Processing Systems*, 2018, pp. 787–797.

[13] G. Bellec, F. Scherr, A. Subramoney, *et al.*, "A solution to the learning dilemma for recurrent networks of spiking neurons," *bioRxiv*, p. 738 385, 2019.

[14] S. G. Wysoski, L. Benuskova, and N. Kasabov, "Evolving spiking neural networks for audiovisual information processing," *Neural Networks*, vol. 23, no. 7, pp. 819–835, 2010.

[15] Q. Sabatier, S.-H. Ieng, and R. Benosman, "Asynchronous event-based fourier analysis," *IEEE Transactions on Image Processing*, vol. 26, no. 5, pp. 2192–2202, 2017.

[16] T. N. Sainath, R. J. Weiss, K. W. Wilson, *et al.*, "Multichannel signal processing with deep neural networks for automatic speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 5, pp. 965–979, 2017.

[17] A. Graves, A.-r. Mohamed, and G. Hinton, "Speech recognition with deep recurrent neural networks," in *Acoustics, Speech and Signal Processing (Icassp), 2013 Ieee International Conference On*, IEEE, 2013, pp. 6645–6649.

[18] Y. Xu, J. Du, L. Dai, and C. Lee, "An Experimental Study on Speech Enhancement Based on Deep Neural Networks," *IEEE Signal Processing Letters*, vol. 21, no. 1, pp. 65–68, 2014.

[19] B. Rueckauer and S.-C. Liu, "Conversion of analog to spiking neural networks using sparse temporal coding," in *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*, 2018, pp. 1–5.

[20] J. M. Cruz-Albrecht, M. W. Yung, and N. Srinivasa, "Energy-efficient neuron, synapse and STDP integrated circuits," *IEEE transactions on biomedical circuits and systems*, vol. 6, no. 3, pp. 246–256, 2012.

[21] M. Seok, D. Jeon, C. Chakrabarti, D. Blaauw, and D. Sylvester, "A 0.27 V 30MHz 17.7 nJ/transform 1024-pt complex FFT core with super-pipelining," in *2011 IEEE International Solid-State Circuits Conference*, IEEE, 2011, pp. 342–344.