Undersampling and Saturation for Impedance Spectroscopy Performance

D.J. De Beer, *Member, IEEE*, and T-H. Joubert, *Senior Member, IEEE*

***Abstract*— The development of low-cost impedance spectroscopy devices is limited by the capabilities of low-cost hardware. Impedance spectroscopy requires high accuracy over a large bandwidth and dynamic range for most applications. In this paper, back-end signal processing solutions are examined and implemented to improve both the bandwidth and dynamic range of a typical low-cost impedance analyser without modifying the hardware. The bandwidth is improved by undersampling the input waveforms and reconstructing the signals based on the known frequency information. The bandwidth can be increased by multiple orders of magnitude and the limiting factor of the system becomes amplifier bandwidth and the Analog-to-Digital Converter (ADC) sample-and-hold circuit. Despite allowing signals to become orders of magnitude larger than the ADC range, the original signal can be reconstructed based on the measured saturated signal. This saturation technique more than doubles the dynamic range for any given accuracy requirement.**

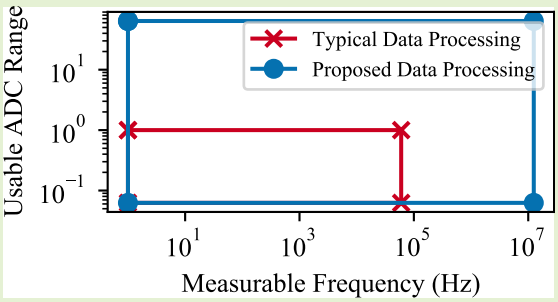

***Index Terms*— Impedance spectroscopy, Low-cost, Saturation, Undersampling**

I. Introduction

The term impedance spectroscopy describes the electrochemical measurement where the electrical impedance of a substance is measured across a set of electrodes [1]. The specific implementation can vary widely across different applications, ranging from highly accurate (and costly) lab-based measurement devices to microcontroller-based measurement devices with low cost and low power draw in mind. Fundamentally, however, the measurement methodology is the same: impedance is measured by applying either an AC voltage or an AC current across a set of electrodes and measuring the resultant current or voltage across a range of frequencies. The majority of methods measure both the voltage and current waveforms across the electrode, because it is not dependent on the ability of the analog front-end to maintain precise AC voltage across the electrodes, regardless of the load impedance.

In order for both signals to be measured by the same device, conversion between current and voltage is required, usually in the form of a transimpedance or transconductance amplifier. In order to determine the impedance, both the magnitude and the phase of each signal needs to be known. The most common way of determining these two metrics is by applying a Discrete Fourier Transform (DFT) to the sampled voltage and current channels which are digitised with an Analog-to-Digital Converter (ADC).

Many recent research efforts have focused on developing low-cost impedance analysers [2]–[9]. The trend in these works is a constant reduction in cost and power requirements. As the cost of these devices continues to be pushed lower and lower, there are performance metrics that need to be traded off with cost in order to achieve a low-cost device.

This paper will focus primarily on signal processing solutions that aim to improve the performance of sub-systems in an impedance analyser without incurring additional cost. The techniques will focus on the bandwidth and the dynamic range of the system and how they can be improved without additional hardware complexity. Some of the techniques and methods are well established in literature, but this work will focus on novel implementations as well as novel uses of established algorithms in low-cost impedance spectroscopy. The techniques will be designed and implemented in simulation, with results being acquired on a custom simulation model of a low-cost impedance device [10], with final validation being done with physical measurements. The block diagram for this system is shown in Fig. 1.

II. Low-cost Design Trade-offs

In order to apply signal processing techniques to the design of low-cost impedance analysers, various design trade-offs must be examined to identify potential gains that can be had by modifying the signal processing without incurring additional hardware cost. To this end, this section will examine the

This work is based on the research supported by the National Research Foundation. Any opinion, finding and conclusion or recommendation expressed in this material is that of the authors and the NRF does not accept any liability in this regard.

D.J. De Beer is with the Department of Electrical, Electronic and Computer Engineering at the University of Pretoria, Gauteng, South Africa (e-mail: dirk.debeer@tuks.co.za).

T-H. Joubert is with the Department of Electrical, Electronic and Computer Engineering at the University of Pretoria, Gauteng, South Africa (e-mail: trudi.joubert@up.ac.za)

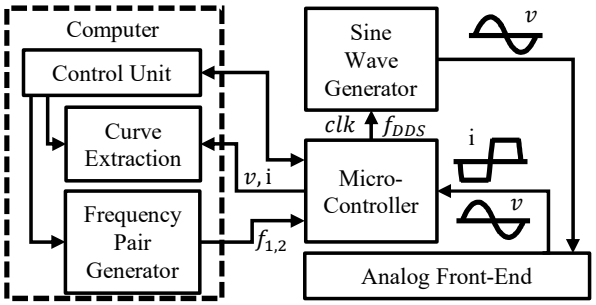

Fig. 1. Block diagram of a low-cost impedance analyser with the proposed back-end signal processing blocks.

various trade-offs that occur in low-cost impedance analyser design.

Low-cost design methodology results in fundamental design and performance limitations because all design requirements need to be carefully balanced with the cost of the device. In the case of low-cost impedance analysers, there are several design criteria that are part of the trade-off with cost namely:

- accuracy and precision,
- stability,
- speed,
- power,
- bandwidth,
- dynamic range, and
- back-end interface.

The first two design criteria, measurement accuracy and precision, are often emphasised in low-cost designs. Since impedance is a complex vector with both a magnitude and phase, there are accuracy and precision metrics for both. With a low-cost system there will be fundamental limitations to both these metrics because cheaper components will not be able to match the performance of expensive laboratory impedance analysers. Therefore, the majority of the cost of low-cost impedance analysers are usually focused on the components that will directly influence the accuracy and precision of the measurement [7].

System stability is a measure of the change in measurement accuracy as a consequence of external stimuli. As the stability increases, the effect of external stimuli decreases. The most prevalent stability consideration is temperature because most electronic components change their characteristics according to temperature. Unless an impedance measurement device has extreme temperature instability, or needs to be used in an extreme environment, it is typically not a focus of the design.

The speed of the measurement refers to the time it takes for the system to measure impedance at a frequency (or range of frequencies). Often, the time a system takes to measure a specific set of frequencies is used to compare between implementations [8].

The power draw of an impedance analyser is only critical for battery-powered (portable) devices as most wired devices would not consume enough power to warrant a specific design consideration.

Bandwidth for low-cost impedance spectroscopy is typically not discussed in depth in articles unless the system is specifically being designed for ultra-low [11] or ultra-high frequencies [12]. A range of 1 Hz to 10 MHz has emerged as a typical target range for impedance spectroscopy [8], with very few low-cost systems exceeding this range.

Dynamic range is another metric that is often not a primary concern in impedance analyser designs, but is vitally important in many applications. A system with high accuracy but a low dynamic range would work well for specific applications but could not be used at all in other applications without changes to the hardware. Thus a system with a high dynamic range can function as a general purpose impedance analyser, meeting the requirements for a variety of applications.

Finally, the back-end interface of an impedance analyser handles the setup and the handling of the data of the measurement. The measurement setup refers to the way in which a user can control the impedance measurements. Some high-end lab systems allow precise and complex control over the impedance measurements with proprietary software (including speed, voltage offsets and frequency resolution) while many low-cost systems require hard-coded changes to the embedded electronics to perform different measurements. Similarly the data handling and storage need to be handled correctly by the device. To this end, most low-cost implementations opt to include a connection to a computer so that the data can be logged and the device can be controlled [2], [3], [8].

From all the metrics that need to be considered for low-cost impedance analyser design, bandwidth and dynamic range are two that can be improved with the correct signal processing techniques after digitisation of the signals. These signal processing techniques and how they can be implemented in the back-end will be discussed.

III. Implementation

A. Bandwidth improvement through undersampling

One of the ways to improve the effective bandwidth of a limited ADC is to use undersampling. This technique ignores the Nyquist criteria of a digitally sampled signal and makes use of precise aliasing to get out-of-band signals to be measurable within the bandwidth of the ADC. The aliased signal can then be perfectly reconstructed with fewer than the required samples from the Nyquist-Shannon theorem due to the sparsity of the signal.

The method to reconstruct the signal is similar to the method of compressive sensing, but due to the simplicity of tones (single frequency signals) one does not need complex algorithms to reconstruct the signal. For most applications a simple DFT is sufficient to extract the magnitude and phase of a tone, even if it has been undersampled. Fig. 2 shows a simple undersampling measurement where two tones are synchronously sampled below their Nyquist rates. The figure shows that the magnitudes of the signals, and the relative phase between them, is preserved through undersampling, both of which are critical for impedance measurements.

Undersampling works because of the way the frequency spectrum of time-continuous signals is changed when sampled in discrete-time intervals. The discrete-time sampling results in the frequency spectrum being repeated in the frequency

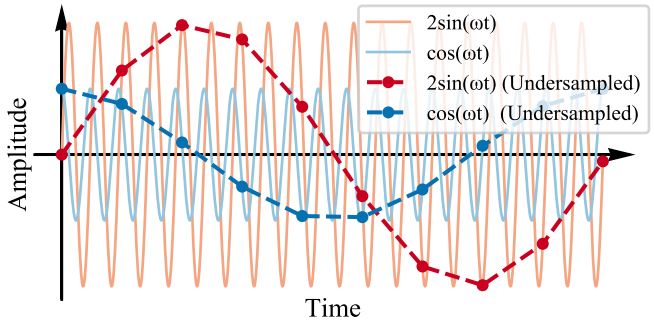

Fig. 2. Amplitude over time for two sinusoids before and after undersampling. The figure shows that undersampling changes the apparent frequency of the sinusoids, but magnitude and phase is preserved.

domain with a period that is equal to the sampling frequency. This results in the Nyquist frequency, which is the highest frequency that can be sampled at a chosen sampling rate without aliasing occurring. Any frequency that is higher than the Nyquist rate will alias down to a proportionally lower frequency in the discrete spectrum. This property allows undersampling to be used for sampling arbitrary high frequency signals, because they can be measured in the sub-Nyquist frequency band.

Undersampling does have limitations, for example the bandwidth of the signal being undersampled needs to be small enough to fit within the sub-Nyquist range without overlap. Consider a out-of-band signal with a bandwidth that is larger than a certain ADCs Nyquist frequency, as seen in Fig. 3. If such a signal is undersampled, then different parts of the signal will overlap in the sub-Nyquist frequency domain and the signal will be distorted, as can be seen in the outlined baseband signal. However, in the case of tone sampling one need not be concerned with some portion of the signal being distorted during aliasing because the bandwidth of the tone is very narrow.

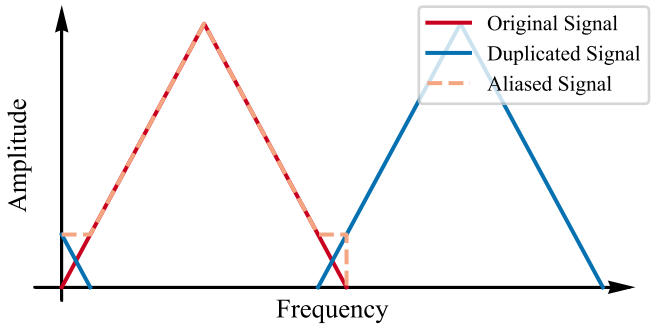

Fig. 3. Frequency spectrum of a sampled signal affected by aliasing. The duplicated signal spectrum which originated from sampling interferes with the spectrum of the original baseband signal.

When undersampling tones, some primary concerns are

- aliased noise,
- limitations of the ADC sample-and-hold circuit, and
- the impact of clock instability or drift.

The first contribution to aliased noise is the accumulation of background noise. This noise is typically reduced when using an anti-aliasing filter, which is primarily used to suppress signals with frequencies beyond the Nyquist frequency to avoid signal aliasing. Since undersampling requires out-of-band signals to be unfiltered, the high frequency background noise is not eliminated. A possible solution to this would be to apply a bandpass filter to the signal to be undersampled. This solution is not practical when the same system needs to sample tones that span large frequency ranges. The second contribution to aliased noise is the aliasing of specific noise sources, for example aliasing a desired signal down to 50/60 Hz would likely result in a noisy measurement as the noise generated by the power grid would overlap with the desired signal. Similarly, any noise that is concentrated in a specific frequency range can end up overlapping with a desired signal in the ADC frequency band if the undersampling ratio is chosen poorly. Therefore, to improve the performance of a system that makes use of undersampling, the band-limited sources of noise in the system needs to be known.

Even though undersampling makes signals appear in the digital space as low frequency signals they remain high frequency in the analog domain. Therefore, the sample-and-hold circuit in the ADC will start failing if the frequencies get too high. Since a sample-and-hold circuit requires capacitors to charge and discharge, if the charge and discharge times exceed the period of the signal, then the ADC will no longer be measuring the signals correctly, even if the undersampling theoretically allows it.

The last primary concern that was mentioned is the impact of clock jitter and drift in the undersampling system. This is not a concern when using an ADC normally, as the frequency of the driving clock of the ADC would be orders of magnitude higher than the signals being sampled by the ADC. When undersampling, however, one might measure signals with frequencies that are much closer to the frequency of the ADC clock. Thus, any instability in the ADC clock could manifest as noise in the measurement. Due to the nature of undersampling, the effect of clock jitter is magnified according to the undersampling ratio: the higher the frequency of the signal being measured, the higher the impact of clock jitter will be. The clock drift over time raises a different concern. Undersampling in impedance spectroscopy requires precise control over the frequency of the impedance signal as well as the sampling frequency. If these two frequencies drift independently then it could have a significant impact on the perceived frequency of the signal once it has aliased to baseband. Since the signal frequency is much higher than the sampling frequency, a small percentage change in the signal frequency causes a very large change at base band. It is therefore very important to ensure these two systems stay synchronised.

B. Precise Undersampling with Frequency Ratios

The target hardware system [13] wherein an undersampling technique will be implemented makes use of an AD9835 Direct Digital Synthesizer (DDS) to generate the waveform for impedance measurement and a dsPIC32 for controlling and sampling the signals. Both of these sub-systems are driven by the same system clock so that the ratios between the

frequencies can precisely be controlled without having to consider clock drift.

The low-cost impedance analyser is connected to a PC with a USB-UART chip. This allows a Python script to communicate with the dsPIC32, and in turn to control the DDS as well. The dsPIC32 can finely control the DDS frequency via a 32-bit number. The sampling frequency of the ADC, on the other hand, is internally modified by the dsPIC32 and cannot be set as precisely as the DDS. Nevertheless, with the numerous clock dividers and period registers, the ADC frequency can be sufficiently controlled to achieve precise undersampling ratios. The dsPIC32 allows simultaneous measurement of the two channels required for impedance measurement, and therefore the relative phase between the two signals can be determined regardless of the undersampling ratio.

To avoid loss of signal information, multiple periods of the wave - at least 20 according to Dudykevych et al. [2] - must be contained in a single measurement. The maximum ADC speed is only 1.2 Mbps, meaning that signals with frequencies up to 60 kHz can be sampled without the use of undersampling. By manipulating the sampling frequency along with the DDS frequency, the system can measure single frequency tones up to the maximum frequency of the DDS. In this case the maximum DDS frequency is 12.5 MHz. Using undersampling for this application increases the bandwidth by a factor of 208.3.

At the lower extreme of signal frequencies the measurement requires at least one full signal period to avoid loss of signal information. The same system of frequency ratios can be used to generate frequency pairs capable of measuring low frequencies beyond the extremely low frequency range (the ITU's lowest frequency designation) of 3 Hz while maintaining one or more period per set of 2048 samples. Some applications require frequencies in this range [14], [15]. Thus the system can take full advantage of the range of the DDS.

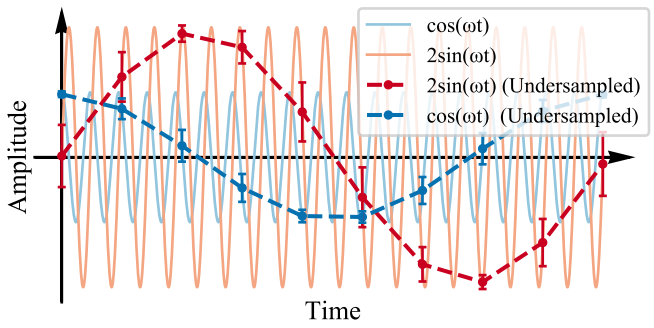

Fig. 4. Amplitude over time for two sinusoids before and after undersampling with the addition of clock jitter. The error bars on the undersampled signals show the standard deviation of the signal measurement when significant clock jitter is introduced.

At high frequencies the impact of clock jitter needs to be considered. The fact that the DDS and ADC are driven by the same clock somewhat reduces the impact of clock jitter due to the fact that the generated and sampled wave will have the same amount of clock jitter affecting the waveforms. Clock jitter can, however, manifest as phase-dependent noise on the signal being measured. Since the jitter is due to phase noise in the oscillator (operating at 50 MHz), the jitter will only manifest as significant signal noise when the signal being measured is at a high enough frequency so that the 50 MHz phase noise causes a significant random phase shift in each ADC sample. An example of the effects of clock jitter at very high undersampling ratios can be seen in Fig. 4.

A constraint of this technique is that precise waveform frequencies are sometimes compromised for better undersampling ratios. From empirical data it was determined that the saturation algorithm requires between 3 and 15 waveform periods in a set of 2048 samples. Therefore, at any given measurement frequency, the sampling frequency needs to be selected to ensure the number of periods is within the range. If no sampling frequency results in a valid number of periods then the measurement frequency is adjusted until a valid frequency pair is found. This constraint is more problematic at low frequencies because larger relative frequency adjustments are required to find valid frequency pairs. The severity of the constraint is affected by the granularity of sampling frequency adjustment.

C. Dynamic range improvement through saturation

Saturation describes the condition in which the amplitude of an input signal to any component exceeds the input range for that component. This is typically a condition to be avoided as any signal that exists beyond the saturation limit of the device is lost. At best it will result in a truncated signal and at worst it will cause damage to the device. Many impedance analysers make use of variable gain amplifiers to keep all signals within a range that can be measured by the ADC to avoid saturation. The downside to variable gain amplifiers is that they make the impedance analyser more expensive. The other issue is that signals will usually not make use of the full ADC range and the gain of the amplifiers will only be changed once the amplitude of the signal is sufficiently close to the limits of the ADC. This means that the ADC is only used with its maximum accuracy for a small fraction of measurements.

The field of compressive sensing covers, among other things, the accurate estimation of signals with significant saturation [16]. These methods require signal sparsity, which means the signal only has few non-zero coefficients in some of its transformation domains [17]. While there have been significant developments in this field, with many new and highly effective algorithms being developed [18]–[21], these methods are unnecessarily complex for saturated sinusoid recovery in impedance spectroscopy due to the nature of the impedance spectroscopy technique.

A pure sinusoid under a Fourier transform has only one (or two with a DC component) non-zero coefficient which definitely makes the signal sparse. The signal has so little information, however, that many of the compressive sensing algorithms would be excessive. At worst, a pure sinusoid has four unknown quantities (amplitude, phase, frequency and offset) which can be solved by a non-linear least-squares algorithm, or more optimally through frequency estimation [22], which makes the problem linear. In impedance spectroscopy however, the frequency is always known. Additionally, the

offset can be constant with certain hardware implementations. Therefore, in the best case, the least-squares algorithm can be limited to only solving for amplitude and phase. Two quantities that can also be initially estimated from the non-saturated region of the sinusoid.

Saturated curve extraction using least-squares regression is therefore uniquely well-suited to be used in impedance spectroscopy for dynamic range improvement. Saturated sinusoid recovery with least-squares regression, even with multiple frequencies [23] is not a novel concept. Leveraging the unique experimental setup of impedance spectroscopy to optimise the least-squares algorithm, however, is novel combination of two complementary techniques.

D. Saturated Curve Extraction

As discussed, despite saturation, the magnitude and phase of a tone can still be extracted accurately by predicting the shape of a sinusoid based on the truncated signal that is measured by the ADC. Thus one could deliberately let the input of the ADC be saturated and still make use of the full ADC range to determine the magnitude and the phase of the signal being measured. As seen in Fig. 5, a DFT can no longer be used for the magnitude and phase extraction because the amplitude of the fundamental will always be limited by the range of the ADC, which distorts the result under saturation. A custom signal processing method in the form of saturated curve extraction is proposed. Fig. 6 shows the curve that can be extracted from a saturated signal.

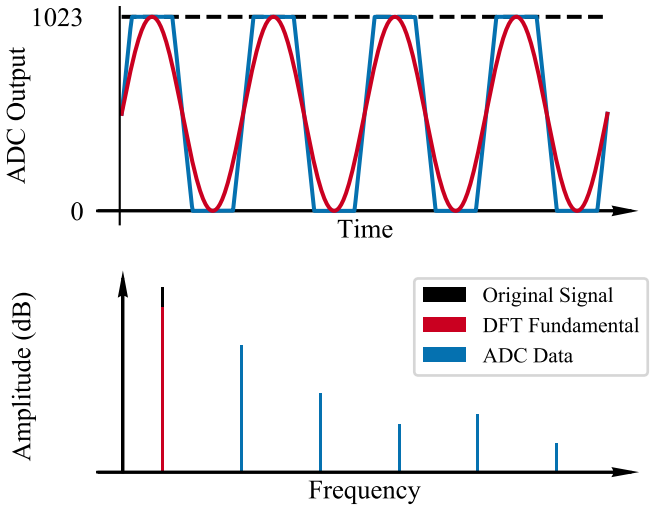

Fig. 5. Time and frequency domain representation of typical DFT feature extraction. The time domain plot shows the saturated signal as seen by the ADC along with the DFT component at the fundamental frequency. The frequency domain plot shows the original input signal, the DFT fundamental and the frequency spectrum of the signal seen by the ADC.

A Python script was developed to accurately determine the magnitude and phase of these truncated sinusoids. At the core of the script is a least-squares curve fitting algorithm that is used in a modified manner to fit sinusoids accurately despite the saturation. The complete algorithm has three distinct phases, namely: data preparation, value estimation, and least-squares regression.

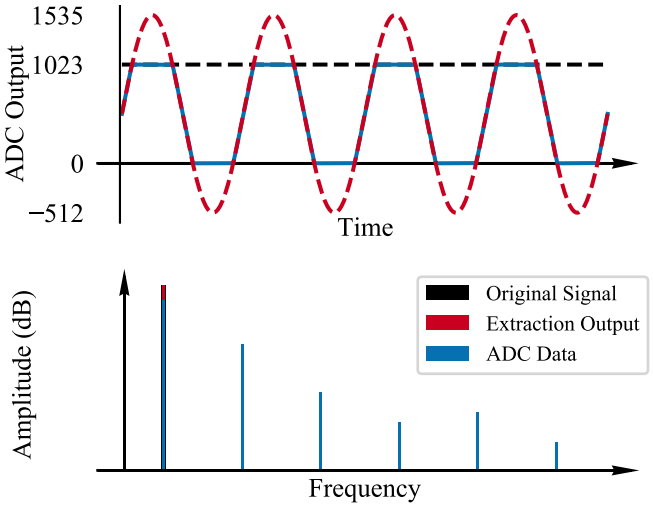

Fig. 6. Time and frequency domain representation of saturated curve extraction. The time domain plot shows the saturated signal as seen by the ADC along with the sinusoid that is extracted with the saturated curve extraction algorithm. The frequency domain plot shows the original input signal, the extracted signal and the frequency spectrum of the signal seen by the ADC.

Before the first step of the saturated curve extraction process, the Python script receives two sets of 2048 data points from the low-cost impedance device. These two sets represent the synchronous voltage and current data across the electrodes. As the frequency of the waves are known, the script can check the number of wave periods in the set to confirm that no data was lost during transmission. Thereafter, the waves are low-pass filtered to eliminate some noise, and truncated to contain an integer number of periods in the set. The truncation ensures the mean of the sinusoidal component of the data set is zero over the set, which improves the ability of the least-squares algorithm to correctly determining the offset.

After the data preparation step, the maximum derivative of the saturated wave is calculated to determine an estimate of the peak amplitude of the wave. This estimate is used as an initial value for the least-squares regression algorithm. In addition to this estimate, the offset and phase of the wave is estimated based on the point where the maximum derivative occurred because it should represent the vertical centre of the sinusoid.

With the amplitude, phase and offset estimated, a least-squares regression algorithm is used to hone in on the exact values of the three parameters simultaneously. The algorithm cannot, however, simply fit a sinusoid because of the saturation. Instead, a custom sinusoid function is used that exhibits the same saturation limits as the measured data. This is important as the least-squares algorithm needs to ensure that both the measured and derived signals saturate in a similar manner to arrive at an accurate estimation of the real values. The results of the least-squares algorithm is then used to determine impedance.

The process demands significantly more processing power than can be achieved with a low-cost microcontroller, but the processing power of a basic laptop or desktop computer is adequate for the application. The process can, however, be optimised for microcontroller implementations if remote

impedance analysers is the design goal.

IV. Results

A 10-bit ADC can measure values between 0 and 1023, meaning it can theoretically achieve a maximum dynamic range of 1:1000 with the smallest resolvable signal having a amplitude of $\frac{1}{2}$LSB. This range is not practically feasible because a signal of that amplitude cannot be measured with accuracy. To achieve an accuracy of 1% for a noiseless signal, the peak-to-peak signal amplitude would need to be 100 LSBs. With the saturation method, the theoretical maximum signal amplitude would be a signal where the difference between two consecutive samples exceeds the full scale ADC range. This amplitude limit is not a constant value because the rate of change in the non-saturated portion of the signal is dependent on frequency. For example, a noiseless low frequency signal could have an arbitrarily large amplitude due the low non-saturated signal slope. Practically, the amplitude of saturated signals is limited by the required accuracy: the noise on any sample used for the saturated curve extraction is multiplied by the ratio between the full ADC range and the amplitude of the saturated signal. Therefore, the dynamic range gain from using the saturated curve extraction is dependent on the noise present in the signal.

Fig. 7 shows the impedance error distribution for measurements using different portions of the ADC full scale range. The graph shows results from $\frac{1}{32}$ to 128 times ADC full-scale range. The simulation model includes noise on both the voltage and current channels. Based on empirical data, the voltage channel has −40 dB of noise while the current channel has −20 dB of noise. Since the noise on the current channel is significantly higher, it is the primary cause of measurement error.

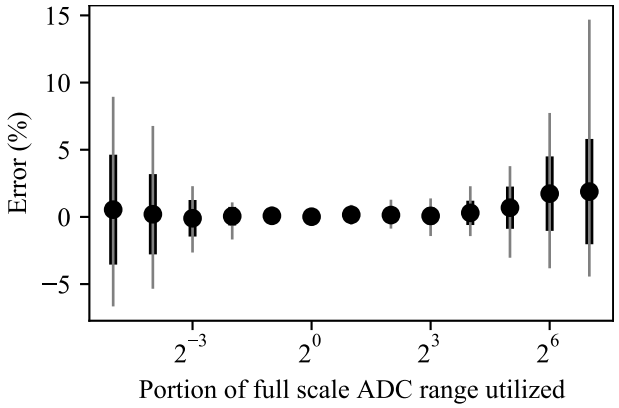

Fig. 7. Measurement accuracy for saturated and non-saturated measurements. The graph shows the average, extremes and standard deviation when measuring above and below the full ADC range with constant noise.

The standard deviation shown in Fig. 7 gives an indication of the level of accuracy achieved when averaging multiple repeated measurements, which is the procedure implemented on the physical device. By comparing the standard deviations of different points on the graph, one can compare the expected system accuracy at these points. For example, the point at 2^{-4} has approximately the same standard deviation as the point at 2^{6}, thus similar measurement accuracy can be achieved at $\frac{1}{16}$ of full scale and at 64 times full scale. Therefore, at these noise levels and this level of accuracy, the technique improved the dynamic range from 1:16 to 1:1024.

On the physical low-cost impedance analyser it was found experimentally that the system could reliably measure signals that were 50 times full scale in amplitude, this would suggest that the noise level in the system is slightly lower than the values used in the simulation model. This increases the dynamic range from 1:10 to 1:500, while maintaining an accuracy of 1%.

As there are no examples of saturated curve extraction for impedance spectroscopy to compare against, the performance of the algorithm is compared against a similar process used in the field of magnetic coupling. Hu et al. [23] achieved measurements up to 10 times full scale, while maintaining an accuracy of 0.6% with a similar noise level (-20 dB) as in this paper. Above this level, however, the results degrade significantly and the error increases to 3.5%. Thus, the algorithm presented here yields 0.4% worse performance up to their 10 times full scale mark. However, this current work is significantly better than the prior work in terms of the dynamic range improvement, because an accuracy of 1% is maintained beyond the 10 times full scale mark - up to 50 times full scale. Despite the reduced accuracy, the work presented here yields significantly improved results in terms of the primary design goal, namely dynamic range improvement.

Impedance measurement that spans a wide bandwidth and large dynamic range was experimentally investigated using the low-cost board for both the proposed saturated curve extraction approach and a typical DFT approach. The load being measured is shown in Fig. 8 and features three resistors and two capacitors, designed such that the impedance varies between three distinct levels as the frequency changes.

Fig. 9 presents the comparative impedance measurement results, relative to the expected impedance obtained by simulation. The key benefit of the newly proposed approach is evident in the increased dynamic range for frequencies larger than 1 kHz. Accurate measurement of the impedance is validated for the curve extraction approach up to frequencies when the on-board amplifier limits become significant, as seen beyond approximately 600 kHz in the figure. The board features low-cost amplifiers that exhibit slew rate limitations when the product of the frequency and the current is high.

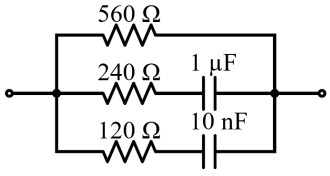

Fig. 8. Resistive and capacitive load for impedance measurement.

V. Conclusion

In this paper it was shown that certain techniques can be incorporated into low-cost impedance analysers in order to

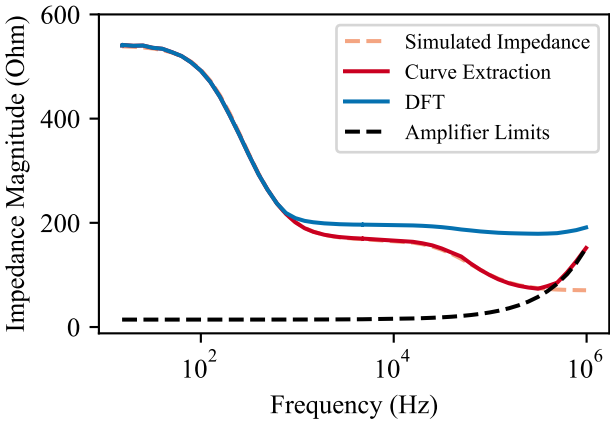

Fig. 9. Impedance measurement comparison between a typical DFT approach and the proposed saturated curve extraction approach.

improve their performance without increasing the production cost of the device.

By making use of undersampling, the bandwidth of the impedance analyser can be extended significantly: with the presented implementation, at least a 200-fold increase in the typical bandwidth of the ADC is achieved. Additionally, by precisely controlling the sampling frequency, more precise measurements at extremely low frequencies are possible, which also improves the bandwidth at the low end.

The paper also showed that by deliberately saturating signals, the full ADC range can be used to measure signals up to 50 times larger than the full scale of the ADC. It was also shown that the presence of noise will limit this improvement factor, nevertheless, a significant increase in the dynamic range of a typical ADC is still attainable.

There are a few limitations to this method:

- One of the most important limitations for a low-cost system is the processing required for the saturated curve extraction. Therefore, an embedded solution would be more expensive than a simple impedance analyser that only makes use of a DFT.
- The second limitation is the fact that the equivalent signal amplitudes are larger than the supply voltage - even though the signals cannot surpass the rails, the ”large” signals have an equivalently large derivative that may cause slew rate problems with the amplifiers in the analog front-end. It was found experimentally that the maximum saturated signal amplitude is inversely proportional to the frequency of the waveform. This directly correlates to the slew rate of the amplifier because it can only handle smaller signals at higher frequencies. The slewing of the amplifier also results in a phase delay that is apparent in measurements. The slew rate issues, however, does show that the maximum amplitude of the signal is not typically limited by the ADC, but rather the other hardware in the system.
- Finally, it is important to note that nowhere on the PCB does a real signal actually exist that is 50 times the ADC range. This is due to the fact that the rails of the analog front-end are approximately equal to the range of the ADC. Thus the front-end amplifier is also saturated.

Future work will include further development of the undersampling technique in order to select better undersampling ratios without having to make adjustments to the measurement frequency. Furthermore, the impact of saturated amplifiers on the analog front-end needs a more in depth investigation to determine whether or not it has a significant impact on the impedance being measured.

Acknowledgment

This work is based on the research supported by the National Research Foundation. Any opinion, finding and conclusion or recommendation expressed in this material is that of the authors and the NRF does not accept any liability in this regard.

References

[1] E. Barsoukov and J. R. Macdonald, *Impedance Spectroscopy: Theory, Experiment, and Applications*, 2nd ed. New Jersey: Wiley-Interscience, 2005.

[2] T. Dudykevych, E. Gersing, F. Thiel, and G. Hellige, “Impedance analyser module for EIT and spectroscopy using undersampling,” in *Physiological Measurement*, vol. 22, no. 1, 2001, pp. 19–24.

[3] M. Michalikova and M. Prauzek, “A hybrid device for electrical impedance tomography and bioelectrical impedance spectroscopy measurement,” in *Canadian Conference on Electrical and Computer Engineering*. Institute of Electrical and Electronics Engineers Inc., Sep. 2014.

[4] J. Ojarand, M. Min, and A. Koel, “Multichannel electrical impedance spectroscopy analyzer with microfluidic sensors,” *Sensors*, vol. 19, no. 8, Apr. 2019.

[5] M. Grossi, C. Parolin, B. Vitali, and B. Riccò, “Electrical impedance spectroscopy (eis) characterization of saline solutions with a low-cost portable measurement system,” *Engineering Science and Technology, an International Journal*, vol. 22, no. 1, pp. 102–108, Feb. 2019.

[6] L. Breniuc, V. David, and C. G. Haba, “Wearable impedance analyzer based on AD5933,” in *EPE 2014 - Proceedings of the 2014 International Conference and Exposition on Electrical and Power Engineering*. Institute of Electrical and Electronics Engineers Inc., Dec. 2014, pp. 585–590.

[7] C. E. F. do Amaral, H. S. Lopes, L. V. Arruda, M. S. Hara, A. J. Gonçalves, and A. A. Dias, “Design of a complex bioimpedance spectrometer using DFT and undersampling for neural networks diagnostics,” *Medical Engineering and Physics*, vol. 33, no. 3, pp. 356–361, Apr. 2011.

[8] R. Munjal, F. Wendler, and O. Kanoun, “Embedded Wideband Measurement System for Fast Impedance Spectroscopy Using Undersampling,” *IEEE Transactions on Instrumentation and Measurement*, vol. 69, no. 6, pp. 3461–3469, Jun. 2020.

[9] D. J. De Beer, T.-H. Joubert, P. H. Bezuidenhout, and M. du Plessis, “CMOS-based impedance spectroscopy for water quality monitoring,” in *Fifth Conference on Sensors, MEMS, and Electro-Optic Systems*, M. du Plessis, Ed., vol. 11043. SPIE, 2019, pp. 449–456.

[10] D. J. De Beer, “Impedance spectroscopy for water quality monitoring,” University of Pretoria, Pretoria, Tech. Rep., Nov. 2018.

[11] J. A. Giacometti, N. Alves, and M. Y. Teruya, “Impedance of Aqueous Solutions of KCl at the Ultra-low Frequency Range: Use of Cole-Cole Impedance Element to Account for the Frequency Dispersion Peak at 20 mHz,” *Brazilian Journal of Physics*, vol. 46, no. 1, pp. 50–55, Feb. 2016.

[12] X. Ma, X. Du, H. Li, X. Cheng, and J. C. Hwang, “Ultra-wideband impedance spectroscopy of a live biological cell,” *IEEE Transactions on Microwave Theory and Techniques*, vol. 66, no. 8, pp. 3690–3696, Aug. 2018.

[13] D. J. De Beer and T. H. Joubert, “Impedance Spectroscopy for Determination of Total Dissolved Solids in Aqueous Solutions of Sodium Chloride and Magnesium Sulphate,” in *IEEE SENSORS*, 2019, pp. 1–4.

[14] M. Murbach, V. Hu, and D. Schwartz, “Nonlinear electrochemical impedance spectroscopy of lithium-ion batteries: Experimental approach, analysis, and initial findings,” *Journal of The Electrochemical Society*, vol. 165, pp. A2758–A2765, Jan. 2018.

[15] S. Feliu, "Electrochemical impedance spectroscopy for the measurement of the corrosion rate of magnesium alloys: Brief review and challenges," *Metals*, vol. 10, no. 6, 2020.
[16] J. Haboba, M. Mangia, F. Pareschi, R. Rovatti, and G. Setti, "A pragmatic look at some compressive sensing architectures with saturation and quantization," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 2, pp. 443–459, 2012.
[17] L. Stanković, E. Sejdić, S. Stanković, M. Daković, and I. Orović, "A tutorial on sparse signal reconstruction and its applications in signal processing," *Circuits, Systems, and Signal Processing*, vol. 38, pp. 1206–1263, Mar. 2019.
[18] Z. Yang, L. Xie, and C. Zhang, "Variational bayesian algorithm for quantized compressed sensing," *IEEE Transactions on Signal Processing*, vol. 61, pp. 2815–2824, 2013.
[19] S. Foucart and J. Li, "Sparse recovery from inaccurate saturated measurements," *Acta Applicandae Mathematicae*, vol. 158, pp. 49–66, Dec. 2018.
[20] F. He, X. Huang, Y. Liu, and M. Yan, "Fast signal recovery from saturated measurements by linear loss and nonconvex penalties," *IEEE Signal Processing Letters*, vol. 25, pp. 1374–1378, Sep. 2018.
[21] Y. Liu, J. J. Zhu, N. Roberts, K. M. Chen, Y. L. Yan, S. R. Mo, P. Gu, and H. Y. Xing, "Recovery of saturated signal waveform acquired from high-energy particles with artificial neural networks," *Nuclear Science and Techniques*, vol. 30, pp. 1–10, Oct. 2019.
[22] J. Q. Zhang, Z. Xinmin, H. Xiao, and S. Jinwei, "Sine wave fit algorithm based on total least-squares method with application to adc effective bits measurement," *IEEE Transactions on Instrumentation and Measurement*, vol. 46, pp. 1026–1030, 1997.
[23] C. Hu, W. Qu, Z. Zhang, Z. Feng, S. Song, and M. Q. Meng, "Recovering amplitudes and phases from saturated multifrequency sinusoid signals," *IEEE Sensors Journal*, vol. 13, pp. 4569–4575, 2013.