IEEE Conf. on Norbert Wiener in 21st Century, Thinking Machines in the Physical World, Melbourne, Australia, 13-15 July 2016.

An Output-Only Nonlinear System Identification Technique Suited to Integer Arithmetic

Lachlan J. Gunn*, François Chapeau-Blondeau†, Andrew Allison*, Derek Abbott*
*School of Electrical and Electronic Engineering, The University of Adelaide, Adelaide, Australia 5005
Email: {lachlan.gunn, andrew.allison, derek.abbott}@adelaide.edu.au
†Laboratoire Angevin de Recherche en Ingénierie des Systèmes (LARIS),
University of Angers, 62 avenue Notre Dame du Lac, 49000 Angers, France
Email: chapeau@univ-angers.fr

***Abstract*—In most designs, residual nonlinearity is considered an inescapable curse—even when it is known to be present, it is often assumed to be too unpredictable or unstable to be dealt with in postprocessing. However, with the aid of output-only system identification, this is no longer the case. We have developed a new technique for compensation of static nonlinear distortion using the internal noise of the device. It improves upon previous approaches by allowing highly-efficient fixed-point implementation, and represents the first step towards direct integration with analog hardware in order to produce an ADC that is blind to its analog frontend.**

I. Introduction

Measurement systems suffer from a wide variety of faults, and a range of techniques have been developed to overcome them. Perhaps the most well-known of these is the Wiener filter [1], which uses linear filtering to provide an estimate of the input signal despite the presence of noise. Despite Wiener's important work in nonlinear system representation [2], nonlinearity is all too often seen as a lost cause—if it cannot be mitigated using feedback, then in most circumstances no attempt is made to correct it in postprocessing owing to a lack of stability or repeatability. Our goal is the same as that of Wiener—to reconstruct a contaminated signal given some knowledge of its basic properties—but extended to nonlinear distortion.

Filtering is often used [3] to mitigate distortion in RF systems, but this is only viable for narrowband signals—in a wideband system, the harmonics and intermodulation products will overlap with the desired signal. In order to overcome these difficulties, some sort of calibration system is necessary. Most techniques for nonlinear system identification [4] require some form of test signal with known properties [5]–[7]; whether deterministic or random, it must be known. However, the use of such an input signal can be difficult and expensive, or for sensor applications completely impossible. Performing output-only identification of the nonlinearity presents a substantial challenge, however opens up a new area of the design space by allowing the use of highly nonlinear devices with otherwise desirable properties.

Others have attempted automatic and signal-independent compensation techniques for analog-to-digital converters (ADCs) [8], however this attempt was suitable only for gross single-bin nonlinearities, attempting to expand or contract certain bins in order to render the signal histogram continuous. This clearly is not completely signal-independent, however for non-degenerate signals this may be a reasonable approach. Another attempt [9] used modifications to the input-output configuration of a pipelined ADC in order to produce uniformly distributed noise—however, this comes with the disadvantage of requiring that the ADC be disconnected from the circuit. In this present paper, we search for an alternative method that can be performed from the measurement data alone, allowing it to be used without taking the measurement device offline, and enabling us to apply the method to data that has previously been collected.

Others have shown that, the presence of noise can have a linearising effect [10] in certain situations. This naturally-occuring linearisation raises the question of whether a more general approach can produce an even greater effect.

We have previously described a technique [11] for the compensation of static nonlinearity using output-only measurements, and succeeded in refining the concept into an algorithm capable of operating in real-time on a microcontroller platform [12]. These successes have motivated us to search for other approaches that will allow us to adapt the technique to smaller devices and to VLSI implementation. We present a technique here that can be efficiently implemented using fixed-point arithmetic without the need for either divisions or square-root operations as required by [12].

A. Principle of operation

Without any *a priori* information on the distorting function, we must make some assumption about the input signal in order to determine its properties. In this case, we make the assumption that the variance of the input's high-frequency content is small and constant with time. This could be some sort of deterministic signal, but in our experiments [11], [12] we have used the input noise of a common-emitter amplifier. By using a high-pass filter to separate this from the large and slowly-varying low-frequency signal, we effectively see the effect that the system has had upon the noise at a variety of operating points.

If the distorting function is smooth, we can consider a linearisation about each operating point; the effect then will be for the noise to be compressed when the system is saturated.

Denoting the input $x(t)$ the output $y(t)$, the input noise $n(t)$, and the distorting function $f(x)$, we may write

$$y(t) = f\left(x(t) + n(t)\right) \tag{1}$$
$$\approx f\left(x(t)\right) + n(t) f'\left(x(t)\right). \tag{2}$$

For notational clarity we assume that $n(t)$ is high-pass and that we can thus separate the second term via filtering, yielding $z(t)$; this gives us

$$E\left[z(t)^2\right] = E\left[n(t)^2\right] f'\left(x(t)\right)^2, \tag{3}$$

and so a means to measure the differential gain $f'(x)$ of the distorting function $f(x)$ at a variety of operating points. We build up a map of differential gain vs. operating point, which we then integrate to find the distorting function $f(x)$ to within an offset factor. In practice, we do not necessarily know the variance of the input noise—just that it is constant—resulting in an ambiguous scale factor at the input. Thus this technique allows us to determine the input signal except for an unknown gain and offset.

In [11] we performed the integration directly using Simpson's rule. This is straightforward, but is inefficient for large numbers of samples and requires that one maintain a large number of samples in memory. Furthermore, it is unable to cope with transfer functions that change over time. We overcame both of these problems in [12] by integrating a curve-fitted version of $(f^{-1})'(x)$; a set of triangular radial basis functions were used to describe $f'(x)$, resulting in a piecewise-linear fitted curve, and a piecewise quadratic fit to the distorting function.

This method of implementation is relatively efficient, however it suffers from some drawbacks when used on resource-constrained platforms. Estimation of the derivative of the inverse distorting function $(f^{-1})'(x)$ requires the computation of $(\sigma^2)^{-\frac{1}{2}}$; this is a relatively expensive operation which must be carried out for every noise estimate, and limits the rate at which the differential gain may be measured.

A further problem with this approach is that it is necessary to explicitly perform the integration—while this can be calculated analytically in terms of basis function coefficients, it is a slow process that cannot be carried out all at once. This raises the question of whether it is possible to perform the transfer function update operation directly in the integrated domain.

II. Feedback-based transfer function estimation

In this paper we introduce a new technique based on negative feedback in order to appropriately determine the nonlinear compensating function. Let us suppose that we have a discrete-time quantized signal $y[n]$, which takes values from $\{0, \ldots, 2^{M-1}\}$, and we wish to compensate this to yield a time series $x[n] \in \{0, \ldots, 2^{N-1}\}$. This latter time-series is found by compensating $y[n]$ with a nonlinear function $f_n(x[n])$ that varies with time. We choose $f_n(\cdot)$ to be piecewise linear with

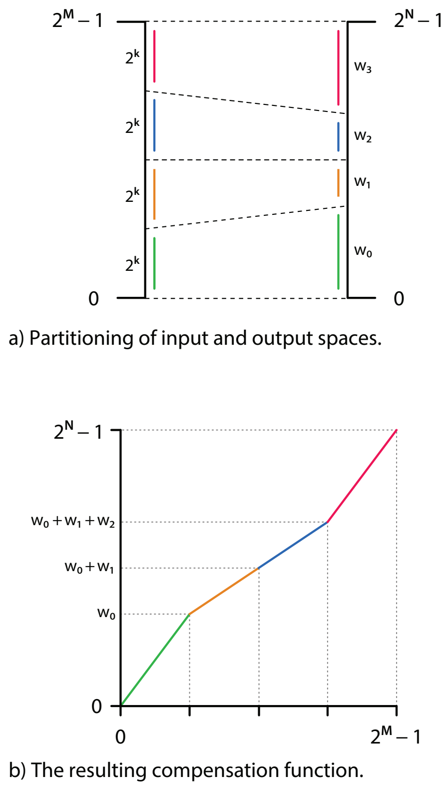

Fig. 1. The construction of the compensating function described in Equation 4. The input and output ranges are partitioned into 2^k segments—the input uniformly, and the output with arbitrary widths, shown in (a). Corresponding partitions are mapped linearly onto one another, resulting in the piecewise linear function shown in (b). By using powers of two for the segment widths, time-consuming division operations are avoided.

segments of size 2^k at the input. We may write this function as

$$f_n(z) = \left[\frac{w_m}{2^k}(z - m2^k)\right] + \sum_{i=0}^{m-1} w_i[n], \tag{4}$$

where

$$m = \left\lfloor \frac{z}{2^k} \right\rfloor \tag{5}$$

is the segment in which z lies, and the segment widths $w_i[n]$ in the output satisfy

$$\sum_{i=0}^{2^{M-k}-1} w_i[n] = 2^N; \tag{6}$$

that is to say, they cover the entire output range. The construction of this function is shown in Figure 1.

Previously we have measured the time-varying noise power at the input and then attempted to calculate the weights that would result in a stationary output process. In order to remove the need to calculate the necessary weights, we instead measure the noise power at the output, which allows us to use a simpler adjustment rule.

```
procedure UpdateSegmentWidths( x[n], σ²_noise[n] )

    if σ²_noise[n] < σ²_avg then
        Δ ⟵ +1
    else if σ²_noise[n] < σ²_avg then
        Δ ⟵ −1
    else
        Δ ⟵ 0
    end if

    UpdateNoiseAverage( x[n], σ²_noise[n] )

    i ⟵ FindSegment( x[n] )
    if w_i[n] + Δ ∉ {1, ..., 2^N−1} then
        return
    end if

    j ⟵ FindSegment( Random({1, ..., 2^N−1}) )
    if w_j[n] + Δ ∉ {1, ..., 2N − 1} then
        return
    end if

    w_i[n+1] ⟵ w_i[n] + Δ
    w_j[n+1] ⟵ w_j[n] − Δ

end procedure
```

Fig. 2. The width-update algorithm.

If the distorting function is exactly equal to the inverse of the compensating function $g(z)$, by our initial assumption the noise at the output will have a constant variance. If the variance is greater than average, this implies that the differential gain is also greater than average, and therefore we must reduce the gain of the compensating function in this region. Conversely, if the variance is less than average, we increase the gain of the compensating function.

We use a simple rule to determine the weight updates—if the noise is greater than average, the corresponding segment width $w_i[n]$ will be reduced by one, and if it is greater then $w_i[n]$ will be increased by one. However, after doing so the output-range constraint no longer satisfies Eqn. 6; if $w_i[n]$ has been reduced, its output bin must be allocated somewhere else, and if it has been increased, its output bin must be taken from somewhere. Our key innovation is, rather than performing a time-consuming global rescaling, to simply give or take a random output bin—a random number is selected from $\{0, \ldots, 2^N\}$, and the width of the corresponding segment is increased or decreased by one. This is shown in greater detail in Figure 2. A proof of convergence for this algorithm will be the subject of a later work.

It is also worth noting that the noise need not be explicitly measured at the output of the compensator; if measured at the input, it may be converted to an output-equivalent noise by scaling its power by w_i^2, where the signal falls within bin i. Doing so allows the use of hardware filters to separate the noise from the low-frequency signal, thereby substantially reducing the computational burden on the processor.

We compute the average noise power $y[n]$ using a simple infinite-impulse-response (IIR) filter of the form

$$y[n] = (1 - 2^{-10})y[n-1] + 2^{-10}x[n],$$

chosen for ease of implementation. It remains to be seen whether more sophisticated techniques, yielding a more representative noise average, will provide a substantial advance in performance. One approach would be to take the median output noise variance, which would cause the random-bin-reallocation process to give and take bins 50% of the time, however it is not yet known whether the removal of this source of bias is important.

III. Results

We test our system based on Equation 4 and UPDATE-SEGMENTWIDTHS. The method used to generate the test signals and measure the total harmonic distortion (THD) [5] of the output is shown in Figure 3. Ten million samples—ten seconds worth—were generated and processed, with the system being allowed five seconds to settle, after which the THD was computed from the remaining data points, up to the tenth harmonic. The inputs to the algorithm were quantised to fourteen bits, and the outputs quantised to twelve. We used 2^7 segments, thus yielding the parameters $M = 14$, $N = 12$, and $k = 7$ in the exposition above. Noise power was measured at the input and scaled by the square of the weights, yielding the output noise.

The total harmonic distortion (THD) of the output of the technique is shown in Figure 4. We see a substantial reduction in THD over a wide range of input distortion levels, however the achievable lower limit is determined by the block size; it is anticipated that refinements to the method will allow a reduction in the required block size whilst maintaining a low THD. The use of large block sizes appears to result in some 'peaking' at high levels of distortion, however at this level of distortion, the algorithm ceases to function, resulting in a net *increase* in distortion, and so would not be used, rendering the point moot. The results therefore indicate the use of larger block sizes. As this test is performed with a highly oversampled signal—by a factor of 5000—this is not a problem here. However, with lower oversampling ratios this design parameter may limit the achievable performance.

IV. Conclusion

We have described a new algorithm for online nonlinear system identification in electronic systems. The proposed system avoids expensive division and square-root operations, and we have produced a fixed-point proof-of-concept implementation. It has been shown in simulation to improve total harmonic distortion by between ten and twenty decibels for a $\tanh$-style distorting function, reaching a floor of almost 0.1% THD. The success of the new technique at low levels of distortion is significant, as it means that this noise-based approach is

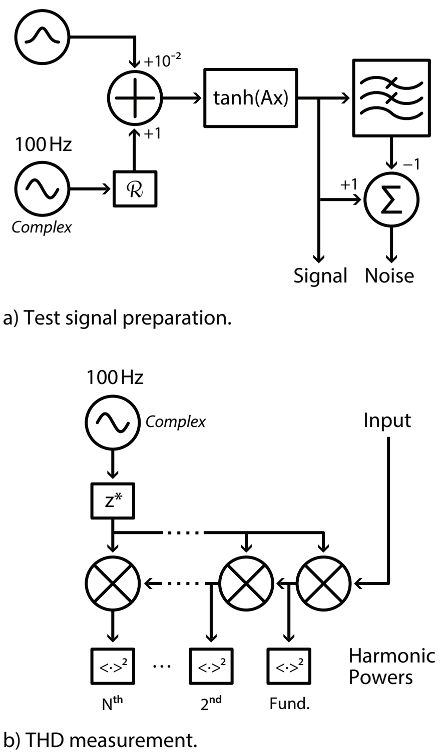

Fig. 3. Test setup for the anti-distortion system; the sampling rate is $1\,\mathrm{MHz}$. The Gaussian noise was generated as a repeating sequence of 2^{16} random variates, formed by the sum of five scaled calls to the POSIX `rand` function. The unusual high-pass filtering approach is inherited from the arrangements in [12], where we required both high-pass and low-pass outputs. In addition to what is shown in (a), we split the input into blocks of variable size—ranging from 2 to 32 in our tests—and compute mean-square value of the noise using the unbiased divide-by-$(N-1)$ formula. The harmonic powers are calculated by repeated complex downconversion; after k downconversions, the DC component will be that originally at fk, and can be found by coherent averaging. We measured up to the sixth harmonic.

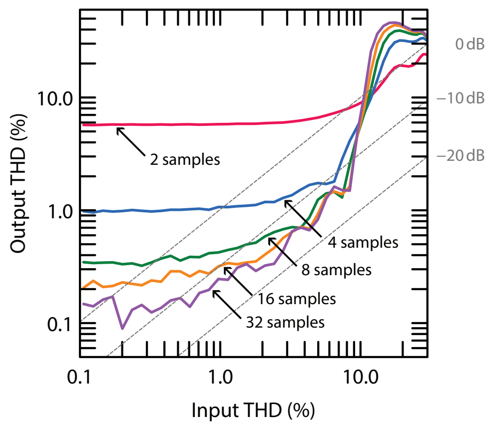

Fig. 4. Input vs. output THD for the proposed method. A $100\,\mathrm{Hz}$ sinusoid had noise added and was distorted by the function $g(x) = \tanh(1.13x)$, with the level of distortion varied by changing the amplitude of the input signal, rescaling the output to match the output range to that corresponding to an input signal of amplitude one. The ultimate floor of the output THD is dependent upon the number of samples used—the block size—for the noise estimate. Each line represents a different block size.

no longer limited to merely gross distortion, but that it can be applied to existing data of already-reasonable quality. This provides substantial expansion of the design space, distortion no longer being the hard limit that it otherwise is.

Acknowledgment

Lachlan Gunn is currently in receipt of an Endeavour Research Fellowship in order to undertake research at the University of Angers, France.

References

[1] N. Wiener, *Extrapolation, Interpolation, and Smoothing of Stationary Time Series*. Wiley, New York, 1949.

[2] ——, *Nonlinear Problems in Random Theory*. Wiley, New York, 1958.

[3] C. Coleman, *An Introduction to Radio Frequency Engineering*. Cambridge University Press, 2004.

[4] L. Ljung, *System Identification: Theory for the User*. Prentice Hall, 1987.

[5] "IEEE standard for digitizing waveform recorders," *IEEE Std 1057-2007 (Revision of IEEE 1057-1994)*, 2008.

[6] R. Martins and A. da Cruz Serra, "Automated ADC characterization using the histogram test stimulated by Gaussian noise," *IEEE Transactions on Instrumentation and Measurement*, vol. 48, no. 2, pp. 471–474, 1999.

[7] F. Alegria, P. Arpaia, A. da Cruz Serra, and P. Daponte, "ADC histogram test by triangular small-waves," in *Proc. 18th IEEE Instrumentation and Measurement Technology Conference*, vol. 3, Budapest, Hungary, 2001, pp. 1690–1695.

[8] T. Rahkonen and M. Kangas, "Histogram based background correction of ADC," in *Proc. Norchip Conference*, Oslo, Norway, 2004, pp. 95–98.

[9] M. Figueiredo, N. Paulino, G. Evans, and J. Goes, "New simple digital self-calibration technique for pipeline ADCs using the internal thermal noise," in *Proc. IEEE Int. Symp. on Circuits and Systems*, May Seattle, USA, 2008, pp. 232–235.

[10] D. Rousseau, J. R. Varela, and F. Chapeau-Blondeau, "Stochastic resonance for nonlinear sensors with saturation," *Physical Review E*, vol. 67, 2003, art. 021102.

[11] L. J. Gunn, A. Allison, and D. Abbott, "Identification of static distortion by noise measurement," *Electronics Letters*, vol. 49, no. 21, pp. 1321–1323, 2013.

[12] ——, "Real-time compensation of static distortion by measurement of differential noise gain," in *Proc. IEEE Workshop on Signal Processing Systems*, 2015, Belfast, United Kingdom, DOI:10.1109/SiPS.2014.6986079.