NOVA

University of Newcastle Research Online

nova.newcastle.edu.au

Eielsen, Arnfinn A.; Fleming, Andrew J. "Existing methods for improving the accuracy of digital-to-analog converters", Published in *Review of Scientific Instruments* Vol. 88, Issue 9, no. 094702. (2017)

Available from: <http://dx.doi.org/10.1063/1.5000974>

Reproduced from Eielsen, Arnfinn A.; Fleming, Andrew J. "Existing methods for improving the accuracy of digital-to-analog converters", *Review of Scientific Instruments* Vol. 88, Issue 9, no. 094702 (2017), with the permission of AIP Publishing.

Accessed from: <http://hdl.handle.net/1959.13/1350134>

Existing Methods for Improving the Accuracy of Digital-to-Analog Converters

Arnfinn A. Eielsen^{a)} and Andrew J. Fleming^{b)}

*School of Electrical Engineering and Computer Science, The University of Newcastle,
Callaghan, NSW 2308, Australia*

(Dated: 17 August 2017)

The performance of digital-to-analog converters is principally limited by errors in the output voltage levels. Such errors are known as element mismatch and are quantified by the integral non-linearity. Element mismatch limits the achievable accuracy and resolution in high-precision applications as it causes gain and offset error, as well as harmonic distortion. In this article, five existing methods for mitigating the effects of element mismatch are compared: physical level calibration, dynamic element matching, noise-shaping with digital calibration, large periodic high-frequency dithering, and large stochastic high-pass dithering. These methods are suitable for improving accuracy when using digital-to-analog converters that use multiple discrete output levels to reconstruct time-varying signals. The methods improve linearity and therefore reduce harmonic distortion, and can be retrofitted to existing systems with minor hardware variations. The performance of each method is compared theoretically and confirmed by simulations and experiments. Experimental results demonstrate that three of the five methods provide significant improvements in the resolution and accuracy when applied to a general-purpose digital-to-analog converter. As such, these methods can directly improve performance in a wide range of applications including nanopositioning, metrology, and optics.

Keywords: digital-to-analog conversion; linearization; metrology; nanopositioning

---

^{a)}Electronic mail: ae840@newcastle.edu.au

^{b)}Electronic mail: andrew.fleming0@newcastle.edu.au

I. INTRODUCTION

Higher accuracy and resolution from digital-to-analog converters (DACs) is called for in applications such as: adaptive optics¹, interferometry^{2,3}, scanning probe microscopy⁴⁻⁸, systems for lithography⁹, and in metrology in general^{10,11}. To meet the increasing requirements posed by precision applications, it is desirable to improve the DAC-performance in an existing system. Hence, there is a demand for methods that can be retrofitted with only minor hardware variations.

The use of a DAC introduces several non-ideal effects into a system, causing unwanted noise and disturbances. The principal source of these effects in modern high-resolution DACs is non-linearity due to element mismatch¹² — described using a static non-linear function called the integral non-linearity (INL). A static non-linearity will generate harmonic distortion if it is excited by a sinusoidal signal¹³.

The theoretical signal-to-noise-and-distortion ratio (SINAD) for a DAC with a word-width of B bits is

$$\text{SINAD} \approx 6.02B - 3.01 + 10 \log_{10}(\text{OSR}) \text{ dBc}, \quad (1)$$

when using a small noise dither to eliminate uniform quantization error¹⁴ and oversampling to reduce the noise density¹⁵. Here, OSR denotes the oversampling ratio. Results using a 16-bit DAC with an OSR of 100 are shown in Fig. 13. The theoretically achievable SINAD is 113 dBc, but the harmonic distortion caused by the INL degrades the SINAD to 93.5 dBc.

Several methods exist to mitigate INL. The INL can be reduced by having more accurate levels. The most accurate levels can be obtained using superconducting Josephson junctions^{10,11}. Using more conventional semiconductors, more accurate levels can be produced using careful component selection¹⁶ or using physical calibration¹⁷⁻²¹. Better accuracy can also be obtained using averaging techniques, such as dynamic element matching^{12,22-25}, large periodic high-frequency dithering²⁶, and large high-pass noise dithering^{27,28}. Noise-shaping (Δ - Σ modulation) with digital calibration²⁹⁻³¹ can improve accuracy due to observer-based feedback control.

These methods have the potential to improve the linearity of a conventional DAC that can switch between a multiple of fixed discrete voltage or current levels. As an alternative to producing multiple levels, time-domain averaged switching methods, such as pulse-width modulation (PWM) with pre-distortion^{32,33} or 1-bit Δ - Σ modulation³⁴, can be used for

accurate time-varying signal reproduction. The latency introduced in switched conversion tends to make such techniques ill-suited to feedback control applications. More specialized methods that focus on generating low-distortion sinusoidal signals also exist. These methods include distortion shaping^{35,36} and harmonic cancellation^{37–39}.

A. Contributions

In this article, five methods for resolution enhancement are identified which can either be retrofitted to existing hardware or applied to off-the-shelf multi-level digital-to-analog converters. A common experimental platform is developed, capable of implementing variants of all the presented methods. Where not already available, analysis methods and physical realizations are developed or improved.

II. NOISE AND DISTORTION IN DACS

There are several non-ideal effects exhibited by a DAC. The fundamental sources are aliasing and quantization, which are due to discretization in both in time and value⁴⁰. Aliasing occurs where sampling a signal in time will generate repeated spectra over the Nyquist-frequency (half the sampling rate)⁴¹. Quantization is the process of mapping a large set of values to a smaller set of values, therefore it discards some values and introduces a signal dependent error⁴².

The main secondary effects include non-linearity due to element mismatch, and thermal and semiconductor noise generated by the components in the DAC. Element mismatch causes the actual output levels of the DAC to deviate from the ideal levels. This generates both a static error as well as harmonic and intermodulation distortion⁴³. The main sources of thermal and semiconductor noise are the resistor network producing the output voltage levels, the voltage reference, and the output buffer^{43–45}. Additional non-linearity is introduced by slew-rate limitations in the output stage and glitches caused by non-ideal transistor switching^{43,46,47}. Gain mismatch and bias voltages also causes errors in the output, but these effects are linear and will not distort signals²⁵.

The methods discussed in this paper are targeting element mismatch, as it has been identified as one of the main contributors to distortion in modern DACs¹². Element mismatch

FIG. 1: Uniform quantizer model.

is modeled as a static, or memoryless, non-linearity $n(w)$. A static non-linearity will generate harmonic distortion if it is excited by a sinusoidal signal. Harmonic distortion is the presence of signal components at multiples of the frequency ω_0 in the output of the function $n(w)$. Element mismatch can be approximated by a Taylor series polynomial. The number of higher order harmonic components is related to the order of the polynomial¹³. If the non-linearity is excited by multiple sinusoidal signals with distinct frequencies, there will be intermodulation components in addition to the harmonic components. The intermodulation components appear at sums and differences of multiples of the input frequencies and can therefore appear below the frequency of the input signal with the lowest frequency. Increasing the order of the polynomial describing the non-linearity or the number of frequency components in the input will generate a higher number of harmonic and intermodulation components^{13,48}.

III. EFFECTS DUE TO QUANTIZATION

A. Uniform Quantization

A quantizer is represented by the block-diagram symbol in Fig. 1a. A quantizer is an operator that takes the input values w from a large set and maps them to discrete values y in a smaller set. A uniform quantizer maps to equidistant values with a step-size δ , called the quantization step-size or the least significant bit (LSB). The ideal uniform quantizer is a discontinuous non-linear function that will generate harmonic and intermodulation distortion^{49,50}.

A DAC typically has 2^B number of levels, where B is the word-size (bits). The quanti-

zation step-size is

$$\delta = \frac{\Delta}{2^B - 1} \quad (2)$$

where Δ is the output range of the DAC. A mid-tread uniform quantizer is defined using the truncation operator $T(w)$

$$k = T(w) \triangleq \left\lfloor \frac{w}{\delta} + \frac{1}{2} \right\rfloor \quad (3)$$

where $\lfloor \cdot \rfloor$ denotes the floor operator. The output of the quantizer k is referred to as the code, which is the input to the DAC. The output y of the quantizer given an input w is

$$y = Q(w) = \delta T(w) = \delta k. \quad (4)$$

For frequency-rich input signals and when using quantizers with $B \geq 7$ bits, the quantization error

$$q(w) \triangleq y - w = Q(w) - w. \quad (5)$$

is often modeled as an additive, zero-mean, and uniformly distributed white-noise signal with variance, power spectral density, and probability density function given by

$$\sigma_q^2 = \frac{\delta^2}{12}, \quad S(\omega) = \sigma_q^2, \quad \text{and} \quad f_q(v) = \frac{1}{\delta} \text{rect}\left(\frac{v}{\delta}\right). \quad (6)$$

This is called Bennett's classical model of quantization, or the pseudo quantization noise (PQN) model^{14,40,43,49}.

If the input signal is narrow-band or small relative to the quantization step-size, e.g. a small-amplitude sinusoidal signal, the model is no longer valid, thus introducing undesirable spurs^{14,51}. This is often the case in technical applications, where signals such as steps, sinusoids, and triangle-waves are common. The PQN model can be made valid by the addition of a dither d , as indicated in Fig. 1b.

The total output error

$$\varepsilon \triangleq y - x, \quad (7)$$

with non-subtractive dithering becomes, using (5),

$$\varepsilon = Q(x + d) - x = d + q(x + d). \quad (8)$$

The total error signal ε , consisting of the dither signal d and the dithered quantization error $q(x + d)$, can be made stationary⁵² with a constant first and second moment that is independent of the signal x , by using a non-subtractive dither d with the triangular

probability distribution function (TPDF) in the interval $[-\delta, \delta]$ ^{14,53}. In effect, the distortion due to uniform quantization is eliminated, but the variance of the total output error ε is three times larger,

$$\sigma_\varepsilon^2 = \sigma_d^2 + \sigma_q^2 = \frac{\delta^2}{4}. \quad (9)$$

For all the results in this paper, a white TPDF noise dither has been used in order to eliminate uniform quantization spuriae.

B. Non-linear Quantizer

All DACs have *element mismatch*. This means that the actual levels deviate from the ideal equidistant levels (4). The element mismatch is typically modeled as an additive static non-linearity. Static non-linearity in conventional DACs is mainly caused by the limited accuracy of resistors and current sources^{22,43,54}. The topology, such as using an R - $2R$ resistor ladder or an array of current sources, will affect the characteristics of the non-linearity^{43,54}. Hence, the output of the quantizer in (4) is modified to be

$$\tilde{y}(k) = y(k) + \delta \text{INL}(k) = \delta k + \delta \text{INL}(k). \quad (10)$$

The static non-linear function $\text{INL}(k)$ is called the integral non-linearity and the standard definition is⁴³

$$\text{INL}(k) \triangleq \frac{\tilde{y}(k) - \delta k}{\delta}. \quad (11)$$

The effect of the two static non-linearities, element mismatch and truncation, due to the input w seen on the output is described by the function $n(w)$:

$$n(w) = \delta \text{INL}(k)|_{k=T(w)} \quad (12)$$

This is a discontinuous function due to the truncation operator $T(w)$, defined in (3). The model of the non-linear quantizer is shown in Fig. 2.

IV. SYSTEM DESCRIPTION

A diagram of the experimental set-up is shown in Fig. 3. All of the experiments utilize either one or two DACs. The DAC is a 16-bit Linear Technology LTC2641 which uses a voltage switched resistor ladder to produce the output. Eight such DACs are present on

FIG. 2: Non-linear quantizer model.

FIG. 3: Experiment set-up.

the National Instruments PCIe-7851R interface card. They are connected to the on-board FPGA by a serial peripheral bus (SPI) and can provide a sampling rate of up to 1 MS/s. The output to the DACs is streamed from the computer (CPU) via direct memory access (DMA) over the peripheral component interconnect express (PCIe) bus to the FPGA. All the methods are implemented using LabView, running on the CPU.

The output of the DACs are connected to a buffer and summing stage, with adjustable gain. One of the inputs to the summing stage can be grounded if not required. There is an optional notch filter, implemented as a buffered, passive Hall-network⁵⁵, with a center frequency at $f_n = 50.0$ kHz. In order to compare the results from all methods, the notch filter was used in all experiments. The low-pass filter is an anti-aliasing filter having a second-order Butterworth response with cut-off frequency at $f_c = 25.0$ kHz and implemented using the unity-gain Sallen-Key topology⁵⁶. To achieve the best noise and linearity performance, high-voltage polypropylene capacitors and low resistor values were required. The Texas Instruments LME49990 operational amplifier was used to implement the circuits, as it has a sufficiently linear response.

The output spectra were measured using a National Instruments USB-6289, which contains an Analog Devices AD7674 18-bit successive approximation analog-to-digital converter

FIG. 4: Measuring the INL.

FIG. 5: Measured INL for two DAC channels.

(ADC). This ADC has sufficiently linear performance, with a spurious-free dynamic range (SFDR) of 120 dBFS⁵⁷ for the input signal frequencies considered here. A sampling rate of 625 kS/s was used. The USB-6289 contains two first-order passive low-pass filters with $f_c = 62.5$ kHz, which constitutes a combined second-order filter with $f_c = 40.0$ kHz, and a variable range that makes it possible to utilize the full range of the ADC depending on the input signal amplitude. All measured power spectra were generated using power spectrum estimation in LabView, using a frequency resolution of 1 Hz, at least 100 averages, and a Kaiser window⁵⁸ with window parameter $\alpha = 38$.

The voltage levels for the two DAC channels used on the National Instruments PCIe-7851R system were measured with an Agilent 34461A precision multimeter, using the set-up shown in Fig. 4. The INL for each channel is plotted in Fig. 5.

FIG. 6: Physical calibration using a secondary DAC.

The maximum input signal amplitude was used to achieve maximum performance. The largest amplitude is limited by saturation; where significant distortion is introduced. The frequencies were 99 Hz and 999 Hz. At 99 Hz the static non-linearity model is accurate. At higher frequencies, additional dynamic non-linear effects will generate distortion, and deteriorated performance is expected at 999 Hz.

V. LINEARIZATION METHODS

A. Physical Level Calibration

Element mismatch is caused by the deviation of the output levels from their ideal values. Adjusting the voltage output of the DAC produces more accurate levels and reduced non-linearity. The output levels $\tilde{y}(k)$ are the levels that can be measured on the output of the DAC, using a voltmeter. The ideal levels are the scaled input codes δk . The output is assumed to be generated according to the model (10). Gain mismatch, meaning inaccuracy in δ , and DC offset also contribute to the error. However, these effects are linear and will not distort signals²⁵. In order to reduce distortion, the objective of physical level calibration is to make (11) constant, by physical means. The calibration can take the form of adjusting component values directly using laser trimming⁵⁹, matching current sources to a common reference^{17,20,21}, or adding or subtracting correction voltages at the output via a secondary DAC^{18,60}. The latter method can be retrofitted to an existing system.

The implemented physical level calibration method is shown in Fig. 6. It works by summing the outputs of a main DAC and a secondary DAC, where the range of the secondary DAC is scaled to, at minimum, 1 LSB of the main DAC. By using a look-up table (LUT) to store correction levels applied to the secondary DAC, this DAC can be used to drive each

level error towards zero in the summing stage.

Two DAC channels can be described by (10)

$$\tilde{y}_1(k) = \delta k + \delta \text{INL}_1(k), \quad (13)$$

and

$$\tilde{y}_2(\tilde{k}) = \delta \tilde{k} + \delta \text{INL}_2(\tilde{k}). \quad (14)$$

The channel gains are α_1 and α_2 , where $\alpha_2 < \alpha_1$ and $\alpha_1\alpha_2\Delta > \delta$. The sum is

$$\begin{aligned} \tilde{y}_s(k, \tilde{k}) &= \alpha_1 \tilde{y}_1(k) + \alpha_2 \tilde{y}_2(\tilde{k}) \\ &= \alpha_1 \delta k + \alpha_1 \delta \text{INL}_1(k) + \alpha_2 \delta \tilde{k} + \alpha_2 \delta \text{INL}_2(\tilde{k}) \\ &\approx \alpha_1 \delta k + \alpha_1 \delta \text{INL}_1(k) + \alpha_2 \delta \tilde{k}, \end{aligned} \quad (15)$$

and if the secondary DAC gain α_2 is small enough: $\alpha_2 \delta \text{INL}_2(\tilde{k}) \approx 0$.

For each main DAC code k , the secondary DAC code $\tilde{k}$ can be found as

$$\arg \min_{\tilde{k}} \left| \alpha_1 \delta \text{INL}_1(k) + \alpha_2 \delta \tilde{k} \right|. \quad (16)$$

Using the secondary DAC, codes that solve this minimization problem will reduce the effect of $\text{INL}_1(k)$, assuming a bipolar DAC output. The result can be stored as a look-up table:

$$\tilde{k} = \text{LUT}(k) \quad (17)$$

The output of the summing stage is now

$$\tilde{y}_s(k) \approx \alpha_1 \delta k + e_{\text{LUT}}, \quad (18)$$

where the residual error e_{LUT} is

$$e_{\text{LUT}} = \alpha_1 \delta \text{INL}_1(k) + \alpha_2 \delta \text{LUT}(k) \quad (19)$$

and if $e_{\text{LUT}} \approx 0$ it means that

$$\text{LUT}(k) \approx -\frac{\alpha_1}{\alpha_2} \text{INL}_1(k), \quad (20)$$

that is, the codes from the LUT should ideally be equal to the scaled INL of the main DAC. Since the LUT codes are integers, and since the INL of the secondary DAC has been neglected, the compensation can not be exact.

Some issues need to be considered in practice: The actual gain for the secondary channel must be measured; as a prediction from component values alone will be inaccurate. The main DAC will likely have a constant output bias voltage, which may be larger than the highest and lowest voltage that can be produced by the secondary DAC. This bias should be neglected, as it does not impact the dynamic performance. Considering an average gain $\theta_{g_i} = \alpha_i \delta$ and bias θ_{b_i} , for a given channel i ; the uniformly spaced, biased response for the DAC channel should be described by

$$\bar{\mathbf{Y}}_i = \begin{bmatrix} \mathbf{K} & \mathbf{1} \end{bmatrix} \begin{bmatrix} \theta_{g_i} \\ \theta_{b_i} \end{bmatrix}, \quad (21)$$

where $\bar{\mathbf{Y}}_i$ is the vector of output values, $\mathbf{K}$ is the vector of possible codes, and $\mathbf{1}$ is a vector of ones. By using a vector of the measured values for a given output, $\tilde{\mathbf{Y}}_i^m$, the least-squares estimates of the gains and biases are found from:

$$\begin{bmatrix} \hat{\theta}_{g_i} \\ \hat{\theta}_{b_i} \end{bmatrix} = \begin{bmatrix} \mathbf{K} & \mathbf{1} \end{bmatrix}^\dagger \tilde{\mathbf{Y}}_i^m \quad (22)$$

Now, by subtracting the estimate of the main DAC bias $\hat{\theta}_{b_1}$ and using the estimated gain for the secondary DAC $\hat{\theta}_{g_2}$, the LUT can be generated by solving

$$\arg \min_{\tilde{k}} \left| \left[\tilde{y}_1^m(k) - \hat{\theta}_{b_1} - \delta k \right] + \hat{\theta}_{g_2} \tilde{k} \right| \quad (23)$$

for each possible main DAC code k . This optimization problem is solved offline, with measured data such as presented in Fig. 5. Alternatively, the secondary DAC codes can be found in-situ while measuring, using the successive approximation technique¹⁸.

B. Dynamic Element Matching

Dynamic element matching (DEM) relies on redundancy in the output elements^{12,22-25} and can be very effective at reducing the effects of element mismatch. If two or more DAC channels are available, element redundancy can be introduced and DEM can be retrofitted to an existing system by summing the channels. A diagram for the implementation of DEM is shown in Fig. 7.

FIG. 7: Dynamic element matching.

FIG. 8: Dynamic element matching block diagram.

The voltage output due to one element j in a DAC, e.g. the output contribution due to one bit in an R-2R ladder⁴³, can be described by

$$z_j = \begin{cases} K_j \delta / 2 + e_{hj}, & \text{if } c_j = 1 \\ -K_j \delta / 2 + e_{lj}, & \text{if } c_j = 0 \end{cases} \quad (24)$$

where $K_i = 2^{i-1}$ is the bit weight, and c_j denotes the bit. The errors e_{hi} and e_{li} denote the mismatches when the bit is turned on (high) and off (low), respectively. For a given set of bits $\{c_j\}$, the output voltage (10) of a DAC channel i is:

$$\tilde{y}_i = \sum_{j=1}^B z_j = \delta k + \text{INL}_i(k) \quad (25)$$

If two DAC channels are summed with an equal gain of $1/2$, then the output voltage is given by

$$\tilde{y}_s = \frac{1}{2} (\tilde{y}_1 + \tilde{y}_2) = \frac{1}{2} (\delta k_1 + \text{INL}_1(k_1) + \delta k_2 + \text{INL}_2(k_2)) , \quad (26)$$

and if the codes, k_1 and k_2 , output to the two DACs satisfy

$$k = k_1 + k_2 , \quad (27)$$

where k is the desired code generated by an ideal uniform quantizer, then the output of the summing stage is

$$\tilde{y}_s = \delta k + \frac{1}{2} (\text{INL}_1(k_1) + \text{INL}_2(k_2)) = \delta k + e_{\text{INL}}. \quad (28)$$

Due to null-space spanned by k_1 and k_2 , the desired code k can be expressed by a large number of combinations⁶¹. How the codes k_1 and k_2 are selected can therefore influence the properties of the error signal e_{INL} . E.g., if k is constant, the sum of the mismatches INL_i can be made time-varying by continuously randomizing the values of k_1 and k_2 that sum to k , thereby e_{INL} can be converted to a white noise signal¹².

DEM is often used when the DAC topology uses a unary coding. That is, the output is generated by a sum of elements of equal but slightly mismatched weights. Newer DEM methods have also been developed for the cases where the weights are different, such as would be the case for an R-2R ladder; which is binary coded. This is called segmented DEM. As the DACs used in the experimental set-up are binary coded, segmented DEM must be used. The segmented DEM method²⁵ is straight-forward to adapt to the case where two binary coded DACs are summed, as it maps directly to the fully segmented case.

If the element selection bits c_i are designed to satisfy

$$k = \sum_{j=1}^{2B} K_j \left(c_j - \frac{1}{2} \right) \quad (29)$$

and the output is described by

$$\tilde{y}_s = \delta k + e_{\text{INL}} = \alpha \delta k + \beta + e_{\text{DEM}}, \quad (30)$$

the stated objective of the segmented DEM method²⁵ is to make e_{DEM} noise-like. However, it does not improve the scaling error α or the bias term β introduced by the INL, but these errors do not influence the dynamic performance.

The DEM encoder is shown in Fig. 8, together with the mapping to generate the output codes to the two DACs. The mapping consists of summing the binary bit weights for a given channel. The encoder consists of a tree of switching blocks $S_{1,r}$ for $r = 1, 2, \dots, 16$, and $S_{k,1}$ for $k = 2, 3, \dots, 16$. Each block produces two output sequences depending on an input sequence and a pseudorandom bit sequence, d_k . The pseudorandom bit sequences are independent and uniformly distributed, i.e., being 0 or 1 with equal probability. A binary coding yields the weights $K_{2l-1} = K_{2l} = 2^{l-1}$ for $l = 1, 2, \dots, 16$.

FIG. 9: Noise-shaping with digital calibration.

The blocks $S_{k,1}$ are called segmenting, and $S_{1,r}$ are called non-segmenting. The top and bottom outputs of the segmenting switching blocks $S_{k,1}$ are

$$\frac{1}{2}(c_{k,1} - 1 - s_{k,1}) \quad \text{and} \quad 1 + s_{k,1} \quad (31)$$

where $c_{k,1}$ is the input and

$$s_{k,1} = \begin{cases} 0, & \text{if } c_i = \text{odd} \\ 1, & \text{if } c_i = \text{even}, d_k = 1 \\ -1, & \text{if } c_i = \text{even}, d_k = 0 \end{cases} \quad (32)$$

and for the non-segmenting blocks $S_{1,r}$

$$\frac{1}{2}(c_{1,r} - s_{1,r}) \quad \text{and} \quad \frac{1}{2}(c_{1,r} + s_{1,r}) \quad (33)$$

where $c_{1,r}$ is the input and

$$s_{k,1} = \begin{cases} 0, & \text{if } c_i = \text{even} \\ 1, & \text{if } c_i = \text{odd}, d_k = 1 \\ -1, & \text{if } c_i = \text{odd}, d_k = 0. \end{cases} \quad (34)$$

This method makes e_{DEM} a white, zero-mean noise signal, independent of the input²⁵.

C. Noise-shaping With Digital Calibration

If the pseudo quantization noise (PQN) model (6) holds, then oversampling can be used to reduce the quantization error in a desired frequency domain¹⁵. The PQN model stipulates

that the quantization error is white, i.e., it has a uniform spectral density in the bandwidth defined by the sampling frequency. This means that if the sampling frequency increases, the error spectral density decreases. The sampling theorem⁴¹ defines the minimum bandwidth required to reconstruct a signal. If a higher bandwidth is used, then this is called oversampling. When using oversampling, if the output is filtered by a filter with a smaller bandwidth than the bandwidth set by the sampling frequency, then the quantization error in the out-of-band frequency range will be removed, and the signal-to-noise ratio at the output of the filter will be improved.

Noise-shaping can be used to move more of the quantization error power into the out-of-band frequency domains, which increases the effect of oversampling⁶². Noise-shaping works by introducing an estimate of the uniform quantization error and using a feedback filter that shapes the noise power at the output of the DAC. Typically, the feedback filter is used to generate a high-pass filter for the quantization error. The noise at high frequencies is then removed by the low-pass filter used for reconstruction at the output of the DAC. A diagram for the implementation of noise-shaping is shown in Fig. 9. In addition to the uniform quantization error, the error due to element mismatch can also be compensated by adding the measured INL to the quantization error estimate^{29,30}, hence digitally calibrating the quantizer model. Generating the error estimate this way is referred to as an open-loop estimator or observer in control theory⁶³.

Noise-shaping with digital calibration can be implemented by the block diagram in Fig. 10. Given the input signal x , an estimate of DAC output (10) is generated using

$$\hat{y} = \tilde{y}^m(k) = \delta k + \delta \text{INL}^m(k). \quad (35)$$

The measured INL is found as

$$\delta \text{INL}^m(k) = \tilde{y}^m(k) - \delta k, \quad (36)$$

where $\tilde{y}^m(k)$ denotes the measurement of the output voltage. Measuring the output levels is done using the method outlined in Fig. 4. The measured INL, $\delta \text{INL}^m(k)$, is saved as a look-up table for a given code k . The generation and application of this look-up table is referred to as *digital calibration*²⁹.

Applying feedback and the noise-shaping filter $W(z)$, it is possible to filter the error $\epsilon = \hat{y} - x$, comparing the estimated non-linear quantizer output and the desired output

FIG. 10: Noise-shaping block diagram.

signal, as

$$\epsilon = \hat{y} - x = (1 - W(z)) \hat{e} = H(z) \hat{e}$$

where $\hat{e} = \hat{y} - w = \hat{q}_e + d$ is the dithered error estimate. The dither d is the TPDF noise dither discussed in Sec. III A, which makes the PQN model valid. $H(z)$ is called the sensitivity function, or noise transfer function, and if it acts as a high-pass filter, then the error $\hat{e}$ is attenuated at low frequencies. A straight-forward choice for $W(z)$ is

$$W(z) = 2z^{-1} - z^{-2}, \quad (37)$$

which yields the high-pass filter $H(z) = (1 - z^{-1})^2$ as the sensitivity function.

Since the estimated quantization and INL error $\hat{q}_e$ is based on a model, it will not perfectly match the actual quantization and INL error q_e . The error introduced is

$$e_q = \hat{q}_e - q_e. \quad (38)$$

The error due to this mismatch propagates to the output as

$$\tilde{y} = x + H(z)e - W(z)e_q \quad (39)$$

where $e = q_e + d$ is the actual dithered quantization and INL error. The actual error e will be high-pass filtered and can be attenuated by the low-pass filter. The model mismatch e_q will be filtered by the noise-shaping filter $W(z)$, which is not a high-pass filter and has unity DC-gain. Hence, it is clear that having a perfect estimate $\hat{q}_e = q_e$ will render the minimum disturbance in the reconstructed signal.

D. Large Periodic High-frequency Dithering

The harmonic distortion due to element mismatch in digital-to-analog converters can be reduced by the application of a large high-frequency periodic dither²⁶. The method is based on the fact that a static non-linear function $n(\cdot)$ can, by the application of a suitable periodic dither, be approximated by a smoothed non-linear function $N(\cdot)$ where $\|N\|_\infty \leq \|n\|_\infty$; hence reducing the effects caused by the non-linearity $n(\cdot)$ ⁶⁴⁻⁶⁷. The smoothed non-linearity $N(\cdot)$ is determined by the non-linearity $n(\cdot)$ and the amplitude distribution function of the dither. The validity of the approximation is mainly dependent on the frequency of the dither, hence it is termed high-frequency (HF) dither.

An implementation of HF dither is shown in the diagram in Fig. 11. The dither signal will cause the desired signal to sweep over several voltage levels of the DAC, effectively averaging out the mismatches. Since the HF dither is unwanted in the output signal, several methods can be used to attenuate it. Under ideal assumptions, summing two identical signals with opposite polarity should cancel the signal in the output. However, there will always be a residual signal in a practical circuit. The dither is both high in frequency, deterministic, and narrow-band, and the residual signal can be efficiently removed by filtering the output with a notch filter with center frequency at the fundamental frequency of the HF dither, as well as low-pass filtering the output.

The definition of the smoothed non-linearity $N(x)$ is:

$$N(x) \triangleq \int_{\mathbb{R}} n(x+v) dF_p(v) \quad (40)$$

Here $F_p(v)$ is the amplitude distribution function of the high-frequency (HF) dither p . This method relies on the equivalence of the above Lebesgue-Stieltjes integral and the time-average over one period τ of the periodic dither, where x is assumed to be constant for the

FIG. 11: Large periodic high-frequency dithering.

duration of the period⁶⁷:

$$\int_{\mathbb{R}} n(x+v) dF_p(v) = \frac{1}{\tau} \int_0^\tau n(x+p(t)) dt \quad (41)$$

The error introduced by the assumption of x being piecewise constant with duration τ goes to zero as $\tau \rightarrow 0$. If the distribution F_p is absolutely continuous, the averaging effect of the dither p on the non-linearity $n(w)$ can be found by evaluation of the Lebesgue-Stieltjes integral of the form:

$$\int_{\mathbb{R}} n(x+v) dF_p(v) = \int_{\mathbb{R}} n(x+v) f_p(v) dv \quad (42)$$

Here, $f_p(v)$ is the amplitude density function, defined as:

$$f_p(v) \triangleq \frac{d}{dv} F_p(v) \quad (43)$$

The smoothed non-linearity $N(x)$ can be found as the cross-correlation of $n(w)$ and the amplitude density function of the dither. A signal with uniform amplitude density

$$f_p(v) = \frac{1}{2A} \text{rect}\left(\frac{v}{2A}\right) = \begin{cases} \frac{1}{2A} & |v| \leq A \\ 0 & |v| > A \end{cases} \quad (44)$$

is an example of a signal with an absolutely continuous amplitude distribution function. One realization is the triangle-wave, which is the dither used in the experiments. If $f_p(v)$ is even, (42) can be found as a convolution product and by using the triangle-wave dither this is equivalent to filtering $\text{INL}(k)$ by a filter with Fourier transform

$$\int_{-\infty}^{\infty} \delta f_p(v) e^{-j2\pi\xi v} dv = \delta \text{sinc}(2A\xi), \quad (45)$$

which has a low-pass characteristic. Hence, the dither attenuates the variations in the INL. Increasing the dither amplitude A increases the smoothing of the INL, but reduces the usable range.

E. Large Stochastic High-pass Dithering

Compared to the application of a large high-frequency periodic dither, the application of a large stochastic noise dither should have a similar effect. Instead of the averaging action of a sweep across several voltage values as in the case of high-frequency periodic dither, the averaging is caused by random selection of voltage levels. However, applying

FIG. 12: Large high-pass noise dithering.

a large dither signal can only be considered useful if the dither can be removed from the output: The application of a noise dither with a high-pass power spectral density has been shown to reduce the distortion due to element mismatch in DACs [Blessner:1987wn](#), [Leyonhjelm:1998ga](#). The low-pass reconstruction filter can then be used to remove the unwanted high-frequency dither noise in the output. A diagram of the implementation is shown in Fig. 12.

Compared to the large high-frequency periodic dither in Sec. VD, the smoothing action of the dither is determined by assuming the input signal x and dither signal h are random variables that are statistically independent^{68–71}, or that the input signal x is a constant⁷². When considering two statistically independent random variables, the time average is^{65,73}

$$\frac{1}{t_2 - t_1} \int_{t_1}^{t_2} n(x(t) + h(t)) dt = \int_{\mathbb{R}} n(x + v) f_h(v) dv, \quad (46)$$

given the two integrals exist. Hence, the effective non-linearity can, similar to (40), be defined as

$$N_h(x) \triangleq \int_{\mathbb{R}} n(x + v) f_h(v) dv, \quad (47)$$

where $f_h(v)$ is the absolutely continuous probability density function for the dither signal h .

A straight-forward method to generate a zero-mean, high-pass power spectral density for the dither h , it is to high-pass filter a white noise signal. No matter what probability density function the signal has, due to the central limit theorem, the linear filtering operation tends to generate an approximate Gaussian probability distribution⁷⁴:

$$f_h(v) \approx \frac{1}{\sqrt{2\pi}\sigma_h} e^{-\frac{v^2}{2\sigma_h^2}} \quad (48)$$

Hence, similar to the expression in (45), the INL will be smoothed by the low-pass characteristic of the Gaussian distribution:

$$\int_{-\infty}^{\infty} \delta f_h(v) e^{-j2\pi\xi v} dv = \delta e^{-\frac{\sigma_h^2 \xi^2}{2}} \quad (49)$$

FIG. 13: Power spectra for a high-performance 16-bit DAC showing harmonic distortion due to element mismatch.

Increasing the variance σ_h^2 of the dither signal h , causes increased smoothing of the INL, but reduces the usable range.

VI. RESULTS AND DISCUSSION

A set of experiments were conducted in order to assess the performance improvement that can be achieved. The measured performance results are summarized in Tab. I. The type of compensation method is recorded in the ‘method’ column, and the ‘enabled’ column indicates if the method was turned on or off. The ‘scale’ column states the maximum input signal amplitude as a percentage of the maximum amplitude for the DACs, and the ‘freq.’ column records the frequency. The ‘rate’ column records the sampling rate for a given experiment. The scale of the input signal indicates the usable range for the DACs when using a given method; for the dithering methods the usable range is reduced by the amplitude of the dither.

The performance results are summarized using the following standard metrics⁴³: The signal-to-noise ratio,

$$\text{SNR} = 20 \log_{10} (\sigma_s / \sigma_n) ,$$

TABLE I: Measured performance results.

Method	Enabl.	Scale	Freq.	Rate	SINAD	THD	SNR
(a) BASELINE		100%	99 Hz	1 MS/s	93.90 dBc	-93.98 dBc	110.5 dB
(b) Physical calibration (PHYSCAL)	NO	100%	99 Hz	1 MS/s	93.87 dBc	-93.96 dBc	110.4 dB
(c) Physical calibration	YES	100%	99 Hz	1 MS/s	104.6 dBc	-105.7 dBc	110.5 dB
(d) Noise-shaping with digital calibration (NSDCAL)	NO	99%	99 Hz	1 MS/s	93.89 dBc	-93.96 dBc	111.3 dB
(e) Noise-shaping with digital calibration	YES	99%	99 Hz	1 MS/s	105.8 dBc	-107.1 dBc	111.5 dB
(f) Periodic high-frequency dither (PHFD)	NO	50%	99 Hz	1 MS/s	90.44 dBc	-90.48 dBc	109.0 dB
(g) Periodic high-frequency dither	YES	50%	99 Hz	1 MS/s	104.8 dBc	-106.6 dBc	109.1 dB
(h) BASELINE		100%	999 Hz	1 MS/s	93.71 dBc	-93.80 dBc	110.4 dB
(i) Physical calibration	NO	100%	999 Hz	1 MS/s	93.62 dBc	-93.71 dBc	110.3 dB
(j) Physical calibration	YES	100%	999 Hz	1 MS/s	96.99 dBc	-97.19 dBc	110.3 dB
(k) Noise-shaping with digital calibration	NO	99%	999 Hz	1 MS/s	93.73 dBc	-93.80 dBc	111.4 dB
(l) Noise-shaping with digital calibration	YES	99%	999 Hz	1 MS/s	96.99 dBc	-97.14 dBc	111.4 dB
(m) Periodic high-frequency dither	NO	50%	999 Hz	1 MS/s	90.12 dBc	-90.17 dBc	109.0 dB
(n) Periodic high-frequency dither	YES	50%	999 Hz	1 MS/s	101.0 dBc	-101.8 dBc	108.7 dB
(o) Dynamic element matching (DEM)	NO	100%	99 Hz	0.5 MS/s	93.81 dBc	-93.95 dBc	108.1 dB
(p) Dynamic element matching	YES	100%	99 Hz	0.5 MS/s	37.41 dBc	-38.23 dBc	44.79 dB
(q) Dynamic element matching	NO	100%	99 Hz	50 kS/s	93.48 dBc	-93.95 dBc	103.0 dB
(r) Dynamic element matching	YES	100%	99 Hz	50 kS/s	51.63 dBc	-57.37 dBc	52.81 dB
(s) Stochastic high-pass dither (SHPD)	NO	70%	99 Hz	400 kS/s	92.78 dBc	-92.88 dBc	108.7 dB
(t) Stochastic high-pass dither	YES	70%	99 Hz	400 kS/s	92.26 dBc	-97.03 dBc	93.85 dB

where σ_s is the standard deviation of the input signal, and σ_n is the the standard deviation of the noise excluding the harmonic components of the input signal. The total-harmonic-distortion,

$$\text{THD} = 20 \log_{10} (\sigma_d / \sigma_s) ,$$

where σ_d is the standard deviation of the harmonic components of the input signal, excluding any other component. When computing the THD, all the harmonic components up to 10 kHz were used. The signal-to-noise-and-distortion ratio,

$$\text{SINAD} = 20 \log_{10} (\sigma_s / \sigma_t) ,$$

where the standard deviation $\sigma_t = \sigma_d + \sigma_n$ accounts for all unwanted components in the output signal. These standard deviations are usually found using a power spectrum^{43,54}, which can computed with the Welch method⁷⁵ using a narrow windowing function, such as the Kaiser window⁵⁸ with a high α coefficient; since the signal components will have a high dynamic range. The MATLAB functions `snr`, `thd`, and `sinad` can be used to produce these numbers, provided a measured time-series or power spectrum. Similar functions are

FIG. 14: Experimental SINAD gains.

available in LabView. The measured power spectra presented here were generated using power spectrum estimation in LabView, which uses the Welch method, using a frequency resolution of 1 Hz, at least 100 averages, and a Kaiser window with window parameter $\alpha = 38$.

The gains in SINAD are summarized in Fig. 14, and the reductions in THD are summarized in Fig. 15. The gains and reductions are computed relative to the experimental configuration, measuring the difference between output with the method is turned off and on, and relative to the baseline, measuring the difference between the best-case unmodified DAC output and the output with the method turned on.

A. Baseline

The performance of a single DAC channel on the National Instruments PCIe-7851R interface used in the experiments is recorded by the power spectra in Fig. 13, and in entry (a) and (h) in Tab. I. The linearization methods should ideally provide an improved performance compared to these results.

FIG. 15: Experimental THD reductions.

B. Physical Level Calibration

The results for physical level calibration (PHYSICAL) are shown in Fig. 16, and entry (b), (c), (i), and (j) in Tab. I. Fig. 16a, shows the inputs to the primary and secondary DAC channels. The primary channel receives the codes for the input signal, while the secondary channel receives the correction codes from the look-up table. Fig. 16b shows the simulation power spectrum when using a full-scale input at 99 Hz. This simulation was performed introducing a random scaling error of up to 10% in the correction codes. Fig. 16c shows the corresponding measured experimental result.

As seen from the results, physical level calibration yields a significant improvement in the THD and the SINAD for a 99 Hz input. The performance improvement is reduced for the 999 Hz input. This is due to additional non-linear behavior that is not related to static element mismatch. As the input frequency increases, the outputs must be switched at a higher rate, which is likely to produce more power in the output due to slew-rate limitations and switching glitches.

At minimum the method requires two DACs, two gain stages, and a summing stage. With regard to computational complexity, the method is lightweight; only requiring the

FIG. 16: Physical level calibration (PHYSCAL).

implementation of a look-up table of integer values. The main disadvantage of the method is the requirement to measure the levels of both channels. This requires a precision voltmeter, and is time-consuming: For a 16-bit converter, if measuring a level takes 0.1 seconds, it takes nearly two hours to measure all levels for one channel. The calibration procedure should be repeated periodically, as the levels are likely to drift over time due to temperature changes and component aging. Any errors in the measured levels directly translate to deteriorated performance.

C. Dynamic Element Matching

The results for dynamic element matching (DEM) are shown in Fig. 17, and entry (o), (p), (q), and (r) in Tab. I. Fig. 17a shows the inputs to the two DACs and the the sum of the two signals. The codes are randomized, but they sum to the desired output. Fig. 17b shows the simulation power spectrum when using a full-scale input at 99 Hz and Fig. 17c shows the corresponding measured experimental result.

(a) Input signals to the DACs.

(b) Simulation power spectra.

(c) Measured power spectra.

FIG. 17: Dynamic element matching (DEM).

Dynamic element matching (DEM) experimental results were poor. The experimental results do not match the simulation results. It should be noted that the performance improves significantly when the sampling rate is reduced from 0.5 MS/s to 50 kS/s, which also reduces the switching rate between the levels. This suggests that additional dynamic non-linearities, slew-rate limitations and switching glitches, are excited by rapid switching and that these effects are more pronounced at higher switching rates. It is apparent that the DAC topology used in the experiments is not ideal for the large and rapid switching required to obtain good performance. The main advantage of DEM is that no knowledge of the INL is necessary in order to compensate for it. In simulations, the effect of DEM can be significantly improved by using noise-shaping. However, experimental results using this combination of methods were dismissed as the increased switching rate due to noise-shaping further deteriorated the performance.

In the implementation, only two DACs and a summing stage are required. However, the method is significantly more computationally intensive compared to the other methods

(a) Actuation signal due to noise-shaping.

(b) Simulation (ideal cases).

(c) Simulation power spectra.

(d) Measured power spectra.

FIG. 18: Noise-shaping with digital calibration (NSDCAL).

presented. If noise-shaping is added, the computational complexity is increased further.

D. Noise-shaping With Digital Calibration

The results for noise-shaping with digital calibration (NSDCAL) are shown in Fig. 18, and entry (d), (e), (k), and (l) in Tab. I. Fig. 18a plots the difference between the input codes with and without compensation, i.e., the actuation signal generated by noise-shaping feedback filter to reduce distortion.

Fig. 18b shows the effect of noise-shaping on an ideal DAC. In the first case, it can be seen that quantization error is shaped according to the sensitivity function (a high-pass filter). The low-frequency performance displayed is, however, not physically attainable. The second case shows the response if realistic measurement noise is added. The effect of low-pass filtering the output is also shown. It is the low-pass filtering of the spectrally shaped error that yields the performance increase. Fig. 18c shows the simulation power spectrum using a 99%-amplitude input signal at 99 Hz. A reduction in input amplitude is

required to accommodate the addition of the actuation signal. The simulation was performed introducing up to 50% randomized error in the table of measured output levels. Fig. 18d shows the corresponding measured experimental result.

As in the case of physical level calibration, noise-shaping with digital calibration provides a significant performance improvement in the THD and the SINAD for a 99 Hz input. Similarly, the performance improvement is reduced for the 999 Hz input. Hence, additional non-linear effects also affects this method as the input frequency is increased.

The method requires only one DAC, but relies on sufficient low-pass filtering of the output. The computational complexity of the method is modest. A look-up table of floating point numbers is required, as well as the implementation of an second-order infinite impulse response (IIR) filter. As in the case of physical level calibration, the main drawback is that the INL must be measured.

Noise-shaping is commonly combined with DEM. As noted above, even though DEM showed improved performance in simulations when combined with noise-shaping, results combining the two methods were not included, as the increased switching rate due to noise-shaping deteriorated the performance when using DEM.

E. Large Periodic High-frequency Dither

The results for large periodic high-frequency dither (PHFD) are shown in Fig. 19, and entry (f), (g), (m), and (n) in Tab. I. Fig. 19a shows both a dither signal, and the DAC input with a dither signal added. The primary channel input is the sum of the input signal and a high-frequency triangle-wave dither. The secondary channel input is the same, but the dither has opposite sign. Fig. 19b shows the result of evaluating (42) for the measured INL and the effect of the dither on the INL; the range is reduced, but the effective INL is also reduced. Fig. 19c shows the simulation power spectrum when using a 50% amplitude input signal (and 50% amplitude dither signal) at 99 Hz. Fig. 19d shows the corresponding measured experimental result.

Applying a large periodic high-frequency dither in the form of a triangle-wave can be seen to yield a significant performance improvement in the THD and the SINAD for both the 99 Hz and 999 Hz input. This method appears to be more robust to non-linearities that become more pronounced as the input frequency increases. This is because the switching

(a) Input signals to the DACs.

(b) Effective INL after dithering.

(c) Simulation power spectra.

(d) Measured power spectra.

FIG. 19: Large periodic high-frequency dither (PHFD).

due the dither signal stays constant and any increase in switching rate is due to the input signal alone. The step between successive switching levels is also constrained, compared to the random levels generated by DEM and stochastic high-pass noise dither. Large switching steps are likely to exacerbate dynamic non-linear effects, especially slew-rate limitations. Also, when using the physical level calibration and noise-shaping methods, the switching introduced due to compensation increases as the input frequency increases. This means that the compensation signals are dependent on the input signal, and more rapid switching will be seen for higher frequency input signals.

The main advantage of this method is the simplicity. It only requires the generation of a triangle-wave signal and no knowledge of the INL is necessary. This makes it the simplest method with regard to computational complexity. The method can be implemented using a single DAC and an effective notch and low-pass filter. For improved suppression of the dither, a second DAC and a summing stage should be used. The main disadvantage of this method is the reduction of the usable output range of the DAC. This causes the noise-floor

(a) Input signals to the DACs.

(b) Simulation (ideal cases).

(c) Simulation power spectra.

(d) Measured power spectra.

FIG. 20: Large stochastic high-pass dither (SHPD).

to be higher. However, an improvement in the noise-floor can be obtained by averaging several channels, if several DACs are available.

F. Large Stochastic High-pass Dither

The results for large stochastic high-pass dither (SHPD) are shown in Fig. 20, and entry (s) and (t) in Tab. I. Fig. 20a shows the input to one of the two DACs used; the input signal and the added stochastic high-pass filtered noise dither. The inputs are identical except for the dither signals which are uncorrelated. The sum of two channels was used to obtain additional attenuation of the noise dither signals due to averaging. Fig. 20b shows the effect of the dither on an ideal DAC. The dither noise is present above the 10 kHz baseband and is severely attenuated by the low-pass filter on the output. Fig. 20c shows the simulation power spectrum when using a 70% amplitude input signal at 99 Hz. Fig. 20d shows the corresponding measured experimental result.

Application of a large stochastic high-pass dither only provided modest improvements

in the THD, and a significant worsening of the noise-floor, hence a poor SINAD. As the high-pass dither signal causes rapid switching between the DACs output elements, some of this noise increase arises from slew-rate limitations and switching glitches. Some of the noise that appears in the baseband is likely due to the whitening of the spectrum that occurs when non-linearities are excited by spectrally shaped Gaussian noise⁷⁶. The sample rate was reduced to 400 kS/s for the experiments. A lower sample rate generates a lower noise-floor, but it also reduces the bandwidth and effectiveness of the high-pass dither. The sample rate of 400 kS/s was chosen as there was a noticeable effect on the INL, and a limited increase in the noise-floor. As was the case for DEM, a DAC topology with better switching performance is needed for this method to work well in practice.

This method requires only one DAC, but sufficient low-pass filtering on the output is needed. The computational complexity is higher than the large periodic high-frequency dither method, as a high-order infinite impulse response (IIR) filter and random number generator are required. An improvement in the noise-floor can be obtained by averaging several channels if several DACs are available. In the experiments two DACs were used, in order to improve the noise-floor.

VII. CONCLUSIONS

Five different methods are investigated for mitigating the effects of non-linearity due to element mismatch in digital-to-analog converters. All of the methods were implemented using a standard off-the-shelf digital-to-analog converter card with custom digital logic and analog circuitry. Experimental results demonstrate that three of the five methods provided significant performance improvements: physical level calibration, noise-shaping with digital calibration, and periodic high-frequency dithering. The two methods that did not perform well, dynamic element matching and stochastic high-pass noise dithering, rely on rapid switching which excites additional dynamic non-linearities. Noise-shaping with digital calibration requires the least custom hardware but relies on precise measurements of the non-linearity. Periodic high-frequency dithering does not require any knowledge of the non-linearity and is the least sensitive to operating conditions. However, additional filtering is required to suppress the dither signal, and the effective range is reduced which increases the noise-floor. Since multiple channels can be averaged to obtain an arbitrarily low noise-floor,

the periodic high-frequency dither method is recommended as the most practical method for significantly improving the accuracy and resolution of digital-to-analog converters.

REFERENCES

- ¹J. M. Beckers, *Astron. Astrophys.* **31**, 13 (1993).
- ²D. Bechtle, *Rev. Sci. Instrum.* **47**, 1377 (1976).
- ³B. P. Abbott, R. Abbott, R. Adhikari, P. Ajith, *et al.*, *Rep. Prog. Phys.* **72**, 076901 (2009).
- ⁴D. R. Baselt, S. M. Clark, M. G. Youngquist, C. F. Spence, and J. D. Baldeschwieler, *Rev. Sci. Instrum.* **64**, 1874 (1993).
- ⁵R. C. Munoz, P. Villagra, G. Kremer, L. Moraga, and G. Vidal, *Rev. Sci. Instrum.* **69**, 3259 (1998).
- ⁶S. M. Salapaka and M. V. Salapaka, *IEEE Control Syst. Mag.* **28**, 65 (2008).
- ⁷G. M. Clayton, S. Tien, K. K. Leang, Q. Zou, and S. Devasia, *J. Dyn. Syst.-T. ASME* **131**, 061101 (19 pages) (2009).
- ⁸Y. K. Yong, S. O. R. Moheimani, B. J. Kenton, and K. K. Leang, *Rev. Sci. Instrum.* **83**, 121101 (22 pages) (2012).
- ⁹A. Biswas, I. S. Bayer, A. S. Biris, T. Wang, E. Dervishi, and F. Faupel, *Adv. Colloid Interface Sci.* **170**, 2 (2012).
- ¹⁰C. A. Hamilton and Y. H. Tang, *Metrologia* **36**, 53 (1999).
- ¹¹M. Nakanishi, *Rev. Sci. Instrum.* **83**, 114701 (2012).
- ¹²I. Galton, *IEEE Trans. Circuits Syst. II* **57**, 69 (2010).
- ¹³J. J. Bussgang, L. Ehrman, and J. W. Graham, *Proc. IEEE* **62**, 1088 (1974).
- ¹⁴R. A. Wannamaker, S. Lipshitz, J. Vanderkooy, and J. N. Wright, *IEEE Trans. Signal Process.* **48**, 499 (2000).
- ¹⁵R. G. Lyons, *Understanding Digital Signal Processing*, 3rd ed. (Prentice Hall, 2010).
- ¹⁶D. R. Zander, *Rev. Sci. Instrum.* **42**, 797 (1971).
- ¹⁷D. W. J. Groeneveld, H. J. Schouwenaars, H. A. H. Termeer, and C. A. A. Bastiaansen, *IEEE J. Solid-State Circuits* **24**, 1517 (1989).
- ¹⁸J. Goes, J. Franca, N. Paulino, J. Grilo, and G. Temes, in *Proc. IEEE Int. Symp. on Circuits and Syst.* (IEEE, London, 1994) pp. 345–348.
- ¹⁹U.-K. Moon, G. C. Temes, and J. Steensgaard, *IEEE Commun. Mag.* **37**, 136 (1999).

- ²⁰A. R. Bugeja and B.-S. Song, IEEE J. Solid-State Circuits **35**, 1841 (2000).
- ²¹G. I. Radulov, P. J. Quinn, H. Hegt, and A. van Roermund, in *Proc. 40th European Solid State Circuits Conf.* (IEEE, Venice, 2005) pp. 169–172.
- ²²R. J. Van De Plassche, IEEE J. Solid-State Circuits **11**, 795 (1976).
- ²³I. Galton and P. Carbone, IEEE Trans. Circuits Syst. II, Analog Digit. Signal Process. **42**, 763 (1995).
- ²⁴R. T. Baird and T. S. Fiez, IEEE Trans. Circuits Syst. II, Analog Digit. Signal Process. **42**, 753 (1995).
- ²⁵K. L. Chan, N. Rakuljic, and I. Galton, IEEE Trans. Circuits Syst. I **55**, 3383 (2008).
- ²⁶A. A. Eielsen and A. J. Fleming, IEEE Trans. Circuits Syst. I **PP**, 1 (2016).
- ²⁷B. A. Blesser and B. N. Locanthi, J. Audio Eng. Soc. **35**, 448 (1987).
- ²⁸S. A. Leyonhjelm, M. Faulkner, and P. Nilsson, Wireless Pers Commun **8**, 31 (1998).
- ²⁹T. Cataltepe, A. R. Kramer, L. E. Larson, G. C. Temes, and R. H. Walden, in *Proc. IEEE Int. Symp. on Circuits and Syst.* (Portland, 1989) pp. 647–650.
- ³⁰J. Arias, P. Kiss, V. Boccuzzi, L. Quintanilla, L. Enriquez, J. Vicente, D. Bisbal, J. S. Pablo, and J. Barbolla, IEEE Trans. Circuits Syst. I **52**, 1033 (2005).
- ³¹M. Frey and H.-A. Loeliger, IEEE Trans. Circuits Syst. I **54**, 229 (2007).
- ³²J. M. Goldberg, *Signal processing for high resolution pulse width modulation based digital-to-analogue conversion*, Ph.D. thesis, King’s College London, King’s College London (1992).
- ³³F. Chierchie and S. O. Aase, IEEE Trans. Circuits Syst. I **62**, 2606 (2015).
- ³⁴R. Schreier and G. C. Temes, *Understanding Delta-Sigma Data Converters* (Wiley-Interscience, 2005).
- ³⁵K. Wakabayashi, K. Kato, T. Yamada, O. Kobayashi, H. Kobayashi, F. Abe, and K. Niitsu, J Electron Test **28**, 641 (2012).
- ³⁶F. Abe, Y. Kobayashi, K. Sawada, K. Kato, O. Kobayashi, and H. Kobayashi, in *Proc. IEEE International Test Conference* (IEEE, 2014) pp. 1–10.
- ³⁷B. K. Vasan, S. K. Sudani, D. J. Chen, and R. L. Geiger, IEEE Trans. Circuits Syst. I **60**, 1122 (2013).
- ³⁸M. J. Barragan, G. Leger, D. Vazquez, and A. Rueda, Analog Integr Circ Sig Process **82**, 67 (2014).
- ³⁹C. Shi and E. Sanchez-Sinencio, IEEE Trans. Circuits Syst. I **62**, 2227 (2015).

- ⁴⁰B. Widrow, I. Kollar, and M.-C. Liu, IEEE Trans. Instrum. Meas. **45**, 353 (1996).
- ⁴¹M. Unser, Proc. IEEE **88**, 569 (2000).
- ⁴²S. P. Lipshitz, R. A. Wannamaker, and J. Vanderkooy, J. Audio Eng. Soc. **40**, 355 (1992).
- ⁴³M. J. Pelgrom, *Analog-to-Digital Conversion*, 2nd ed. (Springer-Verlag, 2013).
- ⁴⁴J. A. Connelly and K. P. Taylor, IEEE Trans. Circuits Syst. **38**, 1133 (1991).
- ⁴⁵P. Horowitz and W. Hill, *The Art of Electronics*, 3rd ed. (Cambridge University Press, 2015).
- ⁴⁶D. M. Freeman, J. Audio Eng. Soc. **25**, 178 (1977).
- ⁴⁷L. Risbo, R. Hezar, B. Kelleci, and H. Kiper, IEEE J. Solid-State Circuits **46**, 2892 (2011).
- ⁴⁸X. Wu, Z. Q. Lang, and S. A. Billings, IEEE Trans. Signal Process. **55**, 3239 (2007).
- ⁴⁹W. R. Bennett, Bell System Technical Journal **27**, 446 (1948).
- ⁵⁰N. M. Blachman, IEEE Trans. Acoust., Speech, Signal Process. **33**, 1417 (1985).
- ⁵¹J. Vanderkooy and S. P. Lipshitz, J. Audio Eng. Soc. **32**, 106 (1984).
- ⁵²L. Ljung, *System Identification: Theory for the User*, 2nd ed. (Prentice Hall, 1999).
- ⁵³B. Widrow and I. Kollar, *Quantization Noise* (Cambridge University Press, 2008).
- ⁵⁴W. Kester, ed., *Analog-Digital Conversion* (Analog Devices, Inc., 2004).
- ⁵⁵H. Hall, Circuit Theory, IRE Transactions on **2**, 283 (1955).
- ⁵⁶R. P. Sallen and E. L. Key, Circuit Theory, IRE Transactions on **2**, 74 (1955).
- ⁵⁷Decibels relative to full scale (dBFS).
- ⁵⁸J. F. Kaiser and R. W. Schafer, IEEE Trans. Acoust., Speech, Signal Process. **28**, 105 (1980).
- ⁵⁹B. J. Tesch and J. C. García, in *BIPOL-96* (IEEE, 1996) pp. 204–207.
- ⁶⁰J. V. Kane, Rev. Sci. Instrum. **36**, 1062 (1965).
- ⁶¹L. Hernandez, IEEE Trans. Circuits Syst. I **45**, 1068 (1998).
- ⁶²H. I. Spang and P. Schultheiss, Communications Systems, IRE Transactions on **10**, 373 (1962).
- ⁶³C.-T. Chen, *Linear System Theory and Design*, 3rd ed. (Oxford University Press, 1999).
- ⁶⁴R. Oldenburger and R. C. Boyer, J Basic Eng-T ASME **84**, 559 (1962).
- ⁶⁵G. Zames and N. A. Shneydor, IEEE Trans. Autom. Control **21**, 660 (1976).
- ⁶⁶S. Mossaheb, Int J Control **38**, 557 (1983).
- ⁶⁷L. Iannelli, K. H. Johansson, U. T. Jönsson, and F. Vasca, Automatica **42**, 669 (2006).
- ⁶⁸I. De Lotto and G. E. Paglia, IEEE Trans. Instrum. Meas. **IM-35**, 170 (1986).

- ⁶⁹M. Bartz, *Microwaves & RF* **32**, 192 (1993).
- ⁷⁰F. Adamo, F. Attivissimo, N. Giaquinto, and A. Trotta, *IEEE Trans. Instrum. Meas.* **52**, 1200 (2003).
- ⁷¹J. Sanjuán, A. Lobo, and J. Ramos-Castro, *Rev. Sci. Instrum.* **80**, 114901 (2009).
- ⁷²P. Carbone, C. Narduzzi, and D. Petri, *IEEE Trans. Instrum. Meas.* **43**, 139 (1994).
- ⁷³K. L. Chung, *A Course in Probability Theory*, 3rd ed. (Academic Press, 2001).
- ⁷⁴R. G. Brown and P. Y. C. Hwang, *Introduction to Random Signals and Applied Kalman Filtering* (Wiley-Interscience, 1997).
- ⁷⁵P. D. Welch, *IEEE Trans. Audio Electroacoust.* **15**, 70 (1967).
- ⁷⁶G. L. Wise, A. P. Traganitis, and J. B. Thomas, *IEEE Trans. Inf. Theory* **23**, 84 (1977).