

Linearisation of Digital-to-Analog Converters by Model Predictive Control[★]

Bikash Adhikari^{*} Raymond van der Rots^{*} John Leth^{**}
Arnfinn A Eielsen^{*}

^{*} *University of Stavanger, Norway (e-mail: {bikash.adhikari, raymond.vanderrots, arnfinn.a.eielsen}@uis.no).*

^{**} *Aalborg University, Denmark (e-mail: jjl@es.aau.dk)*

Abstract: Digital-to-analog converters (DACs) are used to reconstruct digital quantised signals in the analog domain. Practical DACs exhibit several non-ideal effects e.g., principally integral non-linearity (INL). INL is caused by element mismatch, meaning actual output levels deviate from ideal counterparts causing signal distortion. To reduce this error a method employing moving horizon optimal control is proposed where INL has been integrated into the model. The model is built by precisely measuring the INL of the physical device and organising the data into a lookup table. Previous attempts at using moving horizon optimal quantiser (MHOQ) have assumed ideal quantisation, and are therefore unable to reduce the impact of INL. Several other methods exist that can linearise and mitigate the impact of INL but have significant drawbacks such as lack of guaranteed stability, complexity, and the need to use specialised and custom circuit topologies that cannot be replicated with off-the-shelf equipment. The effectiveness of the proposed method is demonstrated via simulations as well as experiments on a common off-the-shelf DAC.

Keywords: Digital-to-analogue converters, predictive control, quantisation, signal processing, moving horizon, constrained set

1. INTRODUCTION

In a typical digital signal processing system an analog-to-digital converter (ADC) performs sampling and quantisation to obtain a digital representation of a signal and a digital-to-analog converter (DAC) is used to recreate the sampled and quantised signal in a physical form, typically a current or voltage. DACs are key to bridging the analog and digital domains. Higher accuracy and resolution from DACs are called for in applications such as digital audio and video recording (Goedhart et al., 1982), adaptive optics (Beckers, 1993), interferometry (Abbott et al., 2009), scanning probe microscopy (Fantner et al., 2005), systems for lithography (Biswas et al., 2012), and metrology (Hamilton and Tang, 1999). To meet the increasing requirements posed by precision applications, it is desirable to improve DAC performance.

The use of a DAC introduces several non-ideal effects into a system, causing unwanted noise and disturbance. The principal source of these effects in modern high-resolution DACs is non-linearity due to element mismatch (Galton, 2010) — 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 signal (Busgang et al., 1974).

Several methods exist to mitigate INL; component selection (Zander, 1971), physical calibration (Groeneveld et al., 1989), special layout techniques, laser trimming,

and distributed biasing (Andersson et al., 2003). However, these methods can not address matching errors that remain after fabrication. To address matching errors, dynamic methods that manipulate the input signal or circuit elements using digital circuitry are employed; dynamic element matching (DEM) (van de Plassche, 1976; Galton, 2010), periodic high-frequency dithering (Eielsen and Fleming, 2017), large high-pass noise dithering (Leyonhjelm et al., 1998) and noise-shaping with digital calibration (Frey and Loeliger, 2007). For a comprehensive overview of existing techniques aimed at improving DAC accuracy, refer to Eielsen and Fleming (2017) and the references provided therein.

For signals at lower frequencies the static non-linearity model is accurate. At higher frequencies, additional dynamic non-linear effects will generate distortion, and deteriorate performance. From the experiment it can be determined that the typical transition to lower frequencies lies within about one-hundredth to ten-thousandth the maximum physical sampling rate of a DAC (Eielsen and Fleming, 2016, 2017). The most pronounced dynamic non-linear effects are glitches (Eielsen et al., 2020; Faza et al., 2023) and slewing (Pelgrom, 2017; Kester, 2005). Various noise sources will also contribute to the error, such as thermal noise and digital cross-talk (Pelgrom, 2017; Kester, 2005).

In this article, a method for linearising a DAC with INL is developed, given that INL is the most significant source of error at lower frequencies. The method is intended to be applied to any DAC where the INL can be measured

[★] This work was supported by the Research Council of Norway, project FRIPRO 313716.

with some accuracy. In this respect, the method is most similar to noise shaping with digital calibration (Frey and Loeliger, 2007). However, a significant drawback of noise-shaping is how to determine the stability of the feedback structure (Neitola, 2017). DEM requires redundancy in output levels, and has severe restrictions in terms of applicable DAC circuit topologies, which makes it difficult to apply in general; an integrated circuit essentially has to be specifically designed to apply the method. Periodic dithering techniques have been shown to be very effective (Eielsen and Fleming, 2017) and can currently provide very high performance, but will impact the signal-to-noise ratio (SNR) as the allowable output range is reduced. Hence it may be difficult to further improve performance using this technique.

In order to address some of these issues a moving horizon optimisation scheme can be utilised as it is possible to determine the stability properties of the method. It can in principle be applied to any DAC, and it does not require a large amplitude control signal that impacts the SNR. This approach has previously been explored in Goodwin et al. (2003) and Quevedo et al. (2004a) for uniform quantisers, i.e. for DACs without INL. The authors introduced the moving horizon optimal quantiser (MHOQ), a method treating signal quantisation as a quadratic optimal control problem. It employs a finite constrained set and a finite prediction horizon while incorporating feedback using the moving horizon principle. This method significantly reduces error variance in a manner that correlates monotonically with the prediction horizon. Due to using uniform quantisation a closed-form solution to the problem can be obtained. However, the assumption of an ideal uniform quantiser in this method is restrictive, given that practical quantisers in DACs are non-uniform. As this paper demonstrates, applying the method to a DAC with INL provides almost no performance gain. Furthermore, the solution space grows exponentially with the number of quantisation levels and the prediction horizon length making the approach impractical for precision DACs with a higher number of quantisation levels.

In this paper, we therefore reformulate the solution by incorporating INL in the model. Modifying the model precludes a closed-form solution, and a suitable mixed-integer optimisation method is proposed to solve the problem, which also to some extent address the issue of the exponentially growing solution space. The paper demonstrates that mixed-integer optimisation provides identical results to the closed-form solution, and furthermore both simulation and experimental results demonstrate that the inclusion of INL in the model provide significant performance gains.

2. ERROR MODELLING

2.1 Uniform Quantisation

A quantiser can be represented by the block diagram in Fig. 1. Let $w = w(t)$ be the input signal, $y = y(t)$ be the



Fig. 1. Quantizer $\mathbf{Q}$ with input w and output y .




Fig. 2. Sinusoidal waveform, quantised waveform and the quantisation error.

quantised output signal, and $\mathbf{Q}$ represent the quantisation operation performed by the quantiser. Moreover, let δ be the quantiser step-size. An ideal DAC is modelled by $\mathbf{Q}$, and δ corresponds to the least significant bit (LSB), as DACs typically use binary encoding. For a word-size $\mathcal{B}$ (number of bits), a DAC has $2^{\mathcal{B}}$ quantisation levels and the step-size is determined by $\delta = \Delta / (2^{\mathcal{B}} - 1)$, where Δ is the full-scale output range of the DAC; see Fig. 2. A higher number of levels translates to a smaller quantisation error, but word-size is limited by several practical aspects. The uniform quantiser is defined as having equidistant levels. A uniform (mid-tread) quantiser $\mathbf{Q}$ can be defined using the truncating operator $\ell = T(w)$ as

$$y = \mathbf{Q}(w) = \delta \ell = \delta T(w) := \delta \left\lfloor \frac{w}{\delta} + \frac{1}{2} \right\rfloor, \quad (1)$$

where $\lfloor \cdot \rfloor$ denotes the floor operator and ℓ is referred to as the level (corresponding to the input w).

The quantisation of a signal introduces quantisation error, defined as

$$q = \mathbf{Q}(w) - w = y - w, \quad (2)$$

which is always constrained to be within $\pm\delta/2$. In Bennett (1948) quantisation is modelled as an additive, zero-mean and uniformly distributed white noise signal, independent of w , illustrated in Fig. 3. This approximation works well for large, frequency-rich signals, but poorly for small signals with few frequency components.




Fig. 3. Additive error model with input w , output y , and the quantisation error q .

2.2 Non-uniform Quantiser

Practical DACs has element mismatch that causes the actual levels to deviate from the ideal, equidistant levels of the uniform quantiser, as illustrated in Fig. 4. This is known as integral non-linearity (INL). Let y be the ideal




Fig. 4. DAC transfer function: Ideal and non-ideal.

and $\tilde{y}$ be the actual level of a DAC, and $\text{INL}(\ell)$ represent the deviation at the level $\ell \in \mathbb{N}$ (a static non-linearity), that is,

$$\tilde{y}(\ell) = y(\ell) + \delta \text{INL}(\ell) = \delta \ell + \delta \text{INL}(\ell) \quad (3)$$

The integral non-linearity $\text{INL}(\ell)$, can be represented as

$$\text{INL}(\ell) := \frac{\tilde{y}(\ell) - \delta \ell}{\delta}. \quad (4)$$

The effect of the non-linearity $n(w)$ on the output y due to the input w can be defined as

$$n(w) = \delta \text{INL}(\ell)|_{\ell=T(w)} = \tilde{y}(T(w)) - \delta T(w). \quad (5)$$

Note that the non-linearity $n(w)$ is a discontinuous function due to the truncation operator. The non-linear quantiser model is shown in Fig. 5.




Fig. 5. Nonlinear quantiser model.

3. SIMILAR SOLUTIONS

3.1 Noise-shaping Quantisers

Noise-shaping quantisers can reduce the effective quantisation error by moving quantisation noise to higher frequencies through oversampling and feedback. The recon-




Fig. 6. Noise-shaping quantiser block diagram.

struction filter is then used to attenuate the frequency-shaped quantisation noise. It operates by estimating the uniform quantisation error and employing a feedback filter to shape the noise power at the output of the DAC. A block diagram for a noise-shaping quantiser is shown in Fig. 6. The feedback filter $F(z)$ is designed such that the transfer function $y = (1 - F(z))\epsilon$ is a high-pass filter. By including the INL in the quantiser model, it is also possible to attenuate element mismatch error. This technique is called digital calibration (Cataltepe et al., 1989; Frey and Loeliger, 2007), and is the same type of modelling as employed in the presented moving horizon optimisation scheme.

It can be noted that in Quevedo et al. (2004a) it was found that the moving horizon optimal quantiser and the noise-shaping quantiser can be made to behave identically in the uniform quantisation case.

4. MOVING HORIZON OPTIMAL QUANTISER

The main contribution of this work is solving the moving horizon optimisation problem when using a quantisation model that includes INL; which allows a more realistic modelling of DACs. In contrast to noise-shaping, the main advantage of this formulation is that it is possible to determine the stability properties of the method, and it allows a more flexible modelling structure. The main disadvantage is a higher computational complexity, but with the rapid growth of computational power in embedded devices, this appears to be a diminishing problem.

In the moving horizon optimal quantiser, the problem is cast into the more general multi-horizon optimisation setting. However, to use such an optimisation framework a filter should be incorporated into the model (Goodwin et al., 2003) as shown in the Fig. 7.




Fig. 7. Moving horizon implementation with H denoting a filter.

In the sequel we will formulate an optimal control problem that will play the role of a uniform quantiser, in that setting the input w will be viewed as a reference signal. Let $\{w(t)\}_{t \in \mathbb{N}}$ be a sequence of sampled reference data, obtained from a continuous signal (in practise a floating-point representation). The quantisation requires this signal to be mapped to a finite signal $\{y(t)\}_{t \in \mathbb{N}}$ where each value of $y(t)$ is restricted to belong to a given finite set $\mathbb{U} = \{y_1, y_2, \dots, y_{n_u}\}$. We assume that the elements of $\mathbb{U}$ are equidistant (uniform quantiser).

Now let

$$H(z) = 1 + C(zI - A)^{-1}B \quad (6)$$

be a stable time-invariant linear filter with the matrices $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times 1}$, and $C \in \mathbb{R}^{1 \times n}$ related to the impulse response $h(i) = CA^{i-1}B$, $i > 0$ of the filter. The output e of the system in Fig. 7 can then be written as the output of the following state-space representation of H

$$\begin{aligned} x(t+1) &= Ax(t) + B(y(t) - w(t)) \\ e(t) &= Cx(t) + y(t) - w(t) \end{aligned} \quad (7)$$

where $x \in \mathbb{R}^n$ is the state vector. The error e corresponds to the difference between the filtered quantised signal and the filtered input signal. The error should be as small as possible; the quantiser should produce an output y optimally approximating the reference input w . To do so we formulate a uniform quantiser in terms of the following optimal control problem.

At time $t = k$ consider the quadratic cost defined by

$$V_N = \sum_{t=k}^{k+N-1} e^2(t) \quad (8)$$

where $e(t)$ is the error defined in equation (7). Then, the optimisation problem can be defined as the problem of finding $y \in \mathbb{U}$ that minimises the cost function (8) while satisfying the state equations (7), in summary,

$$y^*(t) = \arg \min_{y(t)} V_N \quad (9a)$$

subject to

$$x(t+1) = Ax(t) + B(y(t) - w(t)) \quad (9b)$$

$$e(t) = Cx(t) + y(t) - w(t) \quad (9c)$$

$$y(t) \in \mathbb{U}. \quad (9d)$$

Thus, the problem of finding the optimal quantisation level is re-cast into a moving horizon optimal control problem where the input to the system (i.e. the filter) has to be chosen from the finite constrained set i.e., $\mathbb{U}$. Precisely, it corresponds to choosing y^* so that the filtered signal y tracks the filtered reference signal w . The closed-loop system due to the moving horizon implementation is stable for a stable plant (i.e., filter) (Quevedo et al., 2004b).

The closed-form solution for the ideal case of the uniform quantisers is provided in (Goodwin et al., 2003). But it is essential to consider incorporating INL into the model when dealing with real-world DACs, as actual output levels always deviate from ideal counterparts.

Next, the moving horizon implementation where the INL is incorporated into the model is presented. INL is a static property, and it can be measured accurately. The measured data is organised into a lookup table; where the INL corresponding to a specific level is added to the ideal value. Fig. 8 shows a block diagram of the scheme.




Fig. 8. Synthesis and implementation of optimal quantizers in DAC with INL.

In the moving horizon implementation, we optimise over the finite horizon of size N to obtain the optimal values

$y^*(k), \dots, y^*(k+N-1)$. Then we add the INL corresponding to each level from the lookup table to obtain the actual values that are $\tilde{y}(k), \dots, \tilde{y}(k+N-1)$. That is,

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

As with the moving horizon implementation, we apply the first control $\tilde{y}(k)$ and move the optimisation to the next time horizon.

Error variance decreases monotonically with increasing prediction horizon N , but incur increased computational complexity and yields diminishing reduction of the error variance. The simulation and experimental results confirm this behaviour.

5. SIMULATIONS

Numerical simulations are presented for different combinations of methods, the DAC model used in the control scheme, and actual DAC response, as laid out in Tab. 1. The baseline is the case when using direct quantisation, i.e. just passing the input signal through the quantiser, whilst using a randomly generated INL (i.e., DQ-2 configuration). We can then see the effect of removing INL completely in the direct case, and when using noise-shaping with and without digital calibration, as well as MHOQ as formulated in Eq. (9), again with and without the inclusion of INL in the model used in the control scheme.

The MHOQ in equation (9) is formulated as a mixed integer quadratic programming problem (MIQP) where the state variable $x \in \mathbb{R}^n$ is continuous and the control variable $y \in \{0, 1, \dots, 255\}$ is integer for a 8-bit DAC. The MIQP problem is then solved numerically using Gurobi optimizer (Gurobi Optimization, LLC, 2023).

As shown in Kester (2005), the metric of performance for DACs is the signal-to-noise-and-distortion-ratio (SINAD), or equivalently effective number of bits (ENOB). SINAD is defined as the ratio of the standard deviations of the input signal and all other components,

$$\text{SINAD} = 20 \log_{10} \left(\frac{\sigma_s}{\sigma_t} \right)$$

where σ_s is the standard deviation of the input signal and σ_t accounts for all unwanted components in the output signal. The standard deviations are usually found via a power spectral density estimate. Note that the theoretical, lower bound of the SINAD for an oversampled signal is

$$\text{SINAD} = 6.02\mathcal{B} + 1.76 + 10 \log_{10}(\text{OSR}) \text{ dB} \quad (11)$$

where OSR is an oversampling ratio. The $\text{OSR} = f_s/2\text{BW}$ is the ratio of the sampling frequency f_s and desired

Table 1. Abbreviations for different combinations of method, DAC model used in control scheme, and actual DAC response.

Method	DAC Model	DAC Actual	Abbrev.
Direct	Without INL	Without INL	DQ-1
Direct	Without INL	With INL	DQ-2
Noise-shaping	Without INL	Without INL	NSQ-1
Noise-shaping	Without INL	With INL	NSQ-2
Noise-shaping	With INL	With INL	NSQ-INL
Moving horizon	Without INL	Without INL	MHOQ-1
Moving horizon	Without INL	With INL	MHOQ-2
Moving horizon	With INL	With INL	MHOQ-INL

bandwidth BW for the input signal. The ENOB can be obtained from (11) by assuming OSR = 1 and solving for the number of bits $\mathcal{B}$

$$\text{ENOB} = \frac{\text{SINAD} - 1.76}{6.02}.$$

5.1 Comparison to Noise-shaping

First we consider a numerical example from Goodwin et al. (2003), applied to an audio DAC, originally showing the equivalence of noise-shaping and MHOQ in the uniform quantiser case. For the sampling frequency $f_s = 44.1\text{kHz}$, the psycho-acoustic noise-shaping filter $F(z)$ is chosen as

$$F(z) = z^{-1} \frac{2.245 - 0.664z^{-1}}{1 + 0.91z^{-1}}.$$

Using $H(z) = 1/(1 - F(z))$, the equivalent stable time-invariant linear filter (6) is

$$H(z) = 1 + z^{-1} \frac{2.245 - 0.664z^{-1}}{1 - 1.335z^{-1} + 0.664z^{-2}}.$$

The simulation was carried out with a reference signal $w(t) = A \sin(2\pi ft)$ with $f = 3\text{ Hz}$ and $A = 16$ using a 4-bit quantiser, with quantisation levels = $\{0, 1, 2, \dots, 15\}$.

Fig. 9 compares the SINAD gains achieved through different quantisation methods: direct quantisation, noise-shaping and moving horizon optimal quantization (MHOQ) with respect to the baseline. DQ-1, employing direct quantisation, exhibit the poorest performance among all configurations. NSQ-1 and MHOQ-1 assumes the uniform quantiser in the model and implementation, resulting in highest SINAD gain, which, however are impractical and unattainable.




Fig. 9. Simulation: Comparison of the SINAD gain due to direct quantisation, noise-shaping and MHOQ relative to the baseline (SINAD[DQ-2] = 24.968) for the cases in Tab. 1.




Fig. 10. Effect of horizon N on MHOQ performance.

The NSQ-2 and MHOQ-2, designed assuming uniform quantiser but applied to the DAC with INL results in performance loss. The NSQ-INL and MHOQ-INL, which are obtained by incorporating the INL into the model, exhibit significant performance improvements compared to NSQ-2 and MHOQ-2 as shown in the Fig. 9. Moreover, the identical plots for NSQ and MHOQ ($N=1$) in Fig. 9 shows that they are equivalent (Goodwin et al., 2003). Fig. 10 plots SINAD gains for increasing horizons N . It can be seen that increasing N has amplifying effects on the gain.

Next, we present the numerical result for a typical scenario, and using the same configuration, we conduct experiments whose findings are detailed in the following section.

5.2 Moving Horizon Optimal Uniform Quantiser

Numerical results demonstrating the performance gains achieved by employing MHOQ and integrating INL into the model are presented. The simulations were carried out with a reference signal $w(t) = A \sin(2\pi ft)$ with $f = 1\text{ kHz}$ and $A = 256$, a sampling frequency of $f_s = 1\text{ MHz}$, using second-order Butterworth filter with cut-off frequency $f_c = 10\text{ kHz}$, and an 8-bit quantiser. The INL was set to be (nearly) identical to the one in the experimental setup in the next section (see Fig. 12).

Fig. 11 compares the SINAD gain using direct quantisation and MHOQ in different settings. Notably, MHOQ-1 demonstrates the highest SINAD gain, whereas MHOQ-2 shows the lowest performance. However, MHOQ-INL exhibits substantial performance improvement compared to MHOQ-2 and performs closely to the level achieved by MHOQ-1. Next, we present the experimental results under same settings.

6. EXPERIMENTS

6.1 Experimental Set-up

An experimental set-up was devised to verify the proposed method, ensuring the method has practical applicability and confirming a improvement in SINAD does occur when INL is the most significant contribution to the error. In order to minimise dynamic non-linear effects, a high-speed DAC was used in the experiments. A Texas Instruments DAC34SH84EVM evaluation card was modified to provide a voltage output using an OPA1611 op-amp as current-to-voltage converter (IV-converter) (Texas Instruments, 2014, Figure 35). Component values of the IV-converter feedback loop were selected to provide an output swing of $\pm 5\text{ V}$ and 100 kHz bandwidth. The data streams were generated on a Linux machine, and sent to the DAC using an AMD Virtex7 FPGA on the VC707 evaluation kit. The DAC34SH84 is a 16-bit DAC, but in order to ensure that INL is the main error contribution, an INL was emulated by re-quantising it to 8-bit, and using the additional resolution to induce an (almost) arbitrary INL. The INL was verified using a Keysight 34470A multimeter, and the performance was measured using a power spectral density estimate obtained by a National Instruments PXI-5922 oscilloscope. The INL applied in the experiments is shown in Fig. 12. A block diagram of the set-up is shown in Fig. 13. The experiments used identical configurations as in Sec. 5.2 and the results are presented in the Tab. 2.




Fig. 11. Simulation SINAD gains for various methods, relative to the baseline ($\text{SINAD}_{\text{DQ-2}} = 41.236$).




Fig. 12. Integral nonlinearity (INL) applied to the DAC.




Fig. 13. Experimental setup

The results in Fig 14 shows that incorporating the INL into the model in fact leads to the improved performance as the SINAD due to MHOQ-INL is close to the ideal case of MHOQ1. Even though 8-bit has been used, ENOB is 10.24 due to the oversampling. The SINAD gains in the experiments are lower than in the simulations. This is because the DAC has additional noise sources and distortions that are not taken into account in the model, which only includes the effect of INL.




Fig. 14. Experimental SINAD gains for various methods relative to the baseline ($\text{SINAD}_{\text{DQ-2}} = 41.470$).

Table 2. Experimental results.

Method	SINAD	SINAD Gain	ENOB
DQ-1	62.724	21.254	10.13
DQ-2 (Baseline)	41.470		6.59
MHOQ-1	68.562	27.092	11.10
MHOQ-2	47.102	5.631	7.53
MHOQ-INL	67.015	25.545	10.84

7. CONCLUSIONS

Digital-to-analog converters (DACs) suffer from integral non-linearity (INL), causing signal distortion due to output level discrepancies. To address this, we propose a novel method using moving horizon optimal control that integrates INL directly into the model. Unlike other approaches relying on ideal assumptions, our method uses precise INL measurements to create a lookup table within the model. Existing methods to mitigate INL often suffer from instability, complexity, and reliance on specialized hardware. Our approach, validated through simulations and experiments on a standard off-the-shelf DAC, shows promise in reducing INL-induced errors without requiring complex circuitry. Incorporating INL into the model via MHOQ offers a practical solution for improving DAC performance in real-world applications, paving the way for further advancements in DACs.

REFERENCES

- Abbott, B.P., Abbott, R., Adhikari, R., et al. (2009). Ligo: the laser interferometer gravitational-wave observatory. *Reports on Progress in Physics*, 72(7), 076901.
- Andersson, N.U., Andersson, K.O., Vesterbacka, M., and Wikner, J.J. (2003). Models and Implementation of a Dynamic Element Matching DAC. *Analog Integrated Circuits and Signal Processing*, 34(1), 7–16.
- Beckers, J.M. (1993). Adaptive Optics for Astronomy: Principles, Performance, and Applications. *Annual Review of Astronomy and Astrophysics*, 31(1), 13–62.
- Bennett, W.R. (1948). Spectra of quantized signals. *The Bell System Technical Journal*, 27(3), 446–472.
- Biswas, A., Bayer, I.S., Biris, A.S., Wang, T., Dervishi, E., and Faupel, F. (2012). Advances in top-down and bottom-up surface nanofabrication: Techniques, applications & future prospects. *Advances in Colloid and Interface Science*, 170(1–2), 2–27.

- Busgang, J., Ehrman, L., and Graham, J. (1974). Analysis of nonlinear systems with multiple inputs. *Proceedings of the IEEE*, 62(8), 1088–1119.
- Cataltepe, T., Kramer, A.R., Larson, L.E., Temes, G.C., and Walden, R.H. (1989). Digitally corrected multi-bit $\Sigma\Delta$ data converters. In *IEEE International Symposium on Circuits and Systems*, volume 1, 647–650.
- Eielsen, A.A. and Fleming, A.J. (2016). Experimental assessment of dynamic digital-to-analog converter performance for applications in precision mechatronic systems. In *Australian Control Conference*, 329–334.
- Eielsen, A.A. and Fleming, A.J. (2017). Existing methods for improving the accuracy of digital-to-analog converters. *Review of Scientific Instruments*, 88(9), 094702.
- Eielsen, A.A., Leth, J., Fleming, A.J., Wills, A.G., and Ninness, B. (2020). Large-Amplitude Dithering Mitigates Glitches in Digital-to-Analogue Converters. *IEEE Transactions on Signal Processing*, 68, 1950–1963.
- Fantner, G.E., Hegarty, P., Kindt, J.H., Schitter, G., Cidade, G.A.G., and Hansma, P.K. (2005). Data acquisition system for high speed atomic force microscopy. *Review of Scientific Instruments*, 76(2), 026118.
- Faza, A., Leth, J., and Eielsen, A.A. (2023). Criterion for Sufficiently Large Dither Amplitude to Mitigate Non-linear Glitches. In *IEEE Conference on Control Technology and Applications*, 970–977. IEEE.
- Frey, M. and Loeliger, H.A. (2007). On the Static Resolution of Digitally Corrected Analog-to-Digital and Digital-to-Analog Converters With Low-Precision Components. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 54(1), 229–237.
- Galton, I. (2010). Why Dynamic-Element-Matching DACs Work. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 57(2), 69–74.
- Goedhart, D., Plassche, R.J.v.d., and Stikvoort, E.F. (1982). Digital-to-analog conversion in playing a compact disc. *Phillips tech. Rev.*, 40(6), 174–179.
- Goodwin, G.C., Quevedo, D.E., and McGrath, D. (2003). Moving-horizon optimal quantizer for audio signals. *Journal of the Audio Engineering Society*, 51(3), 138–149.
- Groeneveld, D.W.J., Schouwenaars, H.J., Termeer, H.A.H., and Bastiaansen, C.A.A. (1989). A self-calibration technique for monolithic high-resolution D/A converters. *IEEE Journal of Solid-State Circuits*, 24(6), 1517–1522.
- Gurobi Optimization, LLC (2023). Gurobi Optimizer Reference Manual. URL <https://www.gurobi.com>.
- Hamilton, C.A. and Tang, Y.H. (1999). Evaluating the uncertainty of Josephson voltage standards. *Metrologia*, 36(1), 53–58.
- Kester, W. (ed.) (2005). *Data Conversion Handbook*. Elsevier.
- Leyonhjelm, S.A., Faulkner, M., and Nilsson, P. (1998). An Efficient Implementation of Bandlimited Dithering. *Wireless Personal Communications*, 8(1), 31–36.
- Neitola, M. (2017). Lee’s Rule Extended. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 64(4), 382–386.
- Pelgrom, M. (2017). *Analog-to-Digital Conversion*. Springer, 3rd edition.
- Quevedo, D.E., Goodwin, G.C., and De Doná, J.A. (2004a). Finite constraint set receding horizon quadratic control. *International Journal of Robust and Nonlinear Control*, 14(4), 355–377.
- Quevedo, D.E., Goodwin, G.C., and De Doná, J.A. (2004b). Finite constraint set receding horizon quadratic control. *International Journal of Robust and Nonlinear Control*, 14(4), 355–377. doi:10.1002/rnc.887.
- Texas Instruments (2014). OPA161x SoundPlus™ High-Performance, Bipolar-Input Audio Operational Amplifiers. URL <https://www.ti.com/lit/gpn/opa1611>.
- van de Plassche, R.J. (1976). Dynamic element matching for high-accuracy monolithic D/A converters. *IEEE Journal of Solid-State Circuits*, 11(6), 795–800.
- Zander, D.R. (1971). An Ultralinear 16-Bit Digital-to-Analog Current Converter. *Review of Scientific Instruments*, 42(6), 797–800.