Large-amplitude Dithering Mitigates Glitches in Digital-to-analogue Converters

Arnfinn A. Eielsen, John Leth, Andrew J. Fleming, *Member, IEEE,* Adrian G. Wills, and Brett Ninness,

***Abstract*—Glitches introduce impulse-like disturbances which are not be readily attenuated by low-pass filtering. This article presents a model that describes the behaviour of glitches, and a method for mitigation based on a large-amplitude dither signal. Analytical and experimental results demonstrate that a dither signal with sufficient amplitude can mitigate the effect of glitches, when used in conjunction with a low-pass filter. The dither signal in conjunction with low-pass filtering essentially converts a glitch from a high-frequency to low-frequency disturbance.**

***Index Terms*—digital-analogue conversion, glitch reduction, disturbance reduction, convolution, nonlinear distortion, dither**

I. Introduction

Physical implementations of digital-to-analogue converters (DACs) introduce various non-ideal effects, including element mismatch, thermal and semiconductor noise, slew-rate limitations, and glitches caused by non-ideal transistor switching [1]–[7]. These effects are in addition to the fundamental error sources of aliasing and quantisation that occur in a digital signal processing system, due to discretisation in both time and value [8].

DAC-errors deteriorate the achieved performance in various systems, including high-precision motion control [9]. Hence, it is of interest to mitigate them. Repeated spectra and aliasing is reduced by reconstruction filtering and interpolation [6, 10], and quantisation error is eliminated using small-scale dithering [8,11]–[14]. Element mismatch can be compensated for using several methods [15], including dynamic element matching [5,16] and large-scale dithering [17], and slewing can be partially reduced by oversampling [4,6].

This paper targets glitches. A behavioural model is developed and the effect of applying a dither signal is analysed: The glitches are modelled as short pulses and it is shown that it is possible to alter the effective behaviour of these pulses by applying dithering and low-pass filtering. Dithering and low-pass filtering provides mitigation in the sense that it effectively converts a glitch from a large amplitude disturbance with a short duration, to a smaller amplitude disturbance with a longer duration. This effect is established by deriving bounds for the expected value of dithered glitches.

This work was supported in part by the Australian Research Council Discovery Project DP120100487 and the Future Fellowship FT130100543.

A. A. Eielsen (eielsen@ux.uis.no) is with the Dept. of Electrical Engineering and Computer Science, University of Stavanger, Norway

J. Leth (jjl@es.aau.dk) is with the Dept. of Electronic Systems, Automation & Control, Aalborg University, Denmark

A. J. Fleming (andrew.fleming@newcastle.edu.au), A. G. Wills (adrian.wills@newcastle.edu.au), and B. Ninness (brett.ninness@newcastle.edu.au) are with the School of Electrical Engineering and Computer Science, The University of Newcastle, Australia.

By reducing the effective amplitude of glitches, dithering and averaging may mitigate unwanted impulsive loads on a system, and enable further suppression by feedback control. Considering the mitigated glitches as an input disturbance, they can asymptotically be removed via the internal model principle [18].

A glitch is typically described as the transient generated when a DAC switches between two output levels [3,33]. The glitch is commonly associated with timing errors within the DAC: the non-simultaneous action of the switches generating the output levels, which may be exacerbated by using binary weighting of the switches [19]–[22]. Glitches are usually measured as the area of the transient. Triggering of glitches is dependent on the input to the DAC. Glitches are typically reduced in the circuit; the most common methods are summarised in Tab. I. Note that these are hardware techniques that have to be applied during the design of a device, have limited applicability depending on the desired resolution, latency, bandwidth and other specifications, and may introduce additional non-ideal effects.

Many DACs suitable for high-precision motion control applications do not have sufficient glitch compensation [9], since applicable DACs are limited to devices that have both high accuracy and low latency. Low latency is required to avoid compromising the phase margin in closed loop applications. Furthermore, it is of interest to improve the performance in existing systems where replacing the DACs is not an option.

Glitches cause a form of inter-symbol interference (ISI). Variations of dynamic element matching (DEM) techniques have been shown to be effective in reducing ISI [16,21,29]–[32]; essentially by limiting the switching-rate between output elements. DEM techniques are limited to DAC topologies with redundancy in the representation of output values, and techniques for limiting the switching-rate are only applicable to very specific DAC topologies that utilise unary weighting (thermometer code) [16,21,29]–[32]. Furthermore, in segmented topologies [34], using e.g. binary weighing, DEM has demonstrated significant deterioration of performance [15]. Hence, the applicability of these methods is limited.

There are several behavioural models for glitches that occur in DACs [4,33,35]–[38]. While these models can provide excellent results in terms of predicting distortion spectra, power loss and simulating the effect of glitches in general circuits, they are not amenable to mathematical analysis of the effect of adding dithering and averaging, as discussed in this paper.

Dithering is an old technique in control theory [39], and has been used to improve stability in non-linear feedback control

TABLE I: Survey of glitch reducing methods.

Method	Description	Complexity	Performance
Unary (thermometer) weighting [19]–[22] *hardware level*	Binary weighting is area efficient, but prone to major carry glitches. Unary weighting minimises this behaviour, but at the expense of larger circuits and timing errors.	High	Medium
Segmentation [22,23] *hardware level*	A compromise between area usage, timing errors and major carry glitches is to combine unary and binary weighting.	Medium	Medium
Sample-and-hold, multiplexing [3,4,6,20,24] *hardware level*	Disconnecting and holding the output while it changes removes transients such as glitches, but can limit bandwidth and introduce linearity errors.	Medium	Medium
Optimised timing of switches [3,19,20,25] *hardware level*	Glitches are associated with timing errors between switches; optimising the time skew using delay elements reduces the error.	Low	Medium
Return-to-zero [4,6,19,20,24,26] *hardware level*	Switching using return-to-zero can reduce timing errors and suppress signal-dependent phenomena.	Medium	Medium
Differential current switches [19,26,27] *hardware level*	Reduces clock-feedthrough and current spikes, cancelling a current spike by generating an additional, opposite current spike. Enables faster switching.	Medium	Medium
Circuit component optimisation [23,28] *hardware level*	Glitches are affected by transistor sizes, switching voltages, output capacitances and transistor types.	Low	Low
Dynamic element matching [16,21,29]–[32] *firmware level*	Unary weighted DACs have redundancy in the representation of output values. Randomising the output elements for such DACs can minimise the effect of glitches.	High	High

systems [40]–[45]. More recently, dithering has been applied to non-linearities in digital-analogue converters: Mitigation of the harmonic distortion caused by uniform quantisation [8,11]–[14], mitigation of the effects due to the static non-linear transfer-function caused by element mismatch [17,46]–[51], and mitigation of various quantisation and non-linear effects in direct digital synthesis [52]–[54], digital phase-locked loops [55,56] and frequency synthesisers [57].

A. Contribution

A glitch model is developed which is intended to provide a description of glitches of any shape and to predict performance impact. The main differences between previous modelling efforts and the presented model are noted in Sec. II. However, the main feature of the presented model is the structure that makes it amenable to analysis using established theory for dithering non-linear systems [41,43,45]. Hence, the effect of applying stochastic and periodic dither signals is analysed in detail, and bounds for the expected value of the dithered glitch response are derived. It is demonstrated that these bounds can approximately be met by using sufficient low-pass filtering. The mitigation of glitches that occurs is particularly suitable for retrofitting to existing systems as no hardware modifications are required.

B. Notation

A definition is denoted by $\triangleq$ and $*$ denote the convolution product. The Laplace operator is denoted $\mathcal{L}$. Functions of time t are usually denoted by lower case, e.g. $g(t)$, and the Laplace transformed by upper case, e.g. $G(s) = \mathcal{L}[g](s)$. The Heaviside step-function is denoted $H(u)$ defined as

$$H(u) \triangleq \begin{cases} 0\,, & u \leq 0 \\ 1\,, & u > 0 \end{cases} . \tag{1}$$

The notation $L^p(\mathbb{R})$ indicate the L^p-space $p = 1, 2, \ldots, \infty$ and $||g||_p$ the p-norm of $g(t)$. For a stochastic process $d(t, \omega)$ the dependency on the sample variable ω is omitted. Convolution is defined as

$$(f * g)(t) \triangleq \int_{-\infty}^{\infty} f(\tau) g(t - \tau)\, \mathrm{d}\tau\,. \tag{2}$$

C. Outline

The remainder of this article is structured as follows: Sec. II presents the glitch model and an experimental validation of the response. Sec. III frames the analysis of the model in terms of existing theory for dithering (stochastic and periodic), and Sec. IV uses this framework to determine analytical properties of the dither model response for uniformly distributed dither signals (stochastic or periodic). Sec. V describes an experimental set-up used obtain measurements from a digital-to-analogue converter with significant glitches and Sec. VI presents the experimental results in order to validate the model response and to illustrate the character of the mitigation that occurs. A discussion of the results are presented in Sec. VII, and the paper is concluded in Sec. VIII.

II. Glitch Model

There are several models for glitches that occur in DACs. The model in [35] approximates glitch behaviour by opposite ramp-signals that are offset in time, and similar models are presented in [38] using piece-wise linear approximations. In [4,36] glitches are modelled as rectangular pulses, and in [37] a function comprised of hyperbolic and trigonometric expressions is used to approximate the glitch shape. The model in [33] superimposes two arbitrary functions to obtain a good glitch approximation, which also allows Fourier-analysis of the distortion. Unfortunately the formulation or structures of these models are not amenable to analysis of the effects of dithering using existing frameworks; hence a new model is developed in order to achieve this.

In the presented model, glitches are modelled as short, symmetric or asymmetric rectangular pulses driving a linear-time invariant (LTI) filter. If the response of a glitch can be measured with sufficient time resolution, then this approach can describe arbitrary glitch shapes, as an LTI filter can define any causal pulse response [58]. The LTI filter can be found using e.g. system identification techniques [59], or can be determined from first-principles if the components in the signal path are known. With an appropriately chosen LTI filter, the presented model will have an identical response, or be a close approximation, to previously published models.

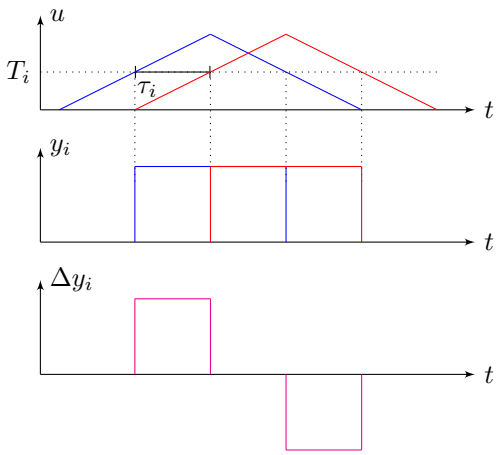

Fig. 1: Symmetric glitch model response to a triangle-wave.

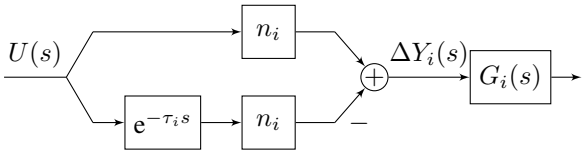

Fig. 2: Glitch sub-model for a symmetric response.

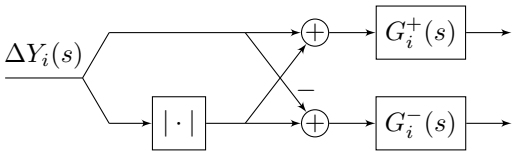

Fig. 3: Glitch sub-model for an asymmetric response.

A. Symmetric Glitch Model

Consider the signal $u(t)$, and the same signal $u(t-\tau_i)$ delayed by $\tau_i > 0$. If a DAC has N_T transition glitches, with polarity depending on the direction of crossing, at the thresholds T_i, and $T_i = T_j$ only when $i = j$, then the effect of the square pulses generated when crossing these thresholds can be described using the symmetric glitch model

$$n_g(u(t)) \triangleq \sum_{i=1}^{N_T} (g_i * \Delta y_i)(t)\,. \tag{3}$$

Here, the square pulses are constructed using

$$\Delta y_i(t) \triangleq y_i\,(t) - y_i\,(t-\tau_i)\ , \tag{4}$$

where

$$y_i(t) \triangleq n_i\,(u(t)) \tag{5}$$

are the outputs of the functions $n_i(u)$. These functions (or non-linearities) are used to generate the pulses triggered at the thresholds T_i as

$$n_i(u) = \frac{1}{\tau} H(u - T_i)\,. \tag{6}$$

The construction of $\Delta y_i(t)$ is illustrated in Fig. 1. If $g_i = 1$ this model is equivalent to the models in [4,36], and since the functions n_i can be chosen arbitrarily, with appropriate choices, the model can be made identical to the models presented in [33,35,37].

Note that in Sec. III and onward the input signal $u(t)$ is the combination of two different signals

$$u(t) = x(t) + d(t)\,, \tag{7}$$

where $x(t)$ is the deterministic signal that should be reconstructed by the DAC. The dither signal $d(t)$ is added to smooth the glitch responses and can be either stochastic or periodic, but is otherwise unwanted in the output.

The rectangular pulse $\Delta y_i(t)$ has unit area. The shape of the glitch (3) is determined by the linear time-invariant (LTI) filter $G_i(s) = \mathcal{L}[g_i](s)$. That is, the shape of a glitch can be described by the response of the filter $G_i(s)$ due to the pulse $\Delta y_i(t)$. Recall that an LTI filter can define any causal pulse response [58]. As the time-delay $\tau_i \to 0$, the response converges to the impulse response. The net area A_i of a glitch,

$$A_i \triangleq \lim_{t\to\infty} \int_0^t \int_{-\infty}^{\infty} g_i(\xi - w)\Delta y_i(w)\mathrm{d}w\mathrm{d}\xi \tag{8}$$

$$= \lim_{s\to 0} s\frac{1}{s}\frac{1}{\tau_i}\frac{1-\mathrm{e}^{-\tau_i s}}{s} G_i(s) = G_i(0)\,, \tag{9}$$

is determined by the DC-gain of $G_i(s)$. The area of this response is typically referred to as the glitch energy [33]. A block diagram of the model is shown in Fig. 2.

B. Asymmetric Glitch Model

If the glitches are asymmetric, i.e. the amplitude is different depending on whether the input is falling or rising, then, with the same assumptions as before, the effect of the square pulses generated when crossing the thresholds can be described using the asymmetric glitch model

$$\tilde{n}_g(u(t)) \triangleq \sum_{i=1}^{N_T} \Big((g_i^+ * y_i^+)(t) + (g_i^- * y_i^-)(t) \Big)\,, \tag{10}$$

where a rising pulse is described by

$$y_i^+(t) \triangleq \frac{1}{2}\left(|\Delta y_i(t)| + \Delta y_i(t)\right) \geq 0\,, \tag{11}$$

and a falling pulse is described by

$$y_i^-(t) \triangleq \frac{1}{2}\left(|\Delta y_i(t)| - \Delta y_i(t)\right) \geq 0\,. \tag{12}$$

A block diagram of the model is shown in Fig. 3 and the construction of $y_i^+(t)$ is illustrated in Fig. 4. The LTI filters $G_i^+(s)$ determine the glitch shape when $u(t)$ is rising, and $G_i^-(s)$ when $u(t)$ is falling. The model is equivalently expressed as

$$\tilde{n}_g(u(t)) = \frac{1}{2}\sum_{i=1}^{N_T} \Big(\left((g_i^+ - g_i^-) * \Delta y_i\right)(t) + \left((g_i^+ + g_i^-) * |\Delta y_i(t)|\right)(t) \Big)\,. \tag{13}$$

If $g_i^-(t) = -g_i^+(t)$, the symmetric model (3) is obtained.

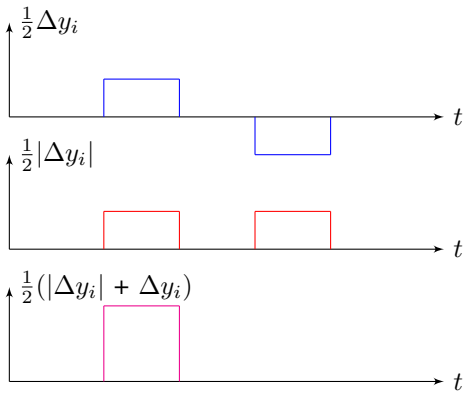

Fig. 4: Asymmetric glitch model response to a triangle-wave.

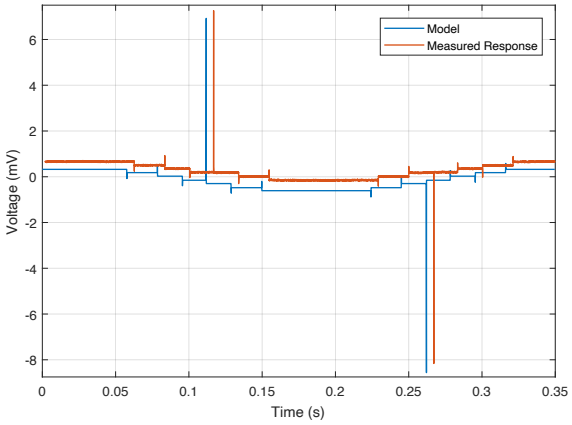

Fig. 5: Measured and model response (13) for a TI DAC8544, using a sinusoidal input signal with an amplitude of 0.01% (6.5535) of full range (65535 → 20 V). Plots offset for clarity.

C. Model Validation

The response of the model is illustrated in Fig. 5 where it is compared to the measured response of a DAC channel in the experimental set-up. The DAC is the 16-bit Texas Instruments DAC8544. The set-up is described in detail in Sec. V. The maximum supported sampling rate of 1 MS/s was used. It has large transition glitches at code intervals of 4096, hence there are 16 transition glitches over the full range. There is a minor glitch for every least-significant bit. The filters $G_i^{\pm}(s)$ are in this case determined to be on the form

$$G_i^{\pm}(s) \triangleq \frac{\hat{A}_i^{\pm}}{\tau} W(s)\,. \tag{14}$$

The areas $\hat{A}_i^{\pm}$ are set to $\hat{A}_i^{+} = -60.6$ nVs for the major glitches with rising input, $\hat{A}_i^{-} = 51.9$ nV for the major glitches with falling input, and $\hat{A}_i^{\pm} = \pm 2.40$ nVs for the minor glitches. The glitches are likely generated during code transitions where the new code signal value appears before or after the signal value of the former code disappears [4]. Depending on the exact implementation, such code dependent transitions may not exhibit symmetry in response.

The glitch energy at mid-code was estimated using the trapezoidal method, as illustrated in Fig. 6, where the model response of a major glitch when the input is rising is compared

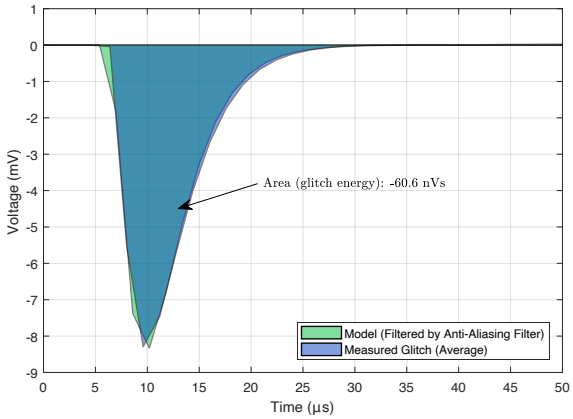

Fig. 6: A measured glitch and the corresponding model response. In this case the model glitch duration τ is set to the sampling time of the ADC. The measurements were averaged over several periods of the input signal.

to the measured response. The areas $\hat{A}_i^{\pm}$ of the remaining 15 glitches were found to be approximately the same. Since the mid-code glitch was primarily excited during experiments, an identical area for all major glitches provided sufficient accuracy. Furthermore, $\tau = 1/625$ μs, which is the sampling time of the analogue-to-digital converter (ADC) (an Analog Devices AD7674) with a maximum supported sampling rate of 625 kS/s. This choice was intended to match the sampling times between measurement and simulation, to avoid downsampling the simulation results. However, running the simulation at 1 MS/s, or higher, and resampling to 625 kS/s yielded near identical results. The filter $W(s)$ is the anti-aliasing filter,

$$W(s) \triangleq \left(\frac{2\pi f_c}{s + 2\pi f_c}\right)^2, \tag{15}$$

with cut-off frequency $f_c = 62.5$ kHz.

As the plot in Fig. 6 indicates, the accuracy of the model is good in terms of predicting the area of the glitch as well as the shape; in this case mainly determined by the energy of the glitch exciting the dynamics of the anti-aliasing filter. Since the response was measured using an ADC with a different sampling rate (and not synchronised with the DAC), the switching times are not predicted exactly. The integral non-linearity (INL) of the DAC was measured and used to predict the static voltage levels, but there is some error in the these measurements which is neglected.

III. Types of Dither

In general, there are two types of dither signals: stochastic and periodic. Both achieve similar effects, and both deliver identical results in terms of the expected value [41]. In the following we determine that the expected value can be expressed via an averaged non-linearity, and consider the error in the output introduced by dithering. By removing the error, the ideal averaged response is obtained. Analysis of the effect of the model response due to dither signals is achieved using theory developed for dithering of non-smooth

systems [45] for periodic signals, and well-known results from probability theory [60] for stochastic signals. As we do not aim for complete generality, we assume in the upcoming Sections the existence of various densities. However, the result can without much difficulty be formulated in terms of the corresponding distribution functions. To avoid unnecessary notation we consider DACs with only one transition glitch in this Section: the subscript i is omitted.

A. Averaging Due to Stochastic Dither

The first case to be considered is where the dither signal in (7) is of a stochastic nature. By a stochastic dither signal $d(t)$ we understand an identically distributed stochastic process having the strictly white noise property, i.e. $d(t_i)$ and $d(t_j)$ are independent for each pair t_i and t_j. We follow [60] on the subject of white noise — for a rigorous mathematical exposition see [61]. In this Section we let $n(u)$ denote a function of bounded variation (on $\mathbb{R}$). It will be thought of as a threshold, e.g. as (6), and referred to as a non-linearity.

Consider the strictly white noise input signal

$$u(t) \triangleq x(t) + d(t)\,, \tag{16}$$

defined as the sum of a deterministic signal $x(t)$ and a stochastic dither signal $d(t)$. Let $f_d(z)$ be the probability density function corresponding to $d(t)$, and $f_x(z,t) = \delta(z - x(t))$ the probability density function corresponding to $x(t)$. Here, $\delta(\cdot)$ denotes the Dirac delta function. The signal $u(t)$ will then have probability density function

$$f_u(z,t) = \int_{\mathbb{R}} f_x(z-w,t) f_d(w)\,\mathrm{d}w = f_d(z - x(t))\,. \tag{17}$$

Consider the strictly white noise output

$$y(t) = n(u(t)) = N(x(t)) + \varepsilon(t)\,, \tag{18}$$

where the application of the dither signal makes it possible to define an averaged[1] non-linearity

$$N(z) \triangleq \int_{\mathbb{R}} n(z+w) f_d(w)\,\mathrm{d}w\,, \tag{19}$$

related to $n(u)$ and the stochastic dither signal $d(t)$, but also introduces an error term $\varepsilon(t)$. Here,

$$\|N\|_\infty \leq \|f_d\|_1 \|n\|_\infty = \|n\|_\infty \tag{20}$$

since $u \in L^\infty(\mathbb{R})$ by the bounded variation property of $n(u)$. The average non-linearity thus reduces the effect caused by the non-linearity. Moreover, if $L_D \triangleq \sup_{z\in\mathbb{R}} |f_d(z)| < \infty$ then by [45, Lemma A.1]

$$L_N \leq L_D TV(n) \tag{21}$$

with L_N a Lipschitz constant of $N(z)$ and $TV(n)$ the total variation of $n(u)$.

Now, by using (19) the expected value of the output is

$$\begin{aligned} \mathrm{E}[y(t)] &= \int_{\mathbb{R}} n(u) f_d(u - x(t))\,\mathrm{d}u && \text{(22a)}\\ &= \int_{\mathbb{R}} n(x(t)+w) f_d(w)\,\mathrm{d}w && \text{(22b)}\\ &= N(x(t))\,, && \text{(22c)} \end{aligned}$$

[1] sometimes also referred to as "effective" or "smoothed"

leading to the output auto-covariance

$$\begin{aligned} C_y(t,s) &\triangleq \mathrm{E}[y(t)y(s)] - \mathrm{E}[y(t)]\mathrm{E}[y(s)] && \text{(23a)}\\ &= \begin{cases} 0 & t \neq s \\ \mathrm{E}[y^2(t)] - N^2(x(t)) & t = s \end{cases} && \text{(23b)}\\ &\triangleq q(t)\chi_{\{t\}}(s) && \text{(23c)} \end{aligned}$$

where $\chi_{\{t\}}(s)$ denotes the indicator function, and the first case in (23b) follows by the strictly white noise property of $y(t)$ and the second case follows by (22c).

Consider the zero-mean strictly white noise error signal

$$\varepsilon(t) = y(t) - N(x(t))\,, \tag{24}$$

having auto-covariance $C_\varepsilon(t,s) = C_y(t,s)$. The time-varying variance of the error signal is thus

$$\begin{aligned} q(t) &= \mathrm{E}[\varepsilon^2(t)] \\ &= \int_{\mathbb{R}} n^2(u) f_d(u - x(t))\,\mathrm{d}u - N^2(x(t))\,. \end{aligned}$$

That is, for the non-linearity (6) it follows that $n^2(u(t)) = n(u(t))/\tau$, and the expression for $q(t)$ can be found as

$$q(t) = \frac{1}{\tau} N(x(t)) - N^2(x(t))\,. \tag{25}$$

In order to approximately obtain the averaged response in the output $y(t)$, the variance $q(t)$ must be reduced.

A well-known result in linear system theory is that for a stationary white noise input signal w, with auto-covariance $C_w(t,s) = {\sigma_w}^2 \chi_{\{t\}}(s)$ where ${\sigma_w}^2$ is a constant, the variance of the response of the output v of an LTI-filter $G(s)$ is [62]

$${\sigma_v}^2 = \|G(s)\|_2 {\sigma_w}^2\,, \tag{26}$$

where $\|G\|_{\mathcal{H}_2}$ denotes the $\mathcal{H}_2$ norm of $G(s)$. Due to the frequency selective properties of an LTI-filter, it is possible to devise a filter that has smaller variance on the output than the input, i.e. reduces the noise. A typical signal chain contains several low-pass filters; in particular the reconstruction filter (needed to satisfy the NyquistShannon theorem). This filtering can be expressed by the filters $g_i(t)$, $g_i^{\pm}(t)$ in the models (3) and (13). These filters can be expected to reduce the unwanted signal component due to the dither signal.

Now, the variance (25) of the error signal (24) is not stationary, and (26) does not apply. However, an LTI-filter will reduce the variance of this signal: Consider the impulse response $g(t)$ corresponding to the LTI filter $G(s)$ and define the filtered error signal

$$\bar{\varepsilon}(t) \triangleq (g * \varepsilon)(t)\,,$$

as the output of the filter $g(t)$ due to the error signal $\varepsilon(t)$. The variance of the filtered error signal is then

$$\begin{aligned} C_{\bar{\varepsilon}}(t,t) &= \mathrm{E}[\bar{\varepsilon}^2(t)] && \text{(27)}\\ &= (q * g^2)(t) && \text{(28)}\\ &= \int_{\mathbb{R}} q(t-w) g^2(w)\,\mathrm{d}w\,, && \text{(29)} \end{aligned}$$

where (28) follows by [60, Theorem 9-3]. Hence if $q \in L^1(\mathbb{R})$ and $g^2 \in L^2(\mathbb{R})$ then

$$\|C_{\bar{\varepsilon}}\|_2 = \|q * g^2\|_2 \tag{30a}$$
$$\leq \|q\|_1 \|g^2\|_2 \tag{30b}$$
$$= \|q\|_1 \|\mathcal{L}g^2\|_{\mathcal{H}_2} \tag{30c}$$

where (30b) follows by Young's convolution theorem [63] and (30c) follows by Parseval's theorem. By (30) the variance of the filtered error signal will be attenuated by the energy of the impulse response squared $\left\|g^2\right\|_2$. For strictly proper, unity gain low-pass filters, the energy of the impulse response is directly related to the cut-off frequency: A lower cut-off frequency yields lower energy, hence smaller variance for the filtered error signal $\bar{\varepsilon}$ and the filtered output will in this sense approach the ideal averaged response, but the effect the filter has on the signal to reconstruct, $x(t)$, must also be taken into account.

B. Averaging Due to Periodic Dither

In terms of smoothing, similar results are obtained using either stochastic or periodic dither signals. In the stochastic case, low-pass filtering on the output reduces the variance of the non-stationary error signal. For the periodic case it is the dither frequency and how well the low-pass filtering on the output resembles a windowed integrator that determines the residual error.

In this case (16) consists of a deterministic signal $x(t)$ and a periodic dither signal $d(t)$ with sufficiently high frequency and amplitude. The averaged non-linearity (19) is also defined for periodic dither signals [41,43,45]. In this case $f_d(z)$ is the amplitude density function of the periodic dither signal $d(t)$, which is the deterministic equivalent to the probability density function. As in the case of stochastic dither, the average non-linearity induced by a periodic dither reduces the effect caused by the non-linearity, as determined by (21).

For each $t \geq 0$ let $J_t = [t-\rho, t]$ denote the interval of length $\rho > 0$. The method of averaging via a periodic dither signal, with period ρ, relies on the equivalence (see [45] for a reference)

$$\int_{\mathbb{R}} n(z+w) f_d(w)\,\mathrm{d}w = \frac{1}{\rho}\int_{J_t} n(z+d(\tau))\,\mathrm{d}\tau \tag{31}$$

between the average non-linearity and a time-average over one period J_t of the periodic dither. Hence, averaging can be implemented by using a windowed integrator (a continuous-time moving average). From a practical point of view this is advantageous, as there is no readily apparent way to implement a physical realisation of the convolution product on the left-hand side of (31), but an integrator like the right-hand side of (31) can be realised using analogue circuitry. This realisation can take the form in (34), shown below, but as the dither frequency is finite, an error is introduced.

To make this more precise recall first from [45] that; if $x(t)$ has a Lipschitz constant $L_X = L_X(J_t)$ on J_t then

$$\left|\int_{J_t} n(x(\tau)+p(\tau))\,\mathrm{d}\tau - \int_{J_t} n(\tilde{x}+p(\tau))\,\mathrm{d}\tau\right| \leq 2L_D L_X TV(n)\rho^2 \tag{32}$$

with $\tilde{x}$ satisfying

$$\min_{t\in J_t} x(t) \leq \tilde{x} \leq \max_{t\in J_t} x(t)\,. \tag{33}$$

As with (31) the error bound (32) holds for any interval of length ρ. Using (19) and (31) and

$$\alpha(t) \triangleq \frac{1}{\rho}\int_{J_t}\gamma(\tau)\mathrm{d}\tau \triangleq \frac{1}{\rho}\int_{J_t} n(x(\tau)+p(\tau))\mathrm{d}\tau, \tag{34}$$

the above result (32) yields

$$|\alpha(t) - N(x(t))| = \frac{1}{\rho}\left|\int_{J_t}\gamma(\tau)\,\mathrm{d}\tau - \int_{J_t} n(x(t)+p(\tau))\,\mathrm{d}\tau\right| \leq 2L_D L_X TV(n)\rho \tag{35}$$

with $L_X = L_X(J_t)$ a Lipschitz constant of $x(t)$ on J_t. If $\bar{L}_X \triangleq \sup_{t\in[0,\infty)} L_X(J_t) < \infty$, the error bound (35) implies the error bound

$$|A(s) - G_N(s)| \leq 2L_D\bar{L}_X TV(n)\rho/\sigma\,,\ \sigma > 0 \tag{36}$$

between the Laplace transformed $A(s)$ and $G_N(s)$ of $\alpha(t)$ and $N(x(t))$, respectively, with $s = \sigma + i\omega$. The error bound (36) can further be elaborated on to obtain a conclusion in analogue with (30). For this, let $r(u) \triangleq H(u+\rho) - H(u)$ denote a rectangle function expressing the interval J_t of width ρ, with $H(u)$ the Heaviside step function (1). The signal $\alpha(t)$ can then be expressed as

$$\alpha(t) = \frac{1}{\rho}\int_0^\infty r(t-\tau)\gamma(\tau)\mathrm{d}\tau = \frac{1}{\rho}(r*\gamma)(t).$$

The Laplace transformed $A(s)$ is thus given by

$$A(s) = \frac{1}{\rho}R(s)\Gamma(s) = \frac{1-\mathrm{e}^{-\rho s}}{\rho s}\Gamma(s) \tag{37}$$

with $R(s)$ and $\Gamma(s)$ the Laplace transformed of $r(t)$ and $\gamma(t)$ respectively.

It is possible to realise a windowed integrator (34), which frequency response is expressed by $R(j\omega)/\rho$, but due to practical limitations of the resulting analogue circuitry, the actual performance will not be satisfactory. However, as noted in Sec. III-A, a typical signal chain contains several low-pass filters. The frequency response $|R(j\omega)/\rho|$ is the sinc-function, which has a low-pass characteristic. If the filters in the circuit resemble the sinc-function, i.e. $\||G(j\omega)| - |R(j\omega)/\rho|\|_2$ is small, then the error bound (36) is a good approximation and

$$A'(s) \triangleq G(s)\Gamma(s) \tag{38}$$

therefore represents an approximation of $G_N(s)$ which is implementable. It can be concluded that to minimise the error when using a periodic dither, the dither frequency should be as high as practically possible, as it makes the period ρ and hence the bound (36) small, and the frequency response of the low-pass filter on the output should resemble the corresponding sinc-function. In applications, high-frequency dither signals may cause additional non-linear distortion, such as slewing. A low-pass filter has limited attenuation, hence it is not possible to completely remove the dither signal. Furthermore, the glitches, as well as other non-linearities such as the integral non-linearity (INL), will produce inter-modulation causing frequency components to appear below

the fundamental frequency of the dither signal. Such effects must be taken into account when determining the potential improvement from dithering.

IV. Properties

In this Section we study the properties of a stochastic and periodic dither signals with uniform density applied to the glitch model. An averaged response is found in the symmetric case, and a response bound is found for the asymmetric case.

A. Dither With Uniform Distribution

Assume that the dither signal $d(t)$ in the stochastic case is uniformly distributed white noise where $d(t) \in [-A, A]$, and in the periodic case is a triangle wave with amplitude A. Then both the probability density function and the amplitude density function will equal a rectangle function $f(z)$ as

$$f(z) = \frac{1}{2A}\Big(H(z/A+1) - H(z/A-1)\Big) \tag{39a}$$

$$= \begin{cases} 0\,, & |z| \geq A \\ \frac{1}{2A}\,, & |z| < A \end{cases}\,, \tag{39b}$$

with Lipschitz constant $L_D = 1/(2A)$. Using (21) for the non-linearity in (6), with total variation $TV(n_i) = 1/\tau_i$, the average non-linearity $N_i(z)$ is bounded in growth by the Lipschitz constant $L_{N_i} = 1/(2A\tau_i)$.

Since $f(z)$ is even, (19) is equivalent to

$$N_i(x(t)) = \int_{\mathbb{R}} n_i(\xi) f(x(t) - \xi) d\xi\,. \tag{40}$$

Hence, for a non-linearity (6) with $T_i = 0$, the average non-linearity is

$$N_i(x(t)) = \frac{1}{2A\tau_i}\int_{\mathbb{R}} H(\xi)\Big(H(x(t) - \xi + A) \tag{41a}$$

$$- H(x(t) - \xi - A)\Big) d\xi \tag{41b}$$

$$= \frac{1}{2A\tau_i}\Big((x(t)+A)H(x(t)+A) \tag{41c}$$

$$- (x(t)-A)H(x(t)-A)\Big) \tag{41d}$$

$$= \begin{cases} 0\,, & x(t) \leq -A \\ \frac{1}{2A\tau_i}(x(t)+A)\,, & |x(t)| \leq A \\ \frac{1}{\tau_i}\,, & x(t) \geq A \end{cases}\,, \tag{41e}$$

and therefore $N_i(x(t)) \in [0, 1/\tau_i]$. Fig. 7 illustrates the effect: The non-linearity $n_i(u(t))$ is shown, as well as $(g_i * n_i(x + d))(t)$ representing the response of a low-pass filter $g_i(t)$ applied to the output $y_i(t) = n_i(x(t) + d(t))$ when using a uniform white-noise dither $d(t)$. Moreover, $N_i(x(t))$, given by (41), is also shown which according to (22c) correspond to the expected value $E[y_i(t)]$. The effect of noise-modulation (25) is noticeable, as the variance is larger around zero, than towards the dither amplitude limits at $-A$ and A. From Fig. 7 it can also be seen that the expected value $E[y(t)]$ of the output is close to the filtered mean $E[(g_i * y_i)(t)] = (g_i * E[y_i])(t)$.

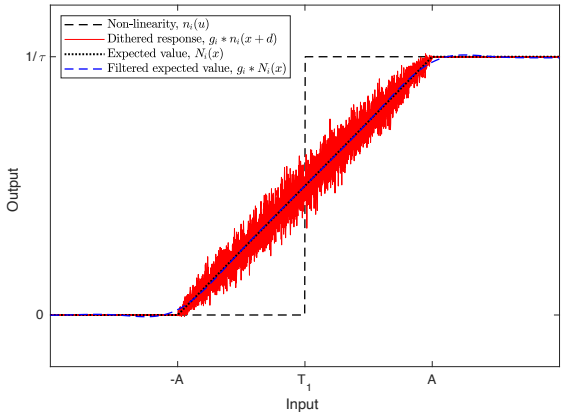

Fig. 7: Average non-linearity: Expected value $E[y_i(t)] = N_i(x(t))$ and the filtered expected value $(g_i * N_i(x))(t)$ compared to simulated response when using uniform white-noise dither.

1) Symmetric Glitches: In the case of symmetric glitches it is straightforward to construct an average model of (3), or equivalently (13) with $g_i^-(t) = -g_i^+(t)$. Consider first the mean value

$$\begin{aligned} E[(g_i * y_i)(t)] &= E\left[\int_{\mathbb{R}} g_i(w) y_i(t-w)\,\mathrm{d}w\right] \\ &= \int_{\mathbb{R}} g_i(w) E[y_i(t-w)]\,\mathrm{d}w \\ &= (g_i * E[y_i])(t) \\ &= (g_i * N_i(x))(t)\,. \end{aligned}$$

Applying this to each term in the model (3) yields

$$E[(g_i * \Delta y_i)(t)] = (g_i * \Delta N_i)(t)$$

with

$$\Delta N_i(t) \triangleq E[y_i(t) - y_i(t-\tau_i)] \tag{42a}$$

$$= E[y_i(t)] - E[y_i(t-\tau_i)] \tag{42b}$$

$$= N_i(x(t)) - N_i(x(t-\tau_i)) \tag{42c}$$

The averaged model corresponding to (3) is therefore

$$E\left[n_g\left(u(t)\right)\right] = \sum_{i=1}^{N_T} (g_i * \Delta N_i)(t). \tag{43}$$

For practical implementation it is important to note that (for a suitable choice of dither amplitude A) the average symmetric glitch model (43) can be approximated by a linear combination of the averaged non-linearities — see (44). To see this assume that for a given time-delay τ_i, the dither amplitude A is chosen such that $E[y_i(t)] = N_i(x(t))$ can be considered approximately constant (for the duration of the time-delay, i.e. slowly varying compared to τ_i). Typically, τ_i is small and A therefore needs to be large, according to (41). It follows that

$$\begin{aligned} E[(g_i * y_i)(t)] &= \int_{\mathbb{R}} g_i(w) E[y_i(t-w)]\,\mathrm{d}w \\ &\approx \int_{\mathbb{R}} g_i(w)\,\mathrm{d}w N_i(x(t)) \\ &= A_i N_i(x(t))\,, \end{aligned}$$

with $A_i = G_i(0)$ as in (8). The mean of each term in the symmetric glitch model (3) can then be approximated as

$$E[(g_i * \Delta y_i)(t)] \approx A_i \Delta N_i(t)$$

where $\Delta N_i(x(t))$ typically is close to zero as the dither amplitude A is chosen very large. The average symmetric glitch model (43) is therefore approximately given by

$$E\left[n_g\left(u(t)\right)\right] \approx \sum_{i=1}^{N_T} A_i \, \Delta N_i(t) \, . \tag{44}$$

2) Asymmetric Glitches: In the case of the asymmetric glitch model (13) it is not possible to find an average model on the same form as in (43) but it is possible to find a bound for the response. This is due to the need to take the absolute value of the square pulses $\Delta y_i(t)$ in (13), which means the output cannot be constructed as a linear combination of the outputs of the non-linearities, as is the case in (3), and we resort to using the triangle inequality. For the first term in (13)

$$\left((g_i^+(t) - g_i^-) * \Delta y_i\right)(t) \triangleq (\tilde{g}_i(t) * \Delta y_i)(t)$$

we proceed as in the symmetric case to obtain

$$E[(\tilde{g}_i * \Delta y_i)](t) = (\tilde{g}_i * \Delta N_i)(t)$$

with the approximation (for large A)

$$E[(\tilde{g}_i * \Delta y_i)](t) \approx \tilde{A}_i \Delta N_i(t)$$

where $\tilde{A}_i = A_i^+ - A_i^-$ and $A_i^\pm = \tilde{G}_i^\pm(0)$. For the second term

$$(g_i^+(t) + g_i^-(t)) * |\Delta y_i(t)| \triangleq \bar{g}_i(t) * |\Delta y_i(t)|$$

we derive two bounds. For the first bound the triangle inequality is applied as

$$|\Delta y_i(t)| = |y_i(t) - y_i(t-\tau)_i| \leq |y_i(t)| + |y_i(t-\tau_i)| \tag{45}$$

to obtain

$$\begin{aligned} E[(\bar{g}_i * |\Delta y_i|)](t) &\leq (\bar{g}_i * E[|y_i|])(t) + (\bar{g}_i * E[|y_i|])(t-\tau_i) \\ &\triangleq B_i^+(t). \end{aligned} \tag{46}$$

Note that since the function $n_i(z)$, given by (6), is positive

$$\begin{aligned} E[|y_i(t)|] &= \int_{\mathbb{R}} |n_i(z)| f(z - x(t)) \, \mathrm{d}z & (47)\\ &= \int_{\mathbb{R}} n_i(z) f(x(t) - z) \, \mathrm{d}z & (48)\\ &= N_i(x(t)) \, , & (49) \end{aligned}$$

with $f(z)$ given by (39). Hence the bound (46) can equivalently be expressed as

$$B_i^+(t) = (\bar{g}_i * N_i(x))(t) + (\bar{g}_i * N_i(x))(t-\tau_i) \tag{50}$$

and for large dither amplitude A be approximated by

$$B_i^+(t) \approx \bar{A}_i \Big(N_i(x(t)) + N_i(x(t-\tau_i)) \Big) \triangleq \bar{B}_i^+(t) \tag{51}$$

where $\bar{A}_i = A_i^+ + A_i^-$.

For the second bound we note that $n(u) = -n(-u) + 1$, hence similar to (45) we apply the triangle inequality to obtain

$$\begin{aligned} |\Delta y_i(t)| &= |n_i(u(t)) - n_i(u(t-\tau_i))| \\ &= |n_i(-u(t)) - n_i(-u(t-\tau_i))| \\ &\leq |n_i(-u(t))| + |n_i(-u(t-\tau_i))| \end{aligned}$$

and therefore also a second bound

$$\begin{aligned} E[(\bar{g}_i * |\Delta y_i|)](t) &\leq (\bar{g}_i * E[|n_i(-u)|])(t) \\ &\quad + (\bar{g}_i * E[|n_i(-u)|])(t-\tau_i) \\ &\triangleq B_i^-(t) \end{aligned} \tag{52}$$

similar to (46). Moreover,

$$B_i^-(t) = (\bar{g}_i * N(-x))(t) + (\bar{g}_i * N(-x))(t-\tau_i) \tag{53}$$

since $E[|n_i(-u)|] = N(-x(t))$, hence for large A

$$B_i^-(t) \approx \bar{A}_i \Big(N_i(-x(t)) + N_i(-x(t-\tau_i)) \Big) \triangleq \bar{B}_i^-(t) \tag{54}$$

Collecting the two bounds (46) and (52) in

$$B_i(t) = \min_t \{B_i^+(t), B_i^-(t)\},$$

an upper bound for the mean of the asymmetric glitch model (13) is

$$\begin{aligned} E[\tilde{n}_g(x(t))] &\leq \frac{1}{2} \sum_{i=1}^{N_T} (\tilde{g}_i * \Delta N_i)(t) + B_i(t) & (55)\\ &\approx \frac{1}{2} \sum_{i=1}^{N_T} \tilde{A}_i \Delta N_i(t) + \bar{B}_i(t) & (56) \end{aligned}$$

with

$$\bar{B}_i(t) = \min_t \{\bar{B}_i^+(t), \bar{B}_i^-(t)\} \, .$$

B. Single Glitch Responses to a Ramp Signal

For simplicity consider a single threshold T_1, which means $N_T = 1$, and suppose $T_1 = 0$. Assume that the signal $x(t)$ can be approximated locally by a ramp

$$x(t) = L_X t \, , \tag{57}$$

where $0 < L_X \tau_1 < A$.

1) Symmetric Glitch Response: From (41e) and (42c) it follows that:

$$\Delta N_1(t) = \frac{1}{2A\tau_1} \begin{cases} 0 \, , & L_X t \in (-\infty, -A) \\ L_X t + A \, , & L_X t \in (-A, -A + L_X \tau_1) \\ L_X \tau_1 \, , & L_X t \in (-A + L_X \tau_1, A) \\ A + L_X \tau_1 - t \, , & L_X t \in (A, A + L_X \tau_1) \\ 0 \, , & L_X t \in (A + L_X \tau_1, \infty) \end{cases} \tag{58}$$

Using a periodic (triangle-wave) dither with density (39), Fig. 8 shows the the response of the symmetric glitch model (3) in blue (with $u(t) = x(t)$), the averaged symmetric glitch model (43) in red and the approximation (44) is indicated by the dotted line. Fig. 9 shows the equivalent for a stochastic dither signal. In this example $L_X = 1000$, $\tau_1 = 10^{-6}$ s = 0.001 ms, and $A = 1$. The output was filtered, using zero-phase filtering [64], by $g_1(t) = k_{g_1} w(t)$,

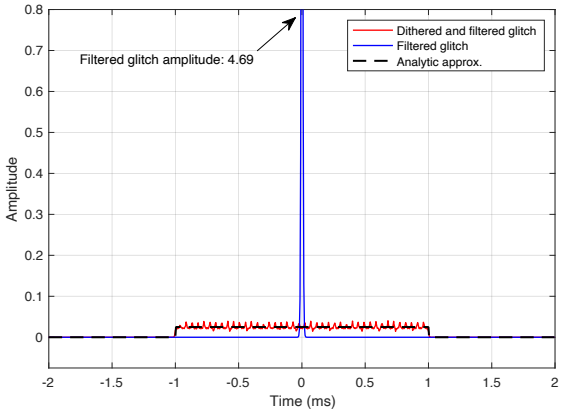

Fig. 8: Symmetric glitch response, using a periodic (triangle wave) dither. Ripple is due to finite dither frequency and imperfect averaging. Note that the plot of the filtered glitch response (no dither) has be clipped in order to obtain a reasonable scale for the dithered response.

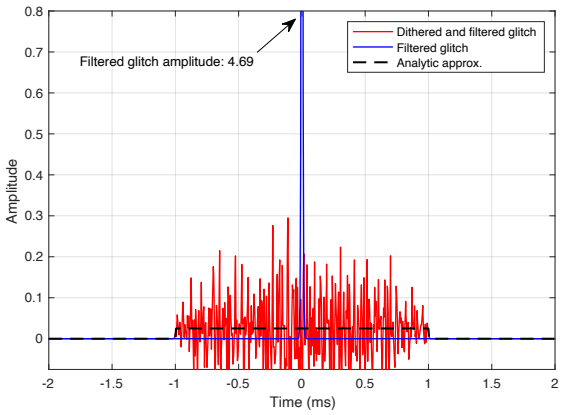

Fig. 9: Symmetric glitch response, using stochastic dither. The variation of the smoothed (dithered and filtered) glitch is due to imperfect filtering on the output. Note that the plot of the filtered glitch response (no dither) has be clipped in order to obtain a reasonable scale for the dithered response.

with $k_{g_1} = 5 \cdot 10^{-5}$ and $w(t) = \mathcal{L}^{-1}[W](t)$, with $W(s)$ given by (15). It can be seen that the glitch is converted from a large amplitude disturbance with a short duration, to a smaller amplitude disturbance with a longer duration, and the maximum slope and value of the averaged glitch response is affected by A, which can be considered a design parameter. For periodic dither the period ρ will affect the result, reducing ripple as the period is reduced, but is limited by the considerations mentioned in Sec. III-B. The time-delay τ_1 and Lipschitz constant L_X are set by the device and specific application. Given (25) and (36), increasing A does not necessarily increase the error variance or ripple in the output of the glitch model. However, the glitch model is linearly combined with the quantised output of the DAC, and the dither will be present as an unwanted disturbance in this signal. Hence, the reconstruction filter should be able to provide sufficient attenuation of the dither signal, depending on the application.

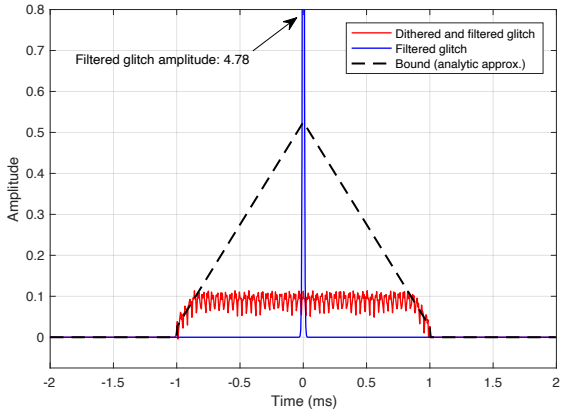

Fig. 10: Asymmetric glitch response, using a periodic (triangle-wave) dither. The maximum amplitude of the smoothed (dithered and filtered) glitch depends on parameters such as dither frequency or input signal amplitude but the expected value will not exceed the bound. Note that the plot of the filtered glitch response (no dither) has be clipped in order to obtain a reasonable scale for the dithered response.

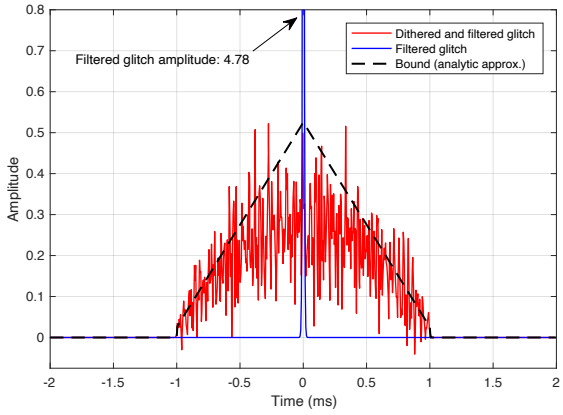

Fig. 11: Asymmetric glitch response, using a stochastic dither. The variation of the smoothed (dithered and filtered) glitch is due to imperfect filtering on the output. A lower cut-off frequency reduces the variance and brings the signal towards the bound. The bound is for the mean value. Note that the plot of the filtered glitch response (no dither) has be clipped in order to obtain a reasonable scale for the dithered response.

2) Asymmetric Glitch Response Bound: In this case the bound $\bar{B}_1(t)$ is:

$$\bar{B}_1(t) = \frac{1}{2A\tau_1} \begin{cases} 0, & L_X t \in (-\infty, -A) \\ L_X t + A, & L_X t \in (-A, -A + L_X\tau_1) \\ 2L_X t + 2A - L_X\tau_1, & L_X t \in (-A + L_X\tau_1, \frac{L_X\tau_1}{2}) \\ -2L_X t + 2A + L_X\tau_1, & L_X t \in (\frac{L_X\tau_1}{2}, A) \\ -L_X t + A + L_X\tau_1, & L_X t \in (A, A + L_X\tau_1) \\ 0, & L_X t \in (A + L_X\tau_1, \infty) \end{cases} \tag{59}$$

The analytic approximate bound (55), as well as a simulated response is shown in Fig. 10 for a periodic (triangle-wave) dither signal, and in Fig. 11 for a stochastic dither signal.

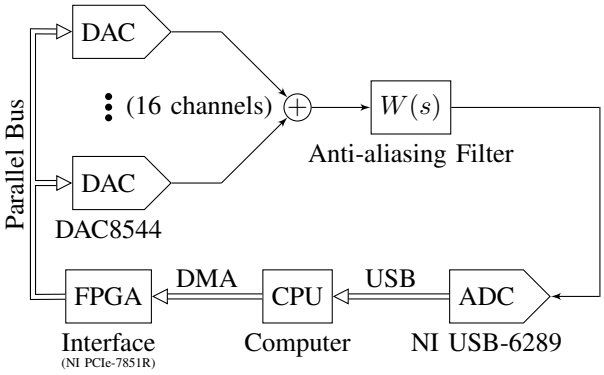

Fig. 12: Experimental set-up.

Again, $L_X = 1000$, $\tau_1 = 0.001$ ms, and $A = 1$. The output was filtered, using zero-phase filtering [64], by

$$\tilde{g}_1 = \frac{1}{2}(k_{g_1}^{+} + k_{g_1}^{-})w(t)\,, \tag{60}$$

with $k_{g_1}^{+} = -5.0 \cdot 10^{-5}$, $k_{g_1}^{-} = 5.1 \cdot 10^{-5}$, and $w(t) = \mathcal{L}^{-1}[W](t)$, where $W(s)$ is as given by (15). Again the glitch is converted from a large amplitude disturbance with a short duration, to a smaller amplitude disturbance with a longer duration. However, due to the asymmetry, the averaged glitch response is larger, but the maximum slope and value of the averaged glitch response is affected by A. Increasing A has the same effect as in the symmetric case in terms of error variance or ripple of the glitch model, but the quantised output of the DAC will also contain the dither signal, necessitating attenuation by the reconstruction filter.

V. Experimental Set-up

The experimental set-up is shown in Fig. 12. The DAC is a Texas Instruments DAC8544, as noted in Sec. II-C. The DAC features a parallel interface with low latency and provides sufficient sampling rate to accommodate dither signals with frequencies well above the mechanical bandwidth of common platforms used for high-precision motion control. The DAC exhibits large transition glitches making it ideal for experimentally validating the model and its response to dithering. Sixteen DAC-channels were used, reducing the stochastic noise-floor to emphasise the deterministic and systematic error introduced by the glitches.

Note that averaging of several channels is not necessary for the method to work. The variance of any deterministic signal is zero at any time instant, and will not be affected by the averaging of several channels; the same signal is reproduced for each channel, summed and scaled, recovering the same signal. Noise, however, is independent for each channel, and summing and scaling will cause the overall variance at any time instant to be reduced by a factor $1/N_C$, where N_C is the number of channels.

In order to minimise the dither in the output of the summing stage, the inverse of the dither signal was applied to half of the DACs, causing the dithers to approximately cancel each other out. When using dither signals with opposite polarity, large voltage differences can occur in the summing stage, potentially drawing higher currents than the DAC is designed for. Hence, the DAC outputs were buffered, using the Texas Instruments LME49990 operational amplifier, which has sufficient bandwidth and linearity.

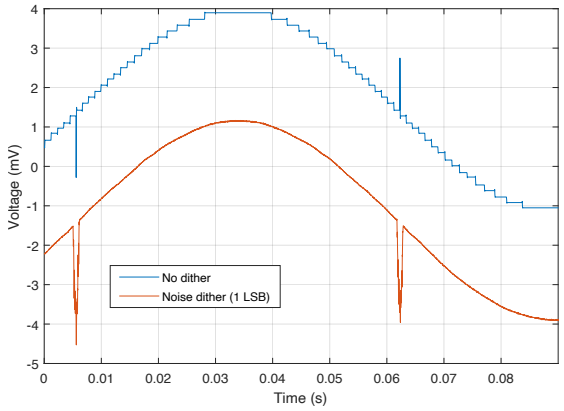

(a) Simulation result (stochastic). Plots offset for clarity.

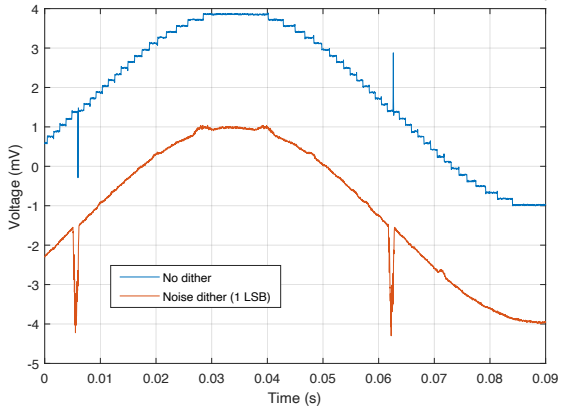

(b) Experimental result (stochastic). Plots offset for clarity.

Fig. 13: Uniform noise with 1 LSB peak-to-peak amplitude.

The DACs were connected via a parallel bus to an FPGA onboard a National Instruments PCIe-7851R interface card. The DAC driver was implemented in VHDL. LabVIEW, running on the computer (CPU), was used to generate and stream signals to the FPGA via direct memory access (DMA) over the peripheral component interconnect express (PCIe) bus. Using the sixteen DACs simultaneously, a sampling rate of 1 MS/s was achieved.

The output was measured using a National Instruments USB-6289, which contains an Analog Devices AD7674 18-bit successive approximation analogue-to-digital converter (ADC). This ADC has sufficient linearity and sampling rate, as well as higher resolution than the DAC, ensuring that it is possible to measure the response of the DAC without additional distortion. A sampling rate of 625 kS/s was used, which is the maximum supported by the interface to USB-6289. As described in Sec. II-C, the USB-6289 contains two first-order passive low-pass filters (15), with $f_c = 62.5$ kHz.

VI. Simulation and Experimental Results

A set of simulations and experiments were conducted to validate the model and evaluate the glitch mitigation method. Since the achievable sampling rate in the experimental set-up

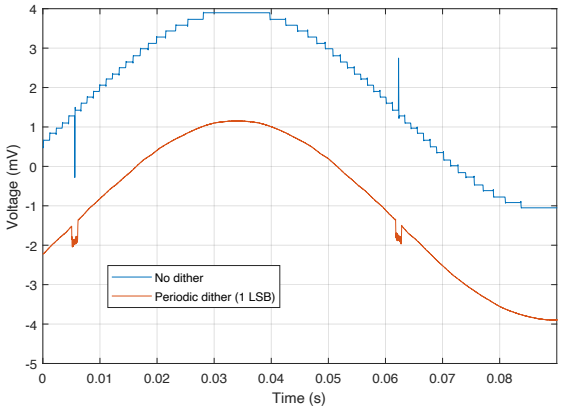

(a) Simulation result (periodic). Plots offset for clarity.

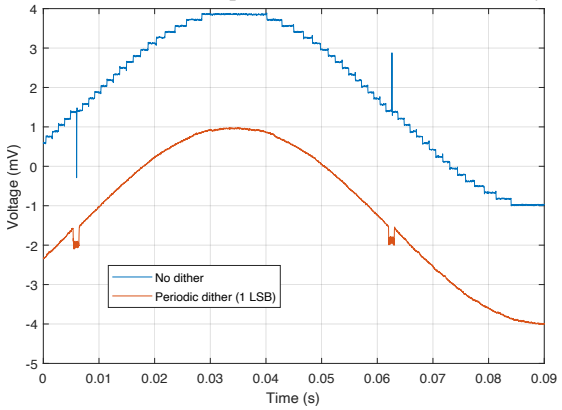

(b) Experimental result (periodic). Plots offset for clarity.

Fig. 14: Triangle-wave at 69 kHz, 1 LSB peak-to-peak.

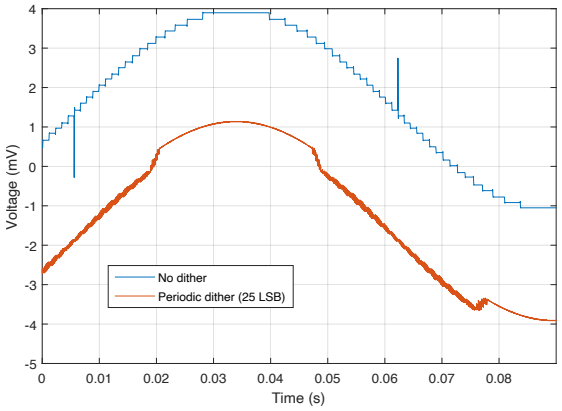

(a) Simulation result (periodic). Plots offset for clarity.

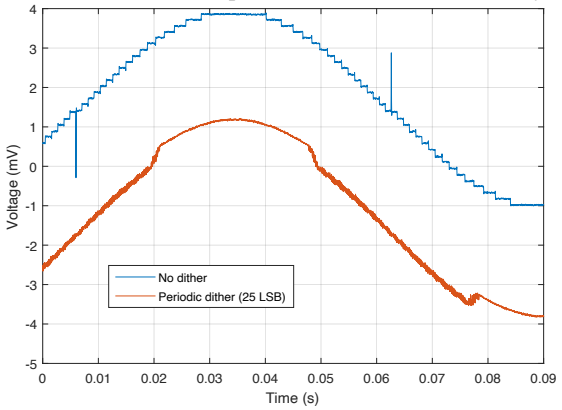

(b) Experimental result (periodic). Plots offset for clarity.

Fig. 15: Triangle-wave at 69 kHz, 25 LSB peak-to-peak.

is 1 MS/s, the glitch duration in the model was set to $\tau = 1$ μs. Furthermore, the model was set up using the glitch areas found in Sec. II-C; that is, the area for the major glitches was set to $\hat{A}_i^+ = -60.6$ nVs for the rising input, and $\hat{A}_i^- = 51.9$ nVs for the falling input, at each major glitch i. The area of the minor glitches was set to ± 2.40 nVs.

The reference signal $x(t)$ was set to be a 9 Hz sine-wave with an amplitude of 0.05% of the full range of the DAC. The stochastic dither signal was produced using a pseudo-random number generator with uniformly distributed samples. The periodic dither signal was set to be a 49 kHz triangle-wave (which has a uniform amplitude distribution).

The ADC input was filtered by the anti-aliasing filter (15), and a second-order low-pass Butterworth filter with cut-off frequency at 10 kHz was used to filter the measured response in order to achieve the desired averaging of the glitches.

Five pairs of configurations were trialled. The first pair was using a noise dither with a 1 least significant bit (LSB) range, that is, the dither was in the range $[-0.5, 0.5]$ LSB, or approximately 0.00153% of the full range of the DAC. The simulation result is shown in Fig. 13a and the experimental result is shown in Fig. 13b. The second pair was using a periodic dither with a 1 LSB peak-to-peak amplitude. The simulation and experimental results are shown in Fig. 14. The third was using a periodic dither with a 25 LSB (0.0381%) peak-to-peak amplitude, with results shown in Fig. 15. The fouth was using a noise dither with a 100 LSB (0.153%) range, with results shown in Fig. 16. The fifth was using a periodic dither with a 100 LSB peak-to-peak amplitude, with results shown in Fig. 17.

VII. Discussion

The main benefit from the averaging effect due to dithering and filtering, is that a glitch can be converted from a large amplitude disturbance with a short duration, to a smaller amplitude disturbance with a longer duration. This is a major benefit in closed-loop control applications, as a control law, which will always have limited bandwidth [65], can attenuate these low-frequency input disturbances but not the high-frequency impulse-like input disturbances.

For symmetric glitches it is apparent from the example summarised in (58) that glitches of equal but opposing amplitudes tend to cancel each other out. The effective glitch response after dithering is small. The actual shape depends not only on the density function of the dither signal but also the input signal $x(t)$. Increasing dither amplitude reduces the effective glitch amplitude and increases the duration of the glitch, whereas an increasing rate of change (Lipschitz constant) for the input signal has the opposite effect. If the input signal is bandwidth limited, there exists a maximum rate of change that can be used as a worst-case, and the amplitude can be chosen accordingly to provide sufficient averaging of the glitches.

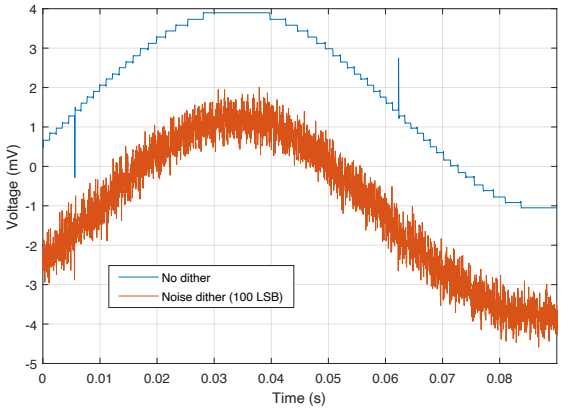

(a) Simulation result (stochastic). Plots offset for clarity.

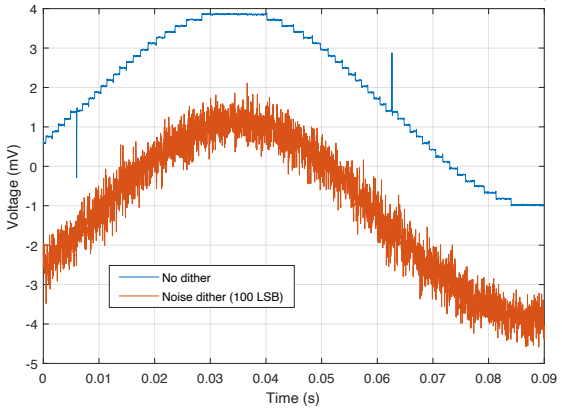

(b) Experimental result (stochastic). Plots offset for clarity.

Fig. 16: Uniform noise with 100 LSB peak-to-peak amplitude.

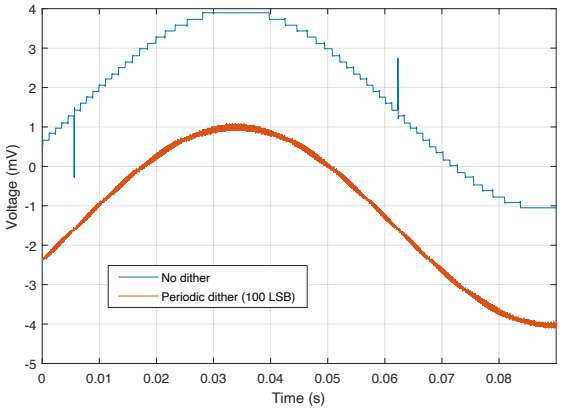

(a) Simulation result (periodic). Plots offset for clarity.

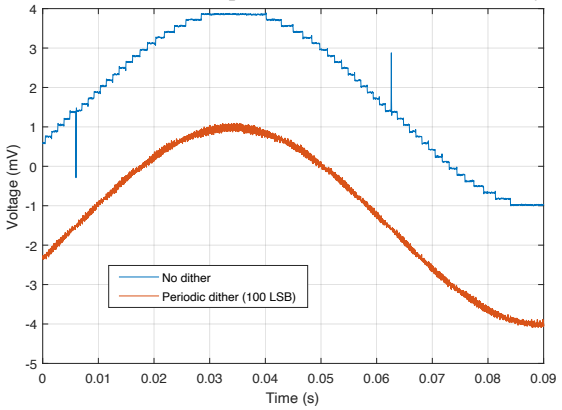

(b) Experimental result (periodic). Plots offset for clarity.

Fig. 17: Triangle-wave at 69 kHz, 100 LSB peak-to-peak.

For asymmetric glitches, the effect of dither amplitude and input signal rate-of-change is the same, but the glitches do not cancel each other out, and a larger excursion for the effective glitch is therefore to be expected, as illustrated by the expression (59). This is a bound for the expected value of the response, and the exact response will depend on dither frequency (for periodic dither signals) and the shape of the input signal. Note that increasing the amplitude of the dither does not necessarily increase the error variance or ripple in the output of the glitch model, due to the results in (25) and (36). However, the glitch model is linearly combined with the quantised output of the DAC, and the dither will be present as an unwanted disturbance in this signal.

Hence, as the dither amplitude becomes larger, it becomes more difficult to attenuate at the output. In Fig. 17 some residual ripple can be seen, when using a periodic dither. Due to the broadband nature of white noise, a significant amount of dither power remains below the cut-off frequency when using a stochastic dither, as can be seen in Fig. 16. Since it is easier to remove the narrow-band ripple, rather than the broadband noise, a periodic dither can be considered to be more practical. However, the dither frequency is limited by the considerations mentioned in Sec. III-B.

The fundamental limitations for the dither signals are set by the word size and sampling rate of the DAC: A smaller word size limits the ability to apply a large dither signal, as the usable range of the DAC is reduced. The sampling rate limits the maximum frequency of the dither signal. For a stochastic dither, a high sampling rate means that more dither power is generated above the cut-off frequency of the reconstruction filter, which reduces overall output error variance. Similarly, for a periodic dither, a high sampling rate accommodates a higher fundamental frequency, reducing residual ripple.

Considering the model validation results in Fig. 5 and 6, as well as the simulation and experimental results in Figs. 13, 14, 15, 16 and 17, it is apparent that the model provides a very good description of the glitches experienced in the DAC used in the experimental set-up. Using a small amplitude dither, as in Figs. 13 and 14, the minor glitches for each step are attenuated but the dither does not provide suppression of the major glitches. Note that the dither also linearises the quantisation steps, as expected [11]. A small amplitude dither is common in audio applications to mitigate quantisation error [66] but can be seen to aggravate the error due to glitches. Increasing the dither amplitude reduces the effect of the glitches significantly, as observed in Fig. 15. In this case the glitch has been effectively converted to a low-frequency disturbance. Further increasing the amplitude causes further suppression of the glitches, as observed in Figs. 16 and 17.

VIII. Conclusions

A new glitch model was demonstrated to provide a good behavioural description of glitches produced by digital-to-analogue converters (DACs). It was further demonstrated both analytically and experimentally that dithering and low-pass filtering can be used to suppress glitches if the dither has sufficiently large amplitude. The improvement is due to the averaging effect on the non-linearity of the DAC. It is straightforward to retrofit the method to an existing, sufficiently fast DAC, as it requires only generation of a dither signal and a suitable analogue filter for attenuation of the dither in the output. The effect of dithering can be viewed as converting glitches from short-duration, high-amplitude disturbances to long-duration, low-amplitude disturbances, which may be further attenuated by an external feedback control law.

References

[1] D. M. Freeman, "Slewing Distortion in Digital-to-Analog Conversion," *J. Audio Eng. Soc.*, vol. 25, no. 4, pp. 178–183, apr 1977.

[2] J. A. Connelly and K. P. Taylor, "An Analysis Methodology to Identify Dominant Noise Sources in D/A and A/D Converters," *IEEE Trans. Circuits Syst.*, vol. 38, no. 10, pp. 1133–1144, 1991.

[3] P. Hendriks, "Specifying Communications DACs," *IEEE Spectrum*, vol. 34, no. 7, pp. 58–69, jul 1997.

[4] R. J. van de Plassche, *CMOS Integrated Analog-to-Digital and Digital-to-Analog Converters*. Kluwer Academic Publishers, 2003.

[5] I. Galton, "Why Dynamic-Element-Matching DACs Work," *IEEE Trans. Circuits Syst. II*, vol. 57, no. 2, pp. 69–74, 2010.

[6] M. J. Pelgrom, *Analog-to-Digital Conversion*, 2nd ed. Springer-Verlag, 2013.

[7] P. Horowitz and W. Hill, *The Art of Electronics*, 3rd ed. Cambridge University Press, 2015.

[8] B. Widrow, I. Kollar, and M.-C. Liu, "Statistical Theory of Quantization," *IEEE Trans. Instrum. Meas.*, vol. 45, no. 2, pp. 353–361, 1996.

[9] A. A. Eielsen, M. Vagia, J. T. Gravdahl, and K. Y. Pettersen, "Damping and Tracking Control Schemes for Nanopositioning," *IEEE/ASME Trans. Mechatronics*, vol. 19, no. 2, pp. 432–444, 2013.

[10] R. E. Crochiere and L. Rabiner, "Interpolation and decimation of digital signals—A tutorial review," in *Proc. IEEE*. IEEE, 1981, pp. 300–331.

[11] J. Vanderkooy and S. P. Lipshitz, "Resolution Below the Least Significant Bit in Digital Systems with Dither," *J. Audio Eng. Soc.*, vol. 32, no. 3, pp. 106–113, 1984.

[12] P. Carbone and D. Petri, "Performance of Stochastic and Deterministic Dithered Quantizers," *IEEE Trans. Instrum. Meas.*, vol. 49, no. 2, 2000.

[13] R. A. Wannamaker, S. Lipshitz, J. Vanderkooy, and J. N. Wright, "A Theory of Nonsubtractive Dither," *IEEE Trans. Signal Process.*, vol. 48, no. 2, pp. 499–516, 2000.

[14] R. Skartlien and L. Øyehaug, "Quantization Error and Resolution in Ensemble Averaged Data With Noise," *IEEE Trans. Instrum. Meas.*, vol. 54, no. 3, pp. 1303–1312, 2005.

[15] A. A. Eielsen and A. J. Fleming, "Existing methods for improving the accuracy of digital-to-analog converters," *Rev. Sci. Instrum*, vol. 88, no. 9, p. 094702, 2017.

[16] J. Remple and I. Galton, "The Effects of Inter-Symbol Interference in Dynamic Element Matching DACs," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 64, no. 1, pp. 14–23, 2017.

[17] A. A. Eielsen and A. J. Fleming, "Improving Digital-to-Analog Converter Linearity by Large High-Frequency Dithering," *IEEE Trans. Circuits Syst. I*, vol. 64, no. 6, pp. 1409–1420, 2017.

[18] B. Francis and W. Wonham, "The internal model principle of control theory," *Automatica*, vol. 12, no. 5, pp. 457–465, Sep. 1976.

[19] D. Mercer, "A 16-b D/A converter with increased spurious free dynamic range," *IEEE Journal of Solid-State Circuits*, vol. 29, no. 10, pp. 1180–1185, 1994.

[20] A. Bugeja, B.-S. Song, P. Rakers, and S. Gillig, "A 14-b, 100-MS/s CMOS DAC designed for spectral performance," *IEEE Journal of Solid-State Circuits*, vol. 34, no. 12, pp. 1719–1732, Dec. 1999.

[21] M. K. Rudberg, M. Vesterbacka, N. Andersson, and J. J. Wikner, "Glitch minimization and dynamic element matching in D/A converters," in *Electronics, Circuits and Systems, 2000. ICECS 2000. The 7th IEEE International Conference on*, vol. 2. IEEE, 2000, pp. 899–902.

[22] M. Vesterbacka, "Linear-coded D/A converters with small relative error due to glitches," in *Proceedings of the 44th IEEE 2001 Midwest Symposium on Circuits and Systems. MWSCAS 2001 (Cat. No.01CH37257)*, vol. 1. Dayton, OH, USA: IEEE, 2001, pp. 280–283.

[23] B. J. Tesch and J. C. Garcia, "A Low Glitch 14-b 100-MHz D/A Converter," *IEEE Journal of Solid-State Circuits*, vol. 32, no. 9, pp. 1465–1469, 1997.

[24] E. Olieman, A.-J. Annema, and B. Nauta, "An Interleaved Full Nyquist High-Speed DAC Technique," *IEEE Journal of Solid-State Circuits*, vol. 50, no. 3, pp. 704–713, 2015.

[25] J. Leijten, J. Van Meerbergen, and J. Jess, "Analysis and reduction of glitches in synchronous networks," in *Proc. Design and Test Conference*. IEEE Computer Society, mar 1995, pp. 398–403.

[26] R. Adams and K. Q. Nguyen, "A 113-dB SNR oversampling DAC with segmented noise-shaped scrambling," *IEEE J. Solid-State Circuits*, vol. 33, no. 12, pp. 1871–1878, 1998.

[27] L. Luh, J. Choma, and J. Draper, "A high-speed fully differential current switch," *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 47, no. 4, pp. 358–363, Apr. 2000.

[28] O. Coudert, "Gate sizing for constrained delay/power/area optimization," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 5, no. 4, pp. 465–472, Dec. 1997.

[29] T. Shui, R. Schreier, and F. Hudson, "Mismatch Shaping for a Current-Mode Multibit DeltaSigma DAC," *IEEE Journal of Solid-State Circuits*, vol. 34, no. 3, pp. 331–338, 1999.

[30] L. Risbo, R. Hezar, B. Kelleci, H. Kiper, and M. Fares, "Digital Approaches to ISI-Mitigation in High-Resolution Oversampled Multi-Level D/A Converters," *IEEE J. Solid-State Circuits*, vol. 46, no. 12, p. 2892, 2011.

[31] V. OBrien, A. G. Scanlan, and B. Mullane, "A Reduced Hardware ISI and Mismatch Shaping DEM Decoder," *Circuits, Systems, and Signal Processing*, vol. 37, no. 6, pp. 2299–2317, 2018.

[32] M.-H. Shen, J.-H. Tsai, and P.-C. Huang, "Random Swapping Dynamic Element Matching Technique for Glitch Energy Minimization in Current-Steering DAC," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 57, no. 5, pp. 369–373, May 2010.

[33] K. Andersson and M. Vesterbacka, "Modeling of Glitches Due to Rise/fall Asymmetry in Current-steering Digital-to-analog Converters," *IEEE Trans. Circuits Syst. I*, vol. 52, no. 11, pp. 2265–2275, 2005.

[34] K. L. Chan, N. Rakuljic, and I. Galton, "Segmented Dynamic Element Matching for High-Resolution Digital-to-Analog Conversion," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 55, no. 11, pp. 3383–3392, 2008.

[35] D. Rabe and W. Nebel, "New approach in gate-level glitch modelling," in *Proc. Design Automation Conference*. Geneva, Switzerland: IEEE Computer Society, 1996, pp. 66–71.

[36] J. J. Wikner and N. Tan, "Modeling of CMOS digital-to-analog converters for telecommunication," *IEEE Trans. Circuits Syst. II*, vol. 46, no. 5, pp. 489–499, 1999.

[37] J. Vandenbussche, G. Van Der Plas, G. Gielen, and W. Sansen, "Behavioral model of reusable D/A converters," *IEEE Trans. Circuits Syst. II*, vol. 46, no. 10, pp. 1323–1326, 1999.

[38] S. Warecki, "Behavioral simulation of digital to analog converters simulation of segmented current steering DAC with utilization of perfect sampling technique," Ph.D. dissertation, The University of Arizona, 2002.

[39] L. R. A. MacColl, *Fundamental Theory of Servomechanisms*. Van Nostrand, 1945.

[40] R. Oldenburger and R. C. Boyer, "Effects of Extra Sinusoidal Inputs to Nonlinear Systems," *J Basic Eng-T ASME*, vol. 84, no. 4, pp. 559–569, 1962.

[41] G. Zames and N. A. Shneydor, "Dither in Nonlinear Systems," *IEEE Trans. Autom. Control*, vol. 21, no. 5, pp. 660–667, 1976.

[42] R. J. Simpson and H. M. Power, "Applications of high frequency signal injection in non-linear systems," *International Journal of Control*, vol. 26, no. 6, pp. 917–943, 1977.

[43] S. Mossaheb, "Application of a method of averaging to the study of dithers in non-linear systems," *Int J Control*, vol. 38, no. 3, pp. 557–576, 1983.

[44] C. A. Desoer and S. M. Shahruz, "Stability of dithered non-linear systems with backlash or hysteresis," *International Journal of Control*, vol. 43, no. 4, pp. 1045–1060, 1986.

[45] L. Iannelli, K. H. Johansson, U. T. Jönsson, and F. Vasca, "Averaging of nonsmooth systems using dither," *Automatica*, vol. 42, no. 4, pp. 669–676, 2006.

[46] I. De Lotto and G. E. Paglia, "Dithering improves A/D converter linearity," *IEEE Trans. Instrum. Meas.*, vol. IM-35, no. 2, pp. 170–177, 1986.

[47] B. A. Blesser and B. N. Locanthi, "The Application of Narrow-Band Dither Operating at the Nyquist Frequency in Digital Systems to Provide Improved Signal-to-Noise Ratio over Conventional Dithering," *J. Audio Eng. Soc.*, vol. 35, no. 6, pp. 446–454, 1987.

[48] S. A. Leyonhjelm, M. Faulkner, and P. Nilsson, "An Efficient Implementation of Bandlimited Dithering," *Wireless Pers Commun*, vol. 8, no. 1, pp. 31–36, 1998.

[49] F. Adamo, F. Attivissimo, N. Giaquinto, and A. Trotta, "A/D converters nonlinearity measurement and correction by frequency analysis and dither," *IEEE Transactions on Instrumentation and Measurement*, vol. 52, no. 4, pp. 1200–1205, 2003.

[50] J. Sanjuán, A. Lobo, and J. Ramos-Castro, "Analog-to-digital converters nonlinear errors correction in thermal diagnostics for the laser interferometer space antenna mission," *Review of Scientific Instruments*, vol. 80, no. 11, p. 114901, 2009.

[51] P. Carbone, C. Narduzzi, and D. Petri, "Dither Signal Effects on the Resolution of Nonlinear Quantizers," *IEEE Trans. Instrum. Meas.*, vol. 43, no. 2, pp. 139–145, 1994.

[52] M. J. Flanagan and G. A. Zimmerman, "Spur-Reduced Digital Sinusoid Synthesis," *IEEE Transactions on Communications*, vol. 43, no. 7, pp. 2254–2262, Jul. 1995.

[53] J. Vankka, "Spur reduction techniques in sine output direct digital synthesis," in *Proceedings of 1996 IEEE International Frequency Control Symposium*. Honolulu, HI, USA: IEEE, 1996, pp. 951–959.

[54] L. Cordesses, "Direct digital synthesis: a tool for periodic wave generation (part 2)," *IEEE Signal Processing Magazine*, vol. 21, no. 5, pp. 110–117, Sep. 2004.

[55] E. Temporiti, C. Weltin-Wu, D. Baldi, M. Cusmai, and F. Svelto, "A 3.5 GHz Wideband ADPLL With Fractional Spur Suppression Through TDC Dithering and Feedforward Compensation," *IEEE Journal of Solid-State Circuits*, vol. 45, no. 12, pp. 2723–2736, Dec. 2010.

[56] C.-R. Ho and M. S.-W. Chen, "A Digital PLL With Feedforward Multi-Tone Spur Cancellation Scheme Achieving <73 dBc Fractional Spur and <110 dBc Reference Spur in 65 nm CMOS," *IEEE Journal of Solid-State Circuits*, vol. 51, no. 12, pp. 3216–3230, Dec. 2016.

[57] V. R. Gonzalez-Diaz, M. A. Garcia-Andrade, G. E. Flores-Verdad, and F. Maloberti, "Efficient Dithering in MASH Sigma-Delta Modulators for Fractional Frequency Synthesizers," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 57, no. 9, pp. 2394–2403, Sep. 2010.

[58] L. Rabiner, "Techniques for Designing Finite-Duration Impulse-Response Digital Filters," *IEEE Transactions on Communications*, vol. 19, no. 2, pp. 188–195, 1971.

[59] J. Schoukens and R. Pintelon, *Identification of Linear Systems*. Pergamon Press, 1991.

[60] A. Papoulis and S. U. Pillai, *Probability, random variables, and stochastic processes*, 4th ed. McGraw-Hill, 2002.

[61] G. Kallianpur, *Stochastic Filtering Theory*. Springer-Verlag, 1980.

[62] S. Boyd and C. Barratt, *Linear Controller Design: Limits of Performance*. Prentice-Hall, 1991.

[63] R. L. Wheeden and A. Zygmund, *Measure and Integral*. Marcel Dekker, Inc., 1977.

[64] F. Gustafsson, "Determining the Initial States in Forward-backward Filtering," *IEEE Trans. Signal Process.*, vol. 44, no. 4, pp. 988–992, 1996.

[65] G. C. Goodwin, S. F. Graebe, and M. E. Salgado, *Control System Design*. Prentice Hall, Inc., 2000.

[66] S. P. Lipshitz, R. A. Wannamaker, and J. Vanderkooy, "Quantization and Dither: A Theoretical Survey," *J. Audio Eng. Soc.*, vol. 40, no. 5, pp. 355–375, 1992.

Arnfinn A. Eielsen was born in 1980 in Stavanger, Norway. He received the MSc degree and the PhD degree in Engineering Cybernetics from the Norwegian University of Science and Technology (NTNU), Trondheim in 2007 and 2012, respectively. He is presently an Associate professor in the Department of Electrical Engineering and Computer Science at the University of Stavanger, Norway. He has previously been employed in the School of Electrical Engineering and Computer Science, The University of Newcastle, Australia.

John J. Leth is an Associate Professor at the Department of Electronic Systems/Automation & Control at the Aalborg University. Research interest are (stochastic) hybrid/switched dynamical systems, as well as mathematical control theory, such as optimal control theory for nonlinear systems, realisation theory for nonlinear systems, and infinite dimensional system theory.

Andrew J. Fleming graduated from The University of Newcastle, Australia with a BE (Elec.) in 2000 and PhD in 2004. He is presently a Professor in the School of Electrical Engineering and Computer Science, The University of Newcastle, Australia. His research includes nanopositioning, high-speed scanning probe microscopy, micro-cantilever sensors, metrological position sensors, and optical nanofabrication. Academic awards include the University of Newcastle Vice-Chancellors Award for Researcher of the Year and the IEEE Control Systems Society Outstanding Paper Award for research published in the IEEE Transactions on Control Systems Technology. He is the co-author of three books, several patent applications and more than 100 Journal and Conference papers.

Adrian G. Wills earned his BE (Elec.) and PhD degrees from The University of Newcastle, Australia in May 1999 and May 2003, respectively. Since then, he has held Postdoctoral research positions at Newcastle and spent three years working in industry. The focus of his research has been in the areas of system identification and model predictive control. In July 2015, he returned to The University of Newcastle to lead the Mechatronics engineering program.

Brett Ninness was born in 1963 in Singleton, Australia and received his BE, ME and PhD degrees in Electrical Engineering from The University of Newcastle, Australia in 1986, 1991 and 1994 respectively. He has stayed with the School of Electrical Engineering and Computer Science at The University of Newcastle since 1993, where he is currently a Professor. His research interests are in system identification and stochastic signal processing, in which he has authored approximately one hundred papers in journals and conference proceedings.