The Impact of SAR-ADC Mismatch on Quantized Massive MU-MIMO Systems

Jérémy Guichemerre and Christoph Studer

Dept. of Information Technology and Electrical Engineering, ETH Zurich, Switzerland
email: jeremyg@iis.ee.ethz.ch and studer@ethz.ch

***Abstract*—Low-resolution analog-to-digital converters (ADCs) in massive multi-user (MU) multiple-input multiple-output (MIMO) wireless systems can significantly reduce the power, cost, and interconnect data rates of infrastructure basestations. Thus, recent research on the theory and algorithm sides has extensively focused on such architectures, but with idealistic quantization models. However, real-world ADCs do *not* behave like ideal quantizers, and are affected by fabrication mismatches. We analyze the impact of capacitor-array mismatches in successive approximation register (SAR) ADCs, which are widely used in wireless systems. We use Bussgang's decomposition to model the effects of such mismatches, and we analyze their impact on the performance of a single ADC. We then simulate a massive MU-MIMO system to demonstrate that capacitor mismatches should *not* be ignored, even in basestations that use low-resolution SAR ADCs.**

I. Introduction

Millimeter-wave frequencies provide access to large portions of bandwidth and are therefore targeted for fifth-generation (5G) and beyond-5G wireless systems [1], [2]. Combined with multi-user (MU) massive multiple-input multiple-output (MIMO), one cannot only achieve high-data-rate communication with multiple user equipments (UEs) in the same frequency band, but also combat the strong path loss at such carrier frequencies by means of beamforming [3], [4]. However, the large number of antenna elements and radio-frequency chains in massive MU-MIMO basestation architectures combined with large bandwidths poses significant practical implementation challenges in terms of system costs and power consumption.

Quantized massive MU-MIMO [5]–[7] has become a popular means to reduce the costs and power consumption in the analog front-ends (FEs). In fact, decreasing the analog-to-digital converter (ADC) resolution relaxes the design requirements of the entire analog FE, thereby reducing power consumption and silicon area for both analog FEs and digital processing. Thus, research has mainly focused on understanding the effects of coarse quantization and on the design of corresponding baseband algorithms that mitigate quantization artifacts. However, the impact of quantization non-idealities that arise in real-world ADCs has been routinely ignored in the literature.

The work of JG and CS was supported in part by an ETH Research Grant. The work of CS was supported in part by the U.S. National Science Foundation (NSF) under grants CNS-1717559 and ECCS-1824379.

The authors would like to thank Dr. Thomas Burger for his support, Gian Marti for comments on an initial version of the paper, as well as Seyed Hadi Mirfarshbafan and Darja Nonaca for their help with the system simulator.

The successive approximation register (SAR) ADC [8] is the data converter of choice for many communications applications supporting bandwidths of tens of megahertz to low gigahertz [9], [10]. The popularity of SAR ADCs is due to several reasons: First, they require low energy per conversion step. Second, they scale well to smaller semiconductor technology nodes. Third, they are suitable for pipelined and interleaved architectures, which enable one to achieve even higher sampling rates [11]–[13]. However, as all real-world data converters, SAR ADCs are *not* ideal quantizers and are affected by imperfections, which already occur during manufacturing. Specifically, the capacitors used for data conversion suffer from mismatches, i.e., the capacitor values of the fabricated circuit deviate from the specified ideal ones, which results in distortions of the quantization function. The impact of such capacitor mismatches on wireless communication systems is, until now, unexplored territory.

A. Contributions

In this paper, we study the impacts of capacitor mismatches in SAR ADCs on massive MU-MIMO wireless systems. To this end, we first propose the effective resolution (EFR), a new figure of merit (FoM) that better reflects the inherent trade-off between quantization noise and clipping artifacts than traditional metrics, such as the effective number of bit (ENOB) [10]. We then use Bussgang's decomposition [14] to model most-significant-bit (MSB) capacitor mismatches, and we analyze the impacts of quantization, clipping, and capacitor mismatch on the EFR of a single SAR ADC. Finally, we present simulation results for a quantized massive MU-MIMO wireless system in order to demonstrate that capacitor mismatches in SAR ADCs can have a significant performance impact and should therefore *not* be ignored.

B. Notation

Matrices are written in bold uppercase, column vectors in bold lowercase, and sets in calligraphic letters; the Frobenius norm of a matrix $\mathbf{A}$ is $\|\mathbf{A}\|_{\mathrm{F}}$. Expectation is denoted by $\mathbb{E}[\cdot]$. We use $\phi_{\sigma_X}(\cdot)$ to refer to the probability density function (PDF) of a zero-mean Gaussian random variable $X \sim \mathcal{N}(0, \sigma_X^2)$ with variance σ_X^2. The PDF of a standard Gaussian $\mathcal{N}(0,1)$ is $\phi(\cdot)$, and $Q(\cdot) = \int_x^\infty \phi(x)\mathrm{d}x$ is the Q-function. We write $\mathrm{u}(\cdot)$ for the unit step function, $\delta(\cdot)$ for the Dirac distribution, and $\mathrm{rect}_{[a,b]}(\cdot)$

for the rectangle function that is 1 on $[a, b]$ and 0 elsewhere. The derivative of a function f is denoted by f'.

II. Prerequisites

We start by analyzing the performance of a single, ideal ADC, where we only model quantization and clipping. To this end, we first summarize Bussgang's decomposition [14], which we use in our performance analysis. We then introduce a meaningful FoM in order to assess the performance of such an ideal ADC. The analysis of a SAR ADC with capacitor mismatch is provided in Sec. III.

A. A Primer on Bussgang's Decomposition

Bussgang's decomposition models the output of a nonlinear input-output transfer function (TF) f applied to a zero-mean Gaussian random variable X with variance σ_X^2 as a superposition of a linear term scaled with Bussgang gain β and a statistically uncorrelated distortion D as [15]

$$f(X) = \beta X + D. \tag{1}$$

Here, f will eventually model the ADC's TF. By multiplying both sides of (1) by X and evaluating its expected value, it can be shown that the Bussgang gain β is given by

$$\beta \overset{(a)}{=} \frac{\mathbb{E}[Xf(X)]}{\sigma_X^2} \overset{(b)}{=} \mathbb{E}[f'(X)]. \tag{2}$$

Here, (a) follows from the requirement of D being uncorrelated with X and (b) from Stein's lemma [16]. Note that (a) holds for any zero-mean random variable X whereas (b) requires X to be Gaussian. In order to assess the performance of a nonlinear function (e.g., the quantizer of an ADC), we are interested in characterizing the power of the distortion:[1]

$$\mathbb{E}\big[D^2\big] = \mathbb{E}\big[(f(X) - \beta X)^2\big] = \mathbb{E}\big[f^2(X)\big] - \beta^2\sigma_X^2. \tag{3}$$

Note that β from (2) also minimizes $\mathbb{E}\big[D^2\big]$. From (2) and (3), we notice that only two quantities need to be known in order to evaluate the power of the distortion D, namely $\mathbb{E}[f'(X)]$ to compute the Bussgang gain β and $\mathbb{E}\big[f^2(X)\big]$.

B. Model for an Ideal ADC

We will now derive Bussgang's decomposition for an ideal N-bit uniform symmetric mid-rise quantizer with an input range of $[-1, 1]$; any value outside this range will be clipped. This ideal ADC model will be used in Sec. II-C to introduce the FoMs and also serve as a baseline. For $x \geq 0$, we have

$$f_{\geq 0}(x) = \frac{\Delta}{2}\mathrm{u}(x) + \Delta\sum_{k=1}^{2^{N-1}-1}\mathrm{u}(x - k\Delta), \tag{4}$$

$$f'_{\geq 0}(x) = \frac{\Delta}{2}\delta(x) + \Delta\sum_{k=1}^{2^{N-1}-1}\delta(x - k\Delta), \tag{5}$$

where $\Delta = 2^{1-N}$ is a least significant bit (LSB) step width. Note that it is sufficient to consider the case $x \geq 0$ as f is an odd function; this also means that f' is an even function. Splitting equations with the sign of x is not necessary for the analysis of an ideal ADC, but will significantly simplify our derivations when considering mismatches in Sec. III.

Using (5), we can write (2) as follows:

$$\begin{aligned}\beta &= \int_{-\infty}^{+\infty} f'(x)\phi_{\sigma_X}(x)\,\mathrm{d}t \overset{(a)}{=} 2\int_0^{+\infty} f'_{\geq 0}(x)\phi_{\sigma_X}(x)\,\mathrm{d}t \\ &\overset{(b)}{=} \frac{\Delta}{\sigma_X}\left(\phi(0) + 2\sum_{k=1}^{2^{N-1}-1}\phi\Big(\frac{\Delta}{\sigma_X}k\Big)\right).\end{aligned} \tag{6}$$

Here, (a) exploits symmetries in the functions f' and ϕ_{σ_X}, and (b) follows from the fact that $\phi_{\sigma_X}(t) = \frac{1}{\sigma_X}\phi\big(\frac{t}{\sigma_X}\big)$.

We note that (6) is a partial Riemann integral of ϕ_{σ_X} on the ADC range $[-1, 1]$. When taking the limit $N \to \infty$, β corresponds to $1 - 2Q(1/\sigma_X)$, which is the Bussgang gain for an ADC with infinite resolution and clipping only [14].

With the expression for β in (6), we can now follow a similar procedure to evaluate $\mathbb{E}\big[f^2(X)\big]$. We first write

$$\begin{aligned}f^2_{\geq 0}(x) = &\sum_{k=0}^{2^{N-1}-1}\left(\frac{\Delta}{2} + k\Delta\right)^2\mathrm{rect}_{[k\Delta,(k+1)\Delta]}(x) \\ &+\left(1 - \frac{\Delta}{2}\right)^2 u(x-1),\end{aligned} \tag{7}$$

from which we obtain the following result:

$$\begin{aligned}\mathbb{E}\big[f^2(X)\big] =& 2\sum_{k=1}^{2^{N-1}-1}\left(\left(\frac{\Delta}{2} + k\Delta\right)^2\int_{k\frac{\Delta}{\sigma_X}}^{(k+1)\frac{\Delta}{\sigma_X}}\phi(t)\mathrm{d}t\right) \\ &+ 2\left(1 - \frac{\Delta}{2}\right)^2 Q\left(\frac{1}{\sigma_X}\right).\end{aligned} \tag{8}$$

Now that we have an expression for the Bussgang gain β and $\mathbb{E}\big[f^2(X)\big]$, we can insert these expressions into (3) to obtain an analytical expression for the distortion power $\mathbb{E}\big[D^2(X)\big]$.

C. Figures of Merit

In order to analyze the effects of non-idealities in an ADC, it is key to utilize a meaningful FoM. A common metric used to analyze the error introduced by a nonlinear distortion and (possibly) noise is the mean-squared error (MSE) defined as $\mathit{MSE} \triangleq \mathbb{E}\big[(f(X) - X)^2\big]$. The MSE directly compares $f(X)$ with X, which suffers from the limitation that $f(X)$ and X remain to be correlated, which is prone to overestimating the effect of the distortion caused by nonidealities.

In order to circumvent the limitations of the MSE, we consider the signal-to-distortion ratio (SDR) defined as

$$\mathit{SDR} \triangleq \frac{\mathbb{E}\big[(\beta X)^2\big]}{\mathbb{E}[(f(X) - \beta X)^2]} = \frac{\beta^2\sigma_X^2}{\mathbb{E}[D^2]}, \tag{9}$$

where the denominator is the error when the scaling effect of Bussgang's decomposition in (1) is taken into account.

Another common FoM used in the ADC literature, known as the effective number of bits (ENOB) [10], measures the effective resolution that the converter allows to reach when quantization, noise, and non-linearities are all taken into

[1] Note that the distortion power $\mathbb{E}\big[D^2\big]$ might include a bias term if $f(X)$ has nonzero mean. This bias could be removed by subtracting $\mathbb{E}[f(X)]^2$ from $\mathbb{E}\big[D^2\big]$, but is ignored throughout the paper for the sake of simplicity.

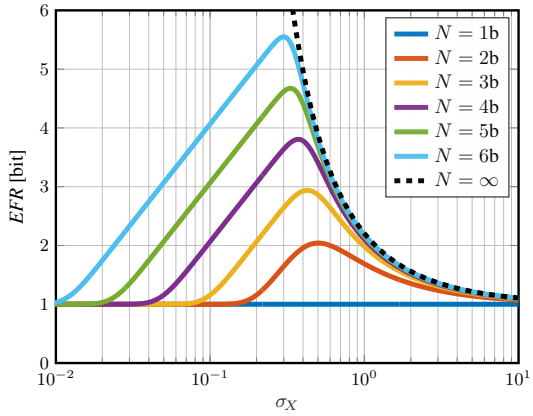

Fig. 1. Effective resolution (EFR) of an ideal N-bit ADC with quantization and clipping only, dependent on the input standard deviation σ_X.

account. The ENOB refers to the resolution an ideal noiseless quantizer needs to have in order to reach the same signal-to-noise-and-distortion ratio (SNDR) as the ADC we want to test when a full-scale sinusoid is applied at its input. Two important assumptions the ENOB makes is that the quantization error is independent of the input and uniformly distributed in $[-\Delta/2, +\Delta/2]$. While the ENOB presents some practical advantages (e.g., it can be measured with a full-scale sinusoid at the input), it is not well-suited as a performance indicator for communication systems for the following two reasons. First, the ADC input signals are typically not sinusoids but rather Gaussian-like [17]. Second, the ENOB lacks a dependency on the input amplitude (or gain), which renders it incapable of characterizing the inherent trade-off between quantization artifacts and amplitude clipping.

In order to circumvent the limitations of ENOB, we can utilize the SDR from (9) to postulate a novel FoM that is related to resolution. Concretely, we propose the effective resolution EFR (measured in bits) defined as follows:

$$EFR \triangleq \frac{SDR_{\text{dB}} + 10\log_{10}(2(\pi-2))}{20\log_{10}(2)} \approx \frac{SDR_{\text{dB}} + 3.59}{6.02} \qquad (10)$$

with $SDR_{\text{dB}} = 10\log_{10}(SDR)$. The denominator ensures that halving the amplitude of the distortion compared to the one of the wanted signal leads to one additional bit of resolution (this principle is also used in the ENOB definition). The offset in the numerator ensures that a 1-bit quantizer has an EFR of 1 b. Note that the SDR of a 1-bit quantizer does not depend on the input amplitude (determined by σ_X), which makes the correction factor meaningful when considering different input gains.

D. EFR of an Ideal ADC

In order to analyze the effect of non-idealities in ADCs, we need an ideal reference. For this purpose, we now briefly discuss the EFR for the N-bit ideal ADC with input range $[-1, +1]$, where we only consider quantization and clipping as in Sec. II-B. In Fig. 1, we see that an N-bit quantizer reaches an EFR close to N only for low resolutions. For higher resolutions, the quantization levels close to the clipping boundaries -1 and $+1$ cannot be fully exploited as an increase in input power would lead to more clipping artifacts. We also see that there exists an optimum input amplitude (determined by σ_X in Fig. 1) that trades off quantization errors versus clipping artifacts—this optimum depends on the ADC's resolution.

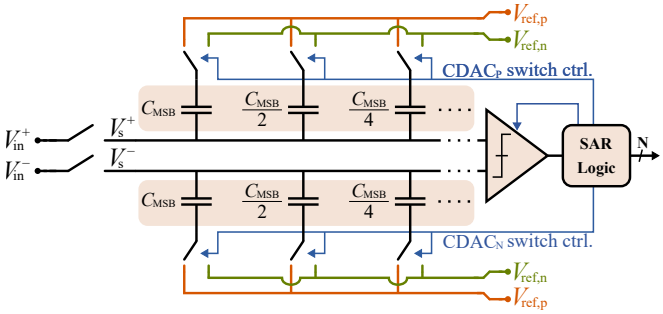

Fig. 2. A high-level block diagram of a SAR ADC.

III. SAR ADC Mismatch Model

We now briefly present the operating principle of a successive approximation register (SAR) ADC and discuss the origin of capacitor mismatches. We then propose a mathematical model to analyze the impact of such mismatches. Finally, we present numerical results for both the model and simulations.

A. Basics of SAR ADC Operation

A SAR ADC operates in two phases. In the first phase (the tracking phase), the ADC follows the input voltage and samples it at the end of this phase when the ADC sampling switches open. In the second phase (the conversion phase), the ADC performs a binary search to sequentially determine the bits associated with the sampled voltage. The way of implementing this binary search is charge-sharing, achieved by capacitor-array (referred to as CDAC for capacitive digital-to-analog converter) switching. Fig. 2 shows a high-level block diagram of a SAR ADC. Like most modern analog circuits, SAR ADCs are implemented differentially (as opposed to single-ended implementation). This means that the voltage we want to digitize is no longer the voltage on a single node, but the voltage difference between a positive and a negative node:

$$V_{\text{in,d}} = V_{\text{in}}^+ - V_{\text{in}}^- \quad \text{and} \quad V_{\text{s,d}} = V_{\text{s}}^+ - V_{\text{s}}^-. \qquad (11)$$

The input range of the ADC is $[-\Delta V_{\text{ref}}, +\Delta V_{\text{ref}}]$ with $\Delta V_{\text{ref}} = V_{\text{ref,p}} - V_{\text{ref,n}}$.

During the tracking phase, the sampling switches (on the left of Fig. 2) are closed and all capacitors are connected to $V_{\text{ref,p}}$. After the sampling switches open, the voltage is captured in the CDAC and conversion begins: $V_{\text{s,d}}$ is a sampled version of $V_{\text{in,d}}$, and the comparator can convert the most significant bit (MSB). If this first conversion yields a 1, i.e, if $V_{\text{s,d}} > 0$, then we need to subtract $\Delta V_{\text{ref}}/2$ from $V_{\text{s,d}}$ to continue with the binary search. To do so, we switch the MSB capacitor of CDAC$_{\text{P}}$ to $V_{\text{ref,n}}$ (CDAC$_{\text{P}}$ is the upper capacitive array on Fig. 2). If the MSB was 0, then we would do the same but on CDAC$_{\text{N}}$ (lower capacitive array on Fig. 2) to add $\Delta V_{\text{ref}}/2$ to $V_{\text{s,d}}$. One can then perform the next comparison, and conditionally add or subtract $\Delta V_{\text{ref}}/4$ by switching $C_{\text{MSB}}/2$ to $V_{\text{ref,n}}$ on the corresponding

CDAC (depending on the comparison result). This process repeats until the required number of bits has been converted. Note that the discussed switching scheme is common in high-speed designs as it only requires one reference voltage ($V_{\text{ref,n}}$ is typically connected to ground), but other possibilities exist.[2]

B. Mismatches in SAR ADCs

When a SAR ADC is manufactured, the capacitors in the CDACs do not have the exact (ideal) value they have been designed for due to process variation [18].[3] Even if such mismatches are a random process, there will be no averaging effect: once a chip has been manufactured, its mismatches are fixed and will affect the system in a deterministic way.

It is common to model the relative variation $\Delta\text{C}/\text{C}$ of a capacitor's capacitance value due to mismatch by a Gaussian process with standard deviation [19], [20]

$$\sigma_{\Delta\text{C}/\text{C}} = \frac{A_\text{C}}{\sqrt{WL}}, \tag{12}$$

where W and L are the physical dimensions of the capacitor, and A_C is a process-dependent constant. To minimize capacitor mismatches, unit-capacitors are typically used, where the smallest capacitor is implemented multiple times and connected in parallel to create larger capacitors. Hence, doubling a capacitor's size is equivalent to doubling its area WL, therefore dividing its relative variations by $\sqrt{2}$.

However, during the binary search performed by a SAR ADC, a capacitor switched from $V_{\text{ref,p}}$ to $V_{\text{ref,n}}$ induces a voltage change to $V_{\text{s,d}}$ proportional to its size. The largest capacitor C_MSB has $\sqrt{2}$ times less variation than its subsequent capacitor $\text{C}_\text{MSB}/2$ on the CDAC, but induces a voltage change twice as large. Therefore, C_MSB mismatch induces a voltage error $\sqrt{2}$ times larger than its subsequent capacitor. The same reasoning holds every time the capacitor size scales down by a factor of two in the CDAC. Due to this behavior, MSB capacitor mismatch has the largest influence on the conversion result.

Fig. 3 shows the input-output TF f of a 5b-SAR ADC with 25% mismatch on its MSBs. Note that this is a large mismatch value used for illustration purposes only. The positive array CDAC_P has a $+25\%$ increase on its C_MSB, while CDAC_N has -25%. Before discussing the shape of the TF f, we make the observation that the MSB mismatch on the CDAC_P and CDAC_N can only influence the conversion outcome for positive and negative input voltages respectively: this is due to the fact that the MSB is the sign of the input.

When the mismatch is positive, the ADC over-corrects during the first quadrant-shift of the binary-search. This leads to input values close to zero being compressed to the closest-to-zero output, creating a saddle-point (see the $x \geq 0$ side in Fig. 3).

[2]The mismatch model presented in Sec. III-C can still be used for other switching strategies or some other types of ADCs. For example, residue amplifier mismatches of a pipelined ADC can also be modeled as capacitor mismatches.

[3]Comparator offset is another typical source of mismatch in SAR ADCs, but its effect only shifts the entire TF by the amount of voltage mismatch the comparator exhibits. This source of mismatch does not cause nonlinear distortion and it is, thus, common to fix its effect with digital post-processing.

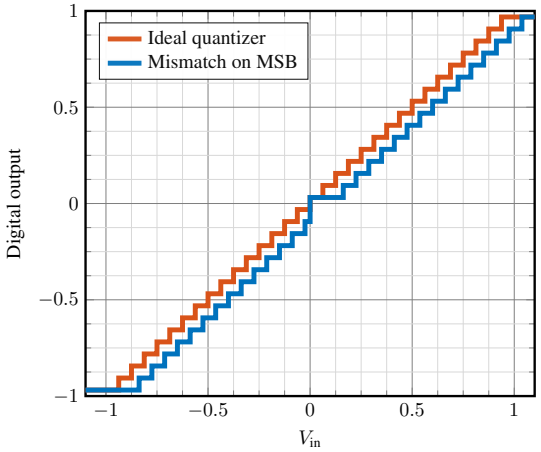

Fig. 3. Input-output function of a 5b-SAR ADC with 25% MSB mismatch.

It also extends the range of the ADC. When the mismatch is negative, the ADC under-corrects the first quadrant-shift. Outputs close to zero are then unreachable as shown on the $x < 0$ side in Fig. 3. Inputs close to the limits of the ADC range will also be clipped as the under-correction makes them larger than they should after the first shift. Overall, digital outputs close to zero are unreachable and the ADC's range is reduced. We also see in Fig. 3 that mismatches can lead to a fixed shift of the TF—to the right in the case illustrated here.

C. Model for SAR ADC with MSB-Mismatch

As explained in Sec. III-B, MSB mismatch has the strongest impact on the final conversion result. To disentangle the effect of the MSB mismatch from the effect of quantization itself, we will model the ADC TF as a piecewise linear function that includes the effect of clipping and the displacement from the ideal TF due to MSB-only mismatch. We will denote m_1 the TF shift induced by mismatch of the MSB on the positive CDAC and m_2 the one on the negative CDAC

$$m_1 = \frac{1}{2}\frac{\Delta\text{C}_{MSB,+}}{\text{C}_{MSB,+}} \quad \text{and} \quad m_2 = \frac{1}{2}\frac{\Delta\text{C}_{MSB,-}}{\text{C}_{MSB,-}}, \tag{13}$$

where the $1/2$ factor comes from the first quadrant shift amplitude of $\Delta V_\text{ref}/2$. For $x \geq 0$, the transfer function f for both possible signs of m_1 are

$$\begin{aligned} f_{m_1\geq 0}(x) &= (x-m_1)\,\text{rect}_{[m_1,1+m_1]}(x)+\text{u}(x-1-m_1) \\ f_{m_1<0}(x) &= (x-m_1)\,\text{rect}_{[0,1+m_1]}(x)+\text{u}(x-1-m_1), \end{aligned} \tag{14}$$

where the resulting functions are shown in Fig. 4 for a 25% capacitor mismatch, i.e., for $|m_1| = 0.125$.

On the $x < 0$ side, we can write the same equations by replacing m_1 by m_2 and mapping $f(x)$ to $-f(-x)$. However, there is no need to write this explicitly as we can simply re-assemble the different component together at the end of the calculation as done for (6) and (8). Now that we have the mismatched TF f, we take its derivative:

$$\begin{aligned} f'_{m_1\geq 0}(x) &= \text{rect}_{[m_1,1+m_1]}(x) \\ f'_{m_1<0}(x) &= \text{rect}_{[0,1+m_1]}(x) - m_1\delta(x). \end{aligned} \tag{15}$$

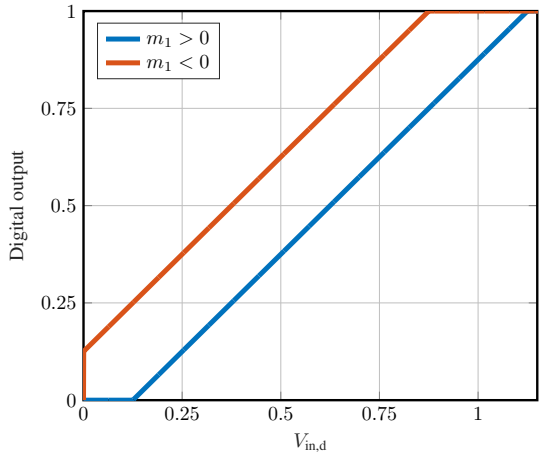

Fig. 4. Transfer functions depending on the sign for MSB mismatch modeling.

After integration, we can calculate the Bussgang gain

$$\beta = \begin{cases} Q(\frac{m_1}{\sigma_X}) - Q(\frac{1+m_1}{\sigma_X}) \\ \quad +Q(\frac{m_2}{\sigma_X}) - Q(\frac{1+m_2}{\sigma_X}), & m_1 \geq 0, m_2 \geq 0 \\ Q(\frac{m_1}{\sigma_X}) - Q(\frac{1+m_1}{\sigma_X}) \\ \quad +\frac{1}{2} - Q(\frac{1+m_2}{\sigma_X}) - \frac{m_2}{\sigma_X}\phi(0), & m_1 \geq 0, m_2 < 0 \\ \frac{1}{2} - Q(\frac{1+m_1}{\sigma_X}) - \frac{m_1}{\sigma_X}\phi(0) \\ \quad +Q(\frac{m_2}{\sigma_X}) - Q(\frac{1+m_2}{\sigma_X}), & m_1 < 0, m_2 \geq 0 \\ 1 - Q(\frac{1+m_1}{\sigma_X}) - Q(\frac{1+m_2}{\sigma_X}) \\ \quad -\frac{m_1+m_2}{\sigma_X}\phi(0), & m_1 < 0, m_2 < 0. \end{cases} \tag{16}$$

To compute the SDR, we also need $\mathbb{E}\big[f^2(X)\big]$. We have

$$\begin{aligned} f^2_{m_1\geq 0}(x) &= (x-m_1)^2\,\text{rect}_{[m_1,1+m_1]}(x) + \text{u}(x-1-m_1) \\ f^2_{m_1<0}(x) &= (x-m_1)^2\,\text{rect}_{[0,1+m_1]}(x) + \text{u}(x-1-m_1) \end{aligned} \tag{17}$$

and can calculate $\mathbb{E}\big[f^2(X)\big]$ by integrating only the $x \geq 0$ side and add together the separate m_1 and m_2 component to obtain the final result. Let us consider the following example with both mismatches being negative and $F(x,m) = f^2_m(x)$ to put the emphasis on the m dependency:

$$\begin{aligned} \mathbb{E}\big[f^2(X)|m_1, m_2 < 0\big] &= \int_{-\infty}^{0} F_{m<0}(-t, m_2)\phi_{\sigma_X}(t)\,\text{d}t \\ &\quad + \int_{0}^{+\infty} F_{m<0}(t, m_1)\phi_{\sigma_X}(t)\,\text{d}t \\ &= \int_{0}^{+\infty} (F_{m<0}(t, m_2) + F_{m<0}(t, m_1))\phi_{\sigma_X}(t)\,\text{d}t. \end{aligned} \tag{18}$$

Hence, we only integrate for positive values of x:

$$\begin{aligned} \mathbb{E}\big[f^2_{m_1\geq 0}(X)\big] &= (\sigma_X^2 + m_1^2)\left(Q\left(\frac{m_1}{\sigma_X}\right) - Q\left(\frac{1+m_1}{\sigma_X}\right)\right) \\ &+ Q\left(\frac{1+m_1}{\sigma_X}\right) + \sigma_X\left((m_1-1)\phi\left(\frac{1+m_1}{\sigma_X}\right) - m_1\phi\left(\frac{m_1}{\sigma_X}\right)\right) \end{aligned} \tag{19}$$

and

$$\mathbb{E}\big[f^2_{m_1<0}(X)\big] = \frac{\sigma_X^2 + m_1^2}{2} + Q\left(\frac{1+m_1}{\sigma_X}\right)(1 - \sigma_X^2 - m_1^2)$$

$$+ \sigma_X\left((m_1-1)\phi\left(\frac{1+m_1}{\sigma_X}\right) - 2m_1\phi(0)\right). \tag{20}$$

With those last results, we obtain an expression for $\mathbb{E}\big[f^2(X)\big]$ by adding the m_1 and m_2 component as shown in (18). Finally, the distortion power $\mathbb{E}\big[D^2\big]$ is obtained by inserting (20) together with the Bussgang gain from (16) into (3).

D. EFR of a Mismatched SAR ADC

With our model from Sec. III-C, we can directly compute the EFR for a given mismatch on the MSB for any input gain σ_X. Fig. 5(a) shows the EFR we can reach without quantization or clipping, and for a fixed mismatch on both MSB capacitors of $|m_{1,2}| = 0.125$ as described in Sec. III-C. Note that the mismatch value corresponds to the size of an LSB for a 4-bit data converter. In general, it is making sense to compare a mismatch value to the size of an LSB step width from the quantizer we consider. This consideration is typical for analog designers while implementing a SAR ADC: one aims at limiting variations to a fraction of an LSB, as it avoids missing output codes and ensures monotonicity in the TF of the data converter [18].

Fig. 5(b) shows the maximum achievable EFR as a function of $|m|$. We also show at which input gain σ_X this peak EFR is achieved. We first observe that the optimal gain σ_X is increasing with $|m|$. This is due to the trade-off between the non-linearity around zero and clipping artifacts: stronger mismatch makes it desirable to allow for larger clipping artifacts to avoid the jump (or saddle point) around zero. Examining Fig. 5(b) closely also reveals that the slope of peak EFR is of -1 bit/octave. This corroborates that comparing $|m|$ with the size of an LSB step is a meaningful procedure: halving the MSB mismatch leads to one additional bit of resolution, which is equivalent to halving the size of the LSB step width.

A 4-bit mid-rise SAR ADC following the model discussed in Sec. II-B with random mismatch on *all* its CDAC-capacitors has been simulated. The standard deviation σ_m of the random mismatch affecting the MSB is defined as in (13). As explained in Sec. III-A, the standard deviation of capacitor variations is divided by $\sqrt{2}$ every time the capacitance value is halved along the CDAC. Results are shown in Fig. 5(c), where we plot the 10%-quantile of the achieved EFR over CDAC-capacitors mismatch realizations. As mismatches are fixed once a chip is fabricated, one needs to consider a worst-case metric that will determine the production yield. We use the 10%-quantile which leads to 90% of fabricated ADCs exhibiting an EFR greater or equal than what is shown in Fig. 5(c).

With the 10%-quantile measure, we notice that capacitor mismatch starts to significantly impact the EFR when σ_m is of the order of a tenth of an LSB ($\Delta = 0.125$ in our 4-bit case). Similar results were obtained for all other resolutions in the simulated range (2 b to 6 b). As for the case of the MSB-mismatch-only model, we notice that the gain σ_X leading to the highest EFR increases with σ_m, but, due to quantization, with a bounded maximum EFR. The maximum value reached at low mismatch standard deviations σ_m corresponds to what was shown theoretically in Sec. II-D (cf. Fig. 1). In summary,

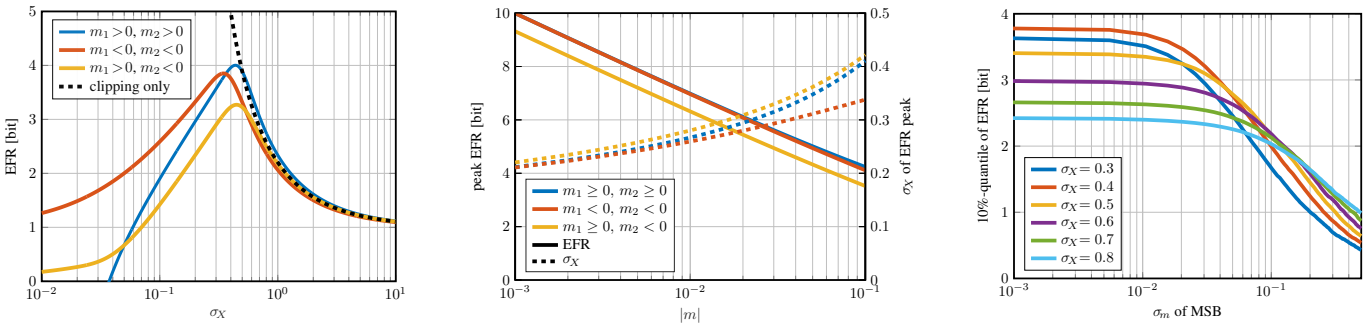

(a) Fixed $m = 0.125$.

(b) Peak achievable EFR as a function of m.

(c) 10%-quantile EFR of a 4 b mismatched ADC.

Fig. 5. EFR numerical results from the model with clipping and $|m_1| = |m_2| = m$ of MSB mismatch (a) and (b); simulation results for a SAR-ADC with mismatches on *all* CDAC-capacitors (c).

our simulation results confirm the intuition gained from the proposed mismatch model, and show the impact of using a worst-case measure like the 10%-quantile.

IV. Impact of SAR ADC Mismatch on Quantized Massive MU-MIMO Systems

As shown in Sec. III, capacitor mismatches can have a considerable impact on the effective resolution of a single SAR ADC. We now show that the same holds true for a quantized massive MU-MIMO wireless system, meaning that SAR ADC mismatches should *not* be ignored. To this end, we first introduce the system model and, then, show simulation results for the case of 4-bit ADCs with mismatches.

A. System Model

We consider a mmWave massive MU-MIMO uplink scenario in which $U = 16$ single-antenna user equipments (UEs) transmit data to a $B = 64$ antenna basestation (BS) with a uniform linear array. We assume that the UEs utilize power control to stay within ± 3 dB from the average power E_s. We use a QuaDRiGa mmMagic channel model [21] with 1° minimum angular separation between the UEs. We assume that the BS has perfect channel state information and uses a linear minimum mean-square error (LMMSE) detector. We consider the following discrete-time frequency-flat system model:

$$\mathbf{y} = f(\mathbf{z}) \quad \text{with} \quad \mathbf{z} = \mathbf{H}\mathbf{s} + \mathbf{n}. \tag{21}$$

Here, $\mathbf{y} \in \mathcal{A}^B$ is the BS receive matrix with output code alphabet $\mathcal{A}$ after the ADCs transfer function f, which includes quantization, mismatch, and clipping applied element-wise to the ideal receive vector $\mathbf{z} \in \mathbb{C}^B$, $\mathbf{H} \in \mathbb{C}^{B \times U}$ is the effective channel matrix including power control, $\mathbf{s} \in \mathcal{S}^U$ contains the transmitted 16-QAM data symbols, and $\mathbf{n} \in \mathbb{C}^B$ models thermal noise and is assumed to contain i.i.d. circularly-symmetric complex Gaussian entries with variance N_0. We define the average receive signal-to-noise ratio (SNR) as

$$SNR \triangleq \frac{\|\mathbf{H}\|_{\mathrm{F}}^2 E_s}{BN_0}. \tag{22}$$

Note that the SAR ADC TF f in (21) is applied entry-wise to the real and imaginary part of the ideal input vector $\mathbf{z}$, which means that for B BS antennas, we model an array of $2B$ SAR ADCs. As discussed in Sec. III-D, we define σ_m as the standard deviation of the Gaussian distribution modeling the effect of the MSB mismatch as defined in (13).

B. Simulation Results

We now show simulation results for the case of 4-bit SAR-ADCs. Different values for σ_m have been simulated, with values being different fractions of an LSB-step $\Delta = 0.125$. We use the uncoded bit error rate (BER) as our performance metric. It is important to note that we report the 10% worst-case uncoded BER, the same way as in Fig. 5(c), because mismatch does not have any averaging property once a chip is fabricated. Since a lower uncoded BER is better than a higher one, this 10% worst case corresponds to the 90%-quantile.

In Fig. 6(a), we see the 90%-quantile BER when mismatch is present on all CDAC capacitors in all SAR ADCs. We see a significant influence of SAR ADC mismatch: A standard deviation σ_m of $\Delta/2$ still leads to an uncoded BER floor of about 1%, which may be unacceptable for real-world wireless communication systems.

In Fig. 6(b), we observe similar behavior from the same simulation setup but where we only consider mismatches in the MSB capacitors. When comparing Fig. 6(a) to Fig. 6(b), we see that considering only MSB mismatch well-approximates a model with mismatches in all CDAC capacitors; this confirms the fact that only modeling the MSB mismatch, as in Sec. III-C, is a sensible approximation. We also notice that, for larger mismatches, the BER difference between the MSB-mismatch-only and the all-capacitor-mismatch models increases. Such large mismatches, however, would anyway be avoided in practical systems and are therefore less relevant.

In Fig. 6(c), we show the cumulative distribution function (CDF) of the uncoded BER at an SNR of 15 dB. Throughout this work, we chose to use σ_m and a 10% worst-case BER performance as our performance metric. Once a target BER specification has been set, all of the fabricated ADCs that are not meeting the specifications have to be discarded. The production yield of an ADC design is therefore defined by the proportion of designs not meeting the target specifications, no matter how good the rest of the designs happen to be.

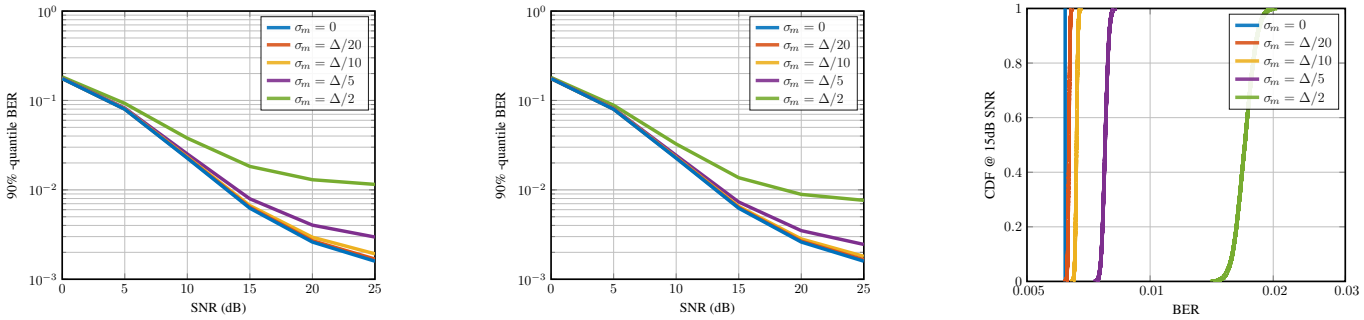

(a) Mismatch in all CDAC-capacitors

(b) Mismatch on MSBs only

(c) CDF at 15 dB SNR with mismatch in all CDAC-capacitors.

Fig. 6. Massive MU-MIMO simulation with 4b SAR ADC showing the 90%-quantile of the achieved uncoded BER.

V. Conclusions

We have analyzed the impact of capacitor mismatches caused by semiconductor process variations in widely-used successive approximation register (SAR) analog-to-digital converters (ADCs). We have developed an analytical model based on Bussgang's decomposition for the most-significant bit (MSB) mismatch of a SAR ADC, and we have analyzed the effects of quantization, clipping, and mismatches on a single ADC with a new figure of merit: the effective resolution (EFR). Our EFR results for a single SAR ADC reveal trade-offs between quantization errors, mismatch distortion, and clipping artifacts that all depend on the ADC's input gain. Finally, we have demonstrated the impact of SAR ADC mismatch on the bit error rate (BER) performance of a quantized massive MU-MIMO communication system. Our simulation results reveal that SAR ADC mismatches should *not* be ignored as they significantly affect the BER, even for low-resolution ADCs.

In practice, ADC designers typically must adhere to stringent specifications in order to limit the impact of process variations on ADC performance. Such strict constraints, however, are against the goal of utilizing simple, low-cost, and energy-efficient analog front-end designs in massive MU-MIMO systems with potentially hundreds of ADCs. Our work reveals that one can take capacitor mismatch into account when designing systems with SAR ADC arrays, which has the potential of extracting relaxed specifications—this paves the way for more efficient ADC array architectures.

References

[1] T. S. Rappaport, S. Sun, R. Mayzus, H. Zhao, Y. Azar, K. Wang, G. N. Wong, J. K. Schulz, M. Samimi, and F. Gutierrez, "Millimeter wave mobile communications for 5G cellular: It will work!" *IEEE Access*, vol. 1, pp. 335–349, May 2013.

[2] A. L. Swindlehurst, E. Ayanoglu, P. Heydari, and F. Capolino, "Millimeter-wave massive MIMO: the next wireless revolution?" *IEEE Commun. Mag.*, vol. 52, no. 9, pp. 56–62, Sep. 2014.

[3] L. Lu, G. Y. Li, A. L. Swindlehurst, A. Ashikhmin, and R. Zhang, "An overview of massive MIMO: Benefits and challenges," *IEEE J. Sel. Topics Signal Process.*, vol. 8, no. 5, pp. 742–758, Apr. 2014.

[4] E. Björnson, J. Hoydis, and L. Sanguinetti, *Massive MIMO Networks: Spectral, Energy, and Hardware Efficiency*. Now Publishers, 2017, vol. 11, no. 3-4, ser. Fundations and Trends in Signal Processing.

[5] C. Mollén, J. Choi, E. G. Larsson, and R. W. Heath, "Uplink performance of wideband massive MIMO with one-bit ADCs," *IEEE Trans. Wireless Commun.*, vol. 16, no. 1, pp. 87–100, Oct. 2017.

[6] S. Jacobsson, G. Durisi, M. Coldrey, U. Gustavsson, and C. Studer, "Throughput analysis of massive MIMO uplink with low-resolution ADCs," *IEEE Trans. Wireless Commun.*, vol. 16, no. 6, pp. 4038–4051, Apr. 2017.

[7] A. Abdallah, M. M. Mansour, L. M. A. Jalloul, and A. Chehab, "When quantized massive MIMO meets large MIMO with higher order modulation," *IEEE Commun. Lett.*, vol. 22, no. 12, pp. 2599–2602, Oct. 2018.

[8] R. Suarez, P. Gray, and D. Hodges, "An all-MOS charge-redistribution A/D conversion technique," in *IEEE Int. Solid-State Circuits Conf. (ISSCC)*, vol. XVII, Feb. 1974, pp. 194–195.

[9] B. Murmann, "ADC Performance Survey 1997-2022," [Online]. Available: https://github.com/bmurmann/ADC-survey.

[10] R. Walden, "Analog-to-digital converter survey and analysis," *IEEE J. Sel. Areas Commun.*, vol. 17, no. 4, pp. 539–550, Apr. 1999.

[11] W. Jiang, Y. Zhu, M. Zhang, C.-H. Chan, and R. P. Martins, "A Temperature-Stabilized Single-Channel 1-GS/s 60-dB SNDR SAR-Assisted Pipelined ADC With Dynamic Gm-R-Based Amplifier," *IEEE J. Solid-State Circuits*, vol. 55, no. 2, pp. 322–332, Nov. 2020.

[12] D. Pfaff, S. Moazzeni, L. Gao, M.-C. Chuang, X.-J. Wang, C. Palusa, R. Abbott, R. Ramirez, M. Amer, M.-C. Huang, C.-C. Lin, F. Kuo, W.-L. Chen, T. Y. Goh, and K. Hsieh, "A 56Gb/s Long Reach Fully Adaptive Wireline PAM-4 Transceiver in 7nm FinFET," in *IEEE Symp. VLSI C.*, Jun. 2019, pp. C270–C271.

[13] J. Lagos, N. Markuli, B. Hershberg, D. Dermit, M. Shrivas, E. Martens, and J. Craninckx, "A 10.1-ENOB, 6.2-fJ/conv.-step, 500-MS/s, Ringamp-Based Pipelined-SAR ADC With Background Calibration and Dynamic Reference Regulation in 16-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 57, no. 4, pp. 1112–1124, Jan. 2022.

[14] J. J. Bussgang, "Crosscorrelation functions of amplitude-distorted Gaussian signals," Massachusetts Inst. Technol., Cambridge, MA, USA, Tech. Rep., Mar. 1952.

[15] O. T. Demir and E. Bjornson, "The Bussgang decomposition of nonlinear systems: Basic theory and MIMO extensions [lecture notes]," *IEEE Signal Process. Mag.*, vol. 38, no. 1, pp. 131–136, Jan. 2021.

[16] C. M. Stein, "Estimation of the mean of a multivariate normal distribution," *Ann. Statist.*, vol. 9, no. 6, pp. 1135–1151, Nov. 1981. [Online]. Available: http://www.jstor.org/stable/2240405

[17] C. Studer, M. Wenk, and A. Burg, "MIMO transmission with residual transmit-RF impairments," in *Proc. Int. ITG Workshop on Smart Antennas (WSA)*, Feb. 2010, pp. 189–196.

[18] R. J. van der Plassche, *Integrated Analog-To-Digital and Digital-To-Analog Converters*. Kluwer Academic Publischers, 1994.

[19] M. Pelgrom, A. Duinmaijer, and A. Welbers, "Matching properties of MOS transistors," *IEEE J. Solid-State Circuits*, vol. 24, no. 5, pp. 1433–1439, Oct. 1989.

[20] M. Marin, S. Cremer, J.-C. Giraudin, and B. Martinet, "Modeling the mismatch of high-k MIM capacitors," in *IEEE Int. Conf. Microelectronic Test Structures*, Mar. 2007, pp. 115–119.

[21] S. Jaeckel, L. Raschkowski, K. Boerner, L. Thiele, F. Burkhardt, and E. Eberlein, "QuaDRiGa - quasi deterministic radio channel generator, user manual and documentation," Fraunhofer Heinrich Hertz Inst., Tech. Rep. v2.2.0, Jun. 2019.