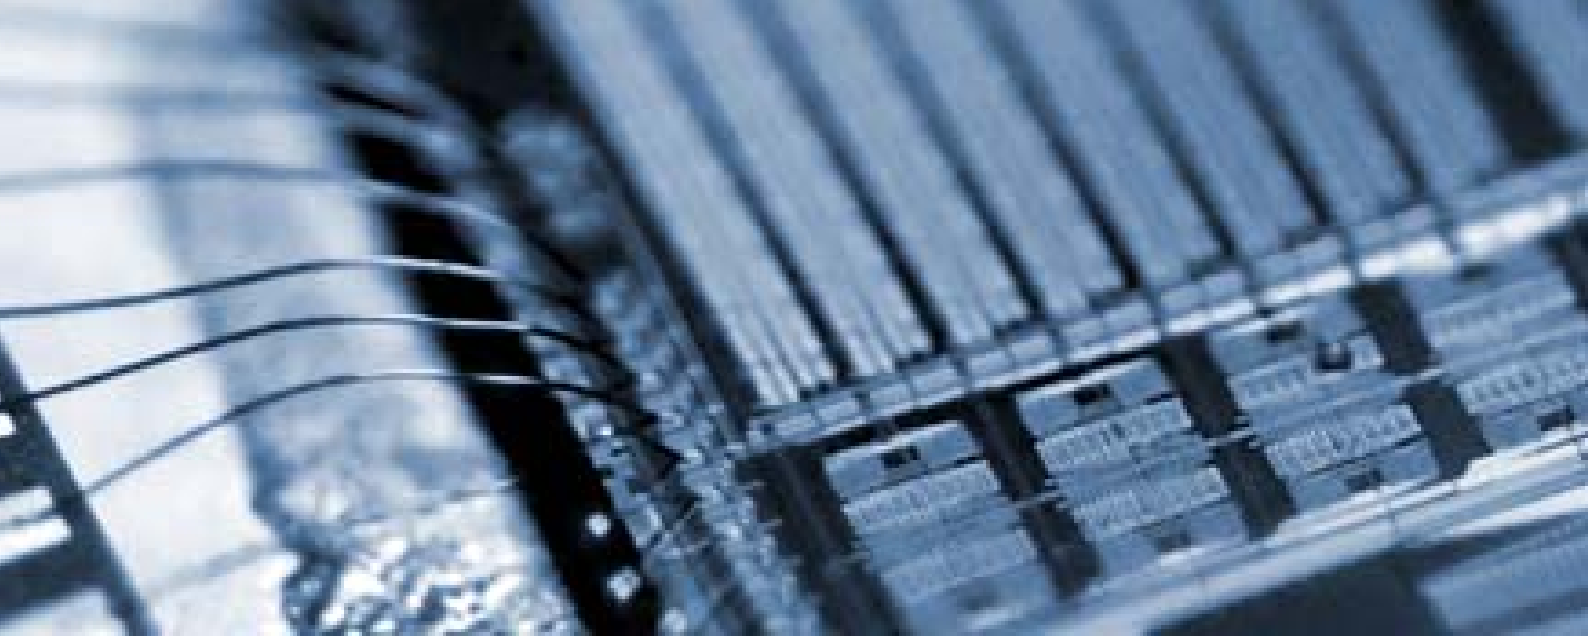

A/D-converter performance evolution

(1974 – 2012)

Bengt E. Jonsson

DOC NO: T-13:001 | REV: 1.1 | DATE: 2013-01-07 | ADMS DESIGN AB

Summary

This work analyzes the performance evolution over time for monolithic A/D-Converter (ADC) implementations reported in scientific publications. The work is based on an exhaustive search of IEEE journals and conferences central to the field from 1974 to spring 2012, and thus represents a near-exhaustive survey of reported scientific ADC data. Using the full set of historical data, empirically observed evolution trends are analyzed for all key performance parameters, including linearity, thermal noise, sampling jitter, and effective resolution. Evolution of power efficiency in terms of two commonly used figures of merit is also investigated. Performance parameters that appear to be in saturation have been identified. The results can be used to predict the state-of-the-art specifications of future ADCs.

About the document

This document is an almost exact replica of a **Converter Passion** article spanning a series of blog posts published in mid 2012. As the intention with this document is to present these blog posts in a more organized and reader-friendly format, *it has not been revised to reflect results reported during the second half of 2012*. It is the author's opinion that most, if not all of the results from this survey still apply, even if the state-of-the-art for some parameters have been nudged slightly. Visit Converter Passion for the latest development of some key parameters. The original online content was published here:

http://converterpassion.wordpress.com/articles/ad-converter-performance-evolution/

Copyright notice

Disclaimer and legal notice

Contents

List of Figures

List of Tables

1. Introduction

Monolithic A/D-converter integration has evolved exponentially over nearly four decades. In the early days it was a major achievement to implement full converter functionality within a single die, and power dissipation was high although resolution and speed were modest in comparison with today's expectations [1]-[7]. Thanks to the early work of pioneers and the continuously increasing effort of the research community, ADC performance has shown an exponential improvement with respect to many key parameters that closely resembles *Moore's Law* for the maximum number of components on a digital IC [8]-[9]. As more and more of electronic systems move into the digital domain, the evolution of A/D-converter technology is an essential enabling factor for a wide range of applications: Low power and small chip area allows the ADC to be integrated into a larger *system-on-chip* (SoC), while advancing the sampling rate at a given resolution or the resolution at a given sampling rate improves system performance or enables entirely new kinds of systems to be built. From a system point-of-view, it is therefore crucial to understand the rate of ADC performance evolution, so that system-level roadmaps can be adapted to what is realistically expected from ADC technology at any given time.

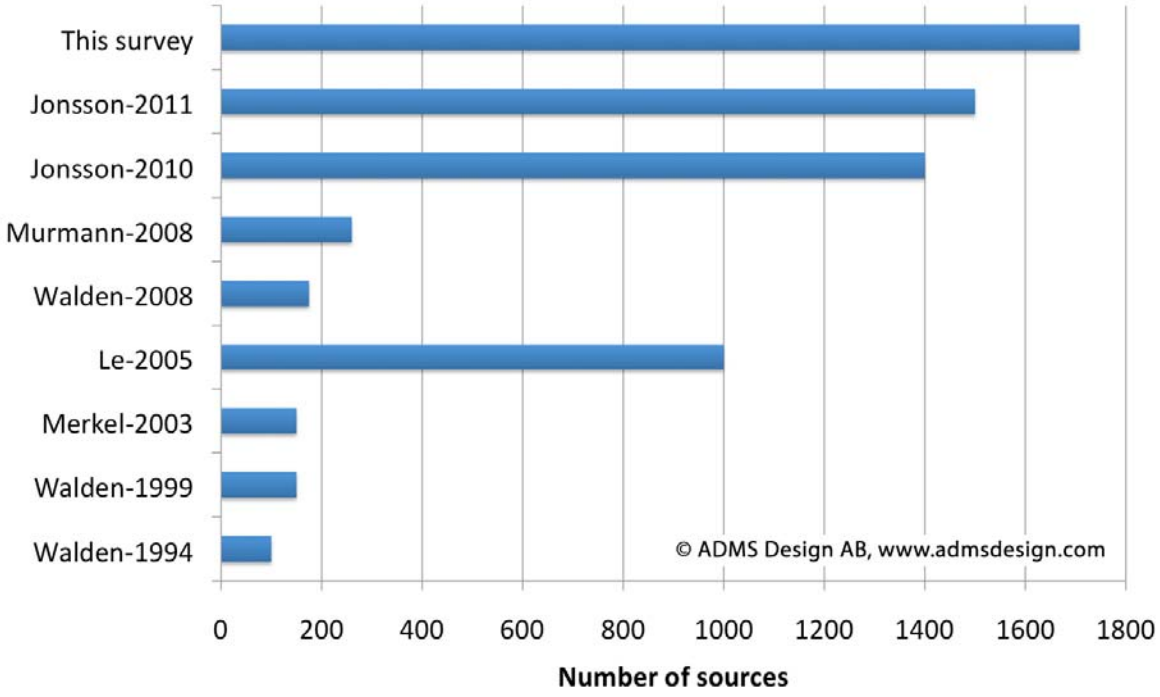

Figure 1.1. Other ADC surveys.

Several authors have presented works that blend evolution surveys and derivations of theoretical physical limits to various degrees. One of the more widely cited publications is the survey by *Walden* [10], examining the performance evolution estimated from 150 commercial and scientific ADCs reported between 1978 and 1997. *Walden's* survey, which had been partially presented in an earlier publication [11], was the first survey based on a larger set of empirical data, and the observed trends for "SNR-bits" was compared to a set of theoretical limits including thermal noise, jitter, comparator metastability, and *Heisenberg's* principle of uncertainty. An average evolution rate of 1.5 "SNR-bits" every 8 years was estimated from

the data set. Approximately one decade later, *Walden* published two updated versions of the survey in [12] and [13] based on slightly larger data sets. *Le* et al. examined the performance of nearly 1000 commercial ADCs released over a period of 20 years [14], and the survey by *Murmann* [15] is a significant recent contribution to the field, including approximately 260 scientific ADCs reported 1997–2008. *Merkel* and *Wilson* surveyed 150 commercial and scientific ADCs with specifications suitable for defense space applications [16]. ADC performance vs. CMOS scaling was investigated by the author in a survey based on 1100 scientific papers in [17], and an analysis of ADC performance evolution was also presented in [18]. Other relevant work is found in [19]-[25]. For more on prior art ADC surveys, see *Appendix: A survey of ADC surveys* on page 35.

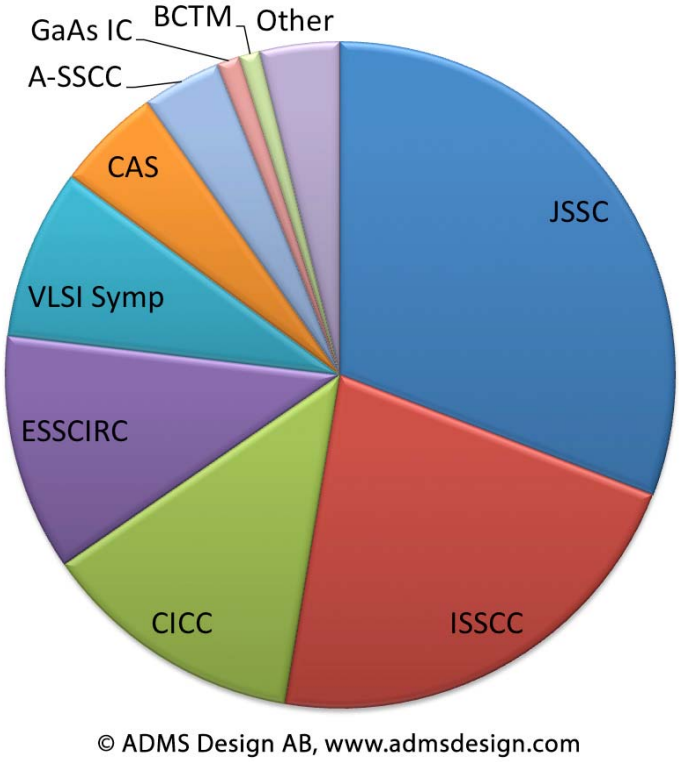

Figure 1.2. ADC paper distribution by source. Only papers reporting measured ADC implementations have been included. Data until Spring 2012.

1.1 About this survey

This work reviews the evolution of several key performance parameters based on a set of over **1700 scientific publications** published from 1974 to spring 2012. It represents a *near exhaustive survey* of experimental ADC performance data reported in scientific publications. A steadily increasing number of ADC implementations have been reported in the scientific literature, growing from a handful to well over 100 publications per year. This trend is an indication of the significance given to ADCs, as well as the effort invested in the field. The distribution of publications over time is shown here, and the distribution by source is shown above. Data was collected from an exhaustive survey of ADC papers in the *IEEE Journal of Solid-State Circuits* and *IEEE Transactions on Circuits and Systems*, as well as

the following conferences: *ISSCC*, *ESSCIRC*, *CICC*, *Symposium on VLSI Circuits*, *A-SSCC*, *GaAs IC Symp*, and *BCTM*. Additional papers from non-exhaustive searches of other sources have also been included. To the best of the author's knowledge, this makes it the largest and most comprehensive survey of reported scientific ADC implementations to this date. The work presented relates to general performance trends over time, and parameters that impose similar performance limits on a broad range of architectures. As an additional dimension, the data is frequently divided into its delta-sigma modulator (DSM) and Nyquist subsets. A deeper analysis of power-performance trade-offs would require a full breakdown by architecture, since there are significant differences between various architectures [14], [25]. Such a study of architectural differences is well beyond the scope of this work and will be treated elsewhere. The overall conclusion of this survey is that many parameters have improved or evolved exponentially, but several noise-related performance measures appear to have saturated.

2. CMOS node adoption

The rate at which scientific ADC implementations migrate to newer CMOS technology is discussed in this section. It was previously observed in [18], using a more one-dimensional approach and data until March 2010. Here, updated ADC survey data is used, and the 2-D distribution of scientific ADC implementation papers over *CMOS node* and *publication year* is analyzed. The result is illustrated by the "heat contours" in Figure 2.1. Starting with dark blue, the colors represent paper counts of 0, 1, 2, 5, 10, 15, … and 40/year, respectively.

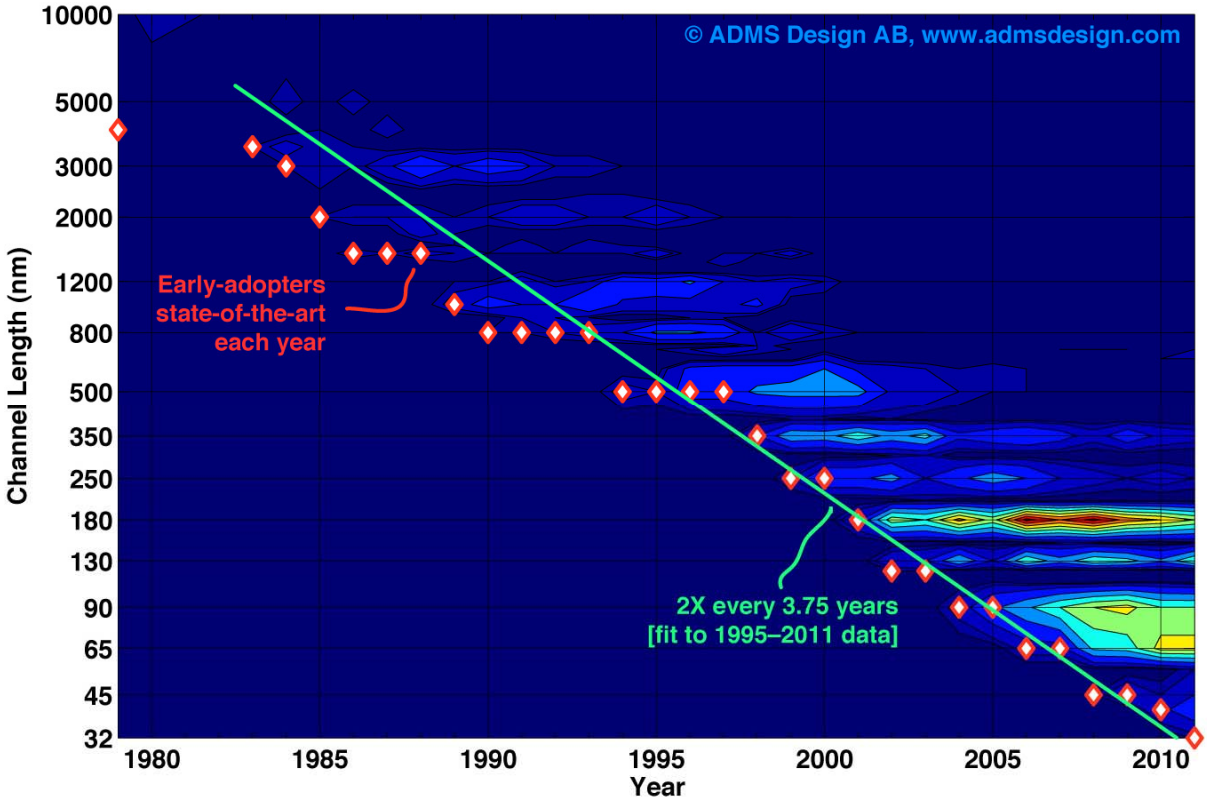

Figure 2.1. Distribution of CMOS nodes used for scientific ADCs over time. Color represents number of publications. The early adopter state-of-the-art data points (◇) are superimposed along with a scaling trend estimated from 1995–2011 data.

2.1 Observation of technology adoption

Figure 2.1 illustrates several key aspects of how the scientific ADC community has adopted new process technology:

- The lower edge of the contours represents the **early adopters**. It defines the *state-of-the-art scaling front* for ADCs. In [18] it was estimated that this front scaled by an average **factor of two every 5.4 years** until 1995. After 1995, the adoption rate increased to **2× every 3.75 years**, which is illustrated by the exponential trend fit. The data points used for the trend estimations are superimposed onto the contour plot.
- The "center-of-mass" illustrates the average node-adoption by the main body of ADC scientists. Although a highly subjective visual estimation, my impression is that the *mainstream adoption rate* is higher from 180 nm and below.

- The horizontal extension of each node reveals *its lifetime in scientific publications*. Popular nodes can remain active for well over a decade. Therefore, the correlation between CMOS node and publication year is weak. In other words: you can't make a good observation of the effects of scaling by simply looking at how something evolves over *time*. Because of the long lifespan of major nodes, they also have time to undergo a *maturing process* as the collective understanding of how to best use the node accumulates. ADC performance vs. scaling and the concept of maturing nodes was treated in [17].
- 180 nm appears to be the *all-time favorite node* for CMOS ADC designs to this date. This was also observed in [17].
- Nodes as old as 0.35 µm are still active in publications.

2.2 Adoption rates, systems on chip, and the scaling gap

Traditionally, there has been a lag – or "*scaling gap*" – between analog/mixed and digital ICs. Digital ASICs have nearly always benefited from using the most recent technology, whereas analog/mixed ICs have faced new design challenges for every step of scaling. Consequentially, ADC designers have lingered in older, or custom, technologies where they knew they could meet the spec, while digital ASIC designers switched to new nodes as soon as possible. This approach was acceptable – perhaps even optimal – as long as ADCs were used as stand-alone components. Moving into **the age of the SoC** (*system on chip*), the scaling gap is increasingly unacceptable. The A/D-converters must be on the same chip as the rest of the circuit, and migrating the digital parts backwards is almost never an option. Therefore the gap must close.

Personally, I believe that this SoC-driven need to close the scaling gap is the most likely explanation for the increased rate of early adoption observed in [18] and mentioned above.

2.3 Future ADC scaling

If the current scaling trend should continue, early adopters would implement ADCs in **5.5 nm** CMOS by 2020. As mentioned in [18], that is well below the technologies predicted available for RF/AMS design by 2020 [26]. As the gap closes between analog and digital, we will therefore see a slowdown in adoption rate. In fact, the RF/AMS data in [26] suggests that the scaling gap is already quite small. A minimum L = 24 nm for HP logic in 2011 should be compared with the most deeply scaled ADC in 32 nm CMOS, presented by a team from Intel [27]. Since a small lag between digital and RF/AMS ASICs (as well as between technology "*year of production*" and ADC "*year of publication*") may be inevitable, ***a slowdown in early-adoption ADC scaling could be just around the corner.***

3. Low-voltage operation vs. scaling

As analog-to-digital converter implementations migrate to scaled-down CMOS technologies, they also face the inevitable downscaling of supply voltages (VDD), and hence signal swing [26]. A signal chain with a weak signal is more likely to suffer from noise than one with a strong signal. The scaling trends for VDD are therefore important as a background to the A/D-converter noise performance trends that will be treated in a few subsequent sections. There are also other reasons for a circuit designer to keep an eye on the evolution of supply voltage, such as the considerable challenges for high-gain OP-amp design or sampling linearity under low-voltage operation, so I hope you'll find this section useful even on its own.

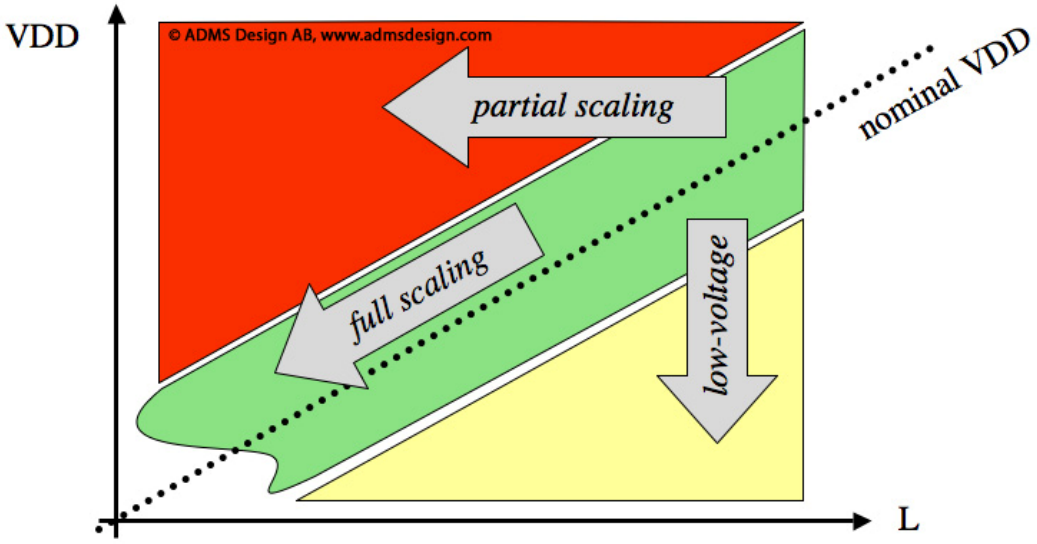

Figure 3.1. Two-dimensional view of CMOS scaling: Channel length and VDD.

3.1 Voltage scaling: Stragglers, mainstream and pioneers

Contrary to the minimum channel length (*L*), we can *choose* to go lower than the nominal supply voltage specified for a process. Possibly at our own risk, but at least it can be done. It has therefore been possible to scale VDD "ahead" of the node you actually use. Unless you have direct access to a semiconductor fab, you can't do that with respect to *L*. This degree of freedom – at least for experimental ADC designs – has lead to the situation illustrated by Figure 3.1:

- Some (or as we shall see below, most) designs use the nominal supply voltage recommended for any given CMOS node.
- Others may use the same node, but the design is not "fully scaled". It relies on higher-than-nominal VDD, and possibly optional process steps that effectively recreate older and less scaled device technology as well.
- A third category not only embraces the full scaling, but actually use a more aggressive scaling of supply voltages. These are **the low-voltage pioneers**.

This section will observe how supply voltage distribute over CMOS node for scientifically reported ADCs, and attempt to extract evolution trends and trajectories for VDD vs. *L*.

3.2 ADC supply voltage vs. CMOS node

As pointed out in [17], the reported supply voltage can vary as much as *one order of magnitude within the same node* for scientific A/D-converters. This is illustrated by the scatter plot of {*L*, VDD} for the entire CMOS ADC data set in Figure 3.2. The VDD used in the plot is the *highest* supply voltage applied to each ADC, and the evolution of low-voltage state-of-the-art over CMOS nodes has been highlighted. *Wismar*, et al., reported a **90 nm** VCO-based Δ-∑ modulator implementation running at **0.2 V** supply voltage (operational @ 0.18 V), which is the **lowest VDD published to this date** [28].

Note that Figure 3.2 differs from a similar graph in [17] in that the graph here is based on two more years of empirical data, and the plot in [17] shows the *lowest* VDD applied to each design instead of the highest. Also, the trajectory for the *de facto* nominal supply voltage vs. CMOS node is overlaid in Figure 3.2. It is not necessarily the "official" VDD, but instead it was derived from the supply used by the majority of designs reported for each node. For nearly all nodes, the choice was abundantly clear. In 65 and 90 nm, however, there were significant subsets of designs using 1 V instead of the 1.2 V used by the majority.

Starting at 1.2 µm, the ultra-low voltage state-of-the-art appears to have followed a distinct trend of evolving as approximately *one fifth of the nominal VDD*. Because of the slowdown in nominal VDD scaling, that trend still holds in relation to the 0.2 V reported by *Wismar*.

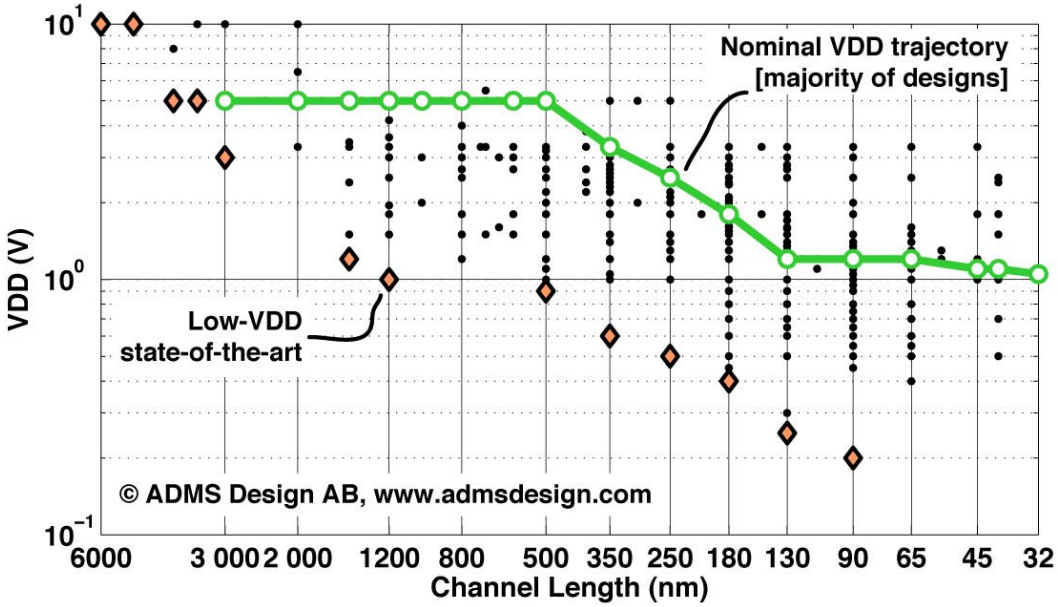

Figure 3.2. Supply voltages used for scientific ADCs vs. CMOS node. The low-voltage state-of-the-art data points (◇) have been highlighted, and the nominal/majority VDD trajectory (○) superimposed.

Although the scatter in Figure 3.2 shows all reported combinations of {*L*, VDD}, it does not reveal the distribution across scientific ADC papers. This is done in Figure 3.3, where color contours represent the number of papers falling into a certain two-dimensional histogram bin. The bins used in this plot have been selected manually in order to create a meaningful, yet readable plot, so that bin centers align with major nodes on the *L* axis, and the most frequently used or otherwise interesting values on the VDD-axis. Furthermore, a non-linear, truncated, ad hoc mapping of contour levels was applied to handle the steep peeks at certain bins while still retaining visibility of all non-zero bins. The contours thus yield a simplified view

of the actual distribution, and cannot be used to derive the actual bin counts or exact distribution. For completeness, Figure 3.4 shows the full distribution of all unique combinations of {*L*, VDD} reported for scientific ADC implementations until Q1-2012.

Figure 3.3 and Figure 3.4 show that *the vast majority of experimental ADCs reported in scientific papers use the nominal supply voltage* for each CMOS node, even if there is a large spread of actual VDD values used in each node. The variation extends significantly below and above the nominal values. We can also observe that the scaling rate for nominal supply voltage with CMOS node appears to have leveled out after 130 nm. Projected VDD for future nodes are found in [26].

In the next section we will look at the trends for VDD over *time*.

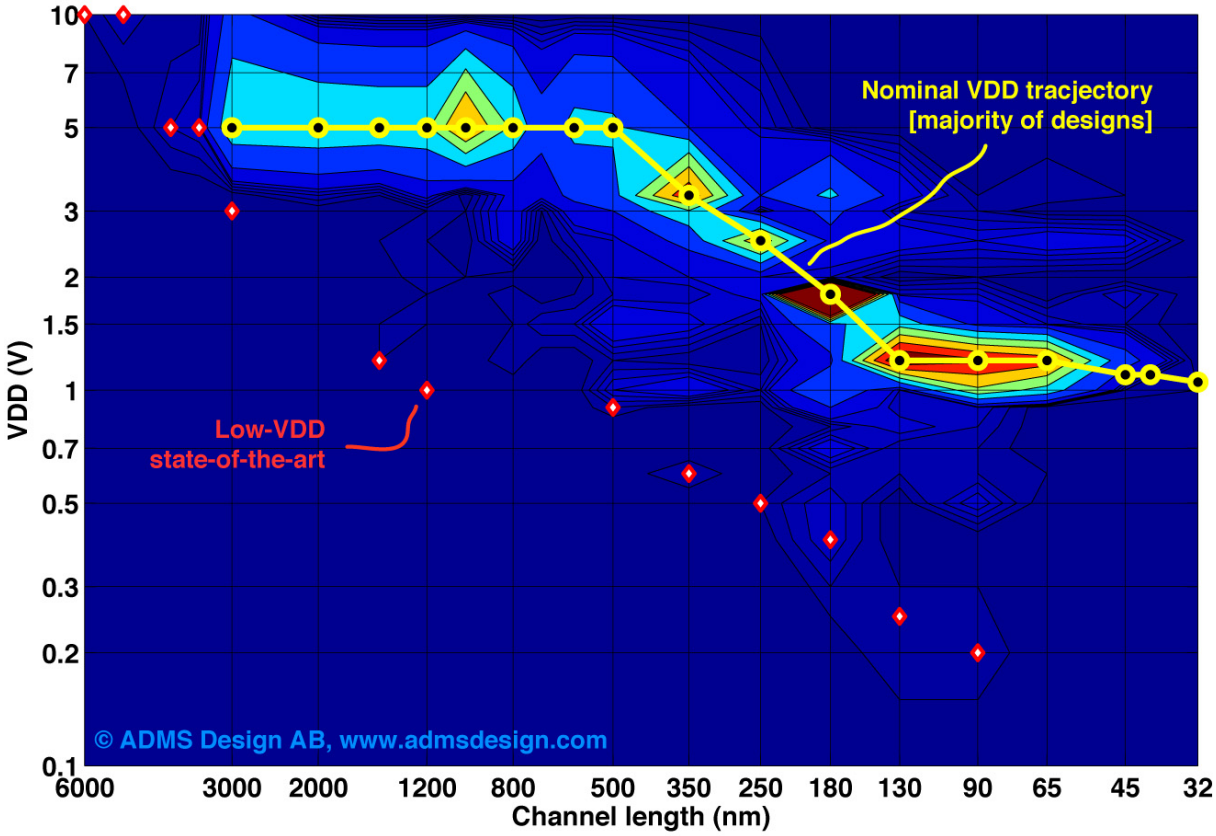

Figure 3.3. Distribution of supply voltage used for scientific ADCs vs. CMOS node. Color contours indicate density of publications. The low-voltage state-of-the-art data points (◇) are superimposed along with the nominal VDD trajectory (○).

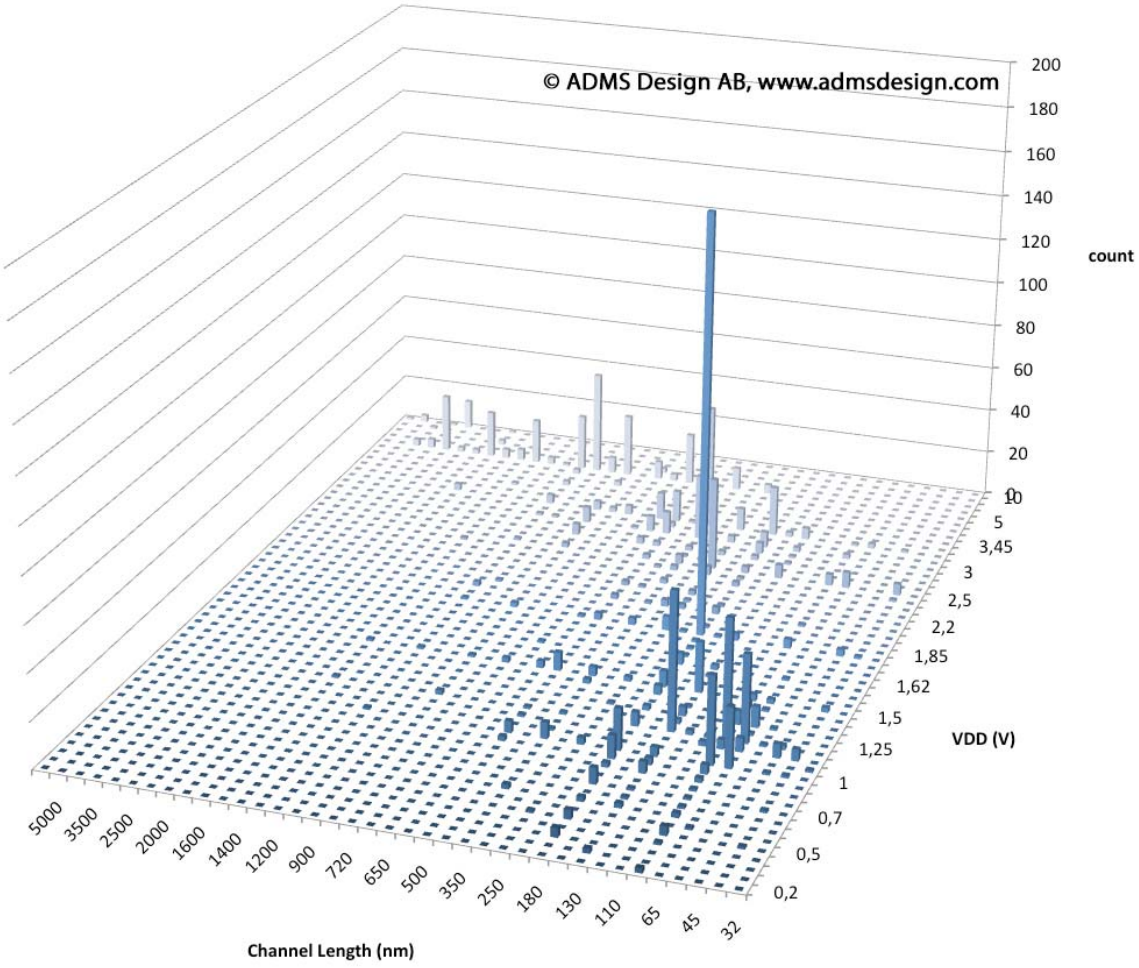

Figure 3.4. Distribution of all unique combinations of VDD and *L* (node) reported for CMOS ADC implementations in scientific papers until Q1-2012. Bin grid is not to scale.

4. Low-voltage operation over time

In the previous section we observed the trends for supply voltage (VDD) vs. process scaling (*L*). In this second part we will complete the picture by looking at VDD trends over time. The timing for introduction of new process technology, and the nominal supply voltage for future nodes, are reasonably well defined through the continuously updated *International Technology Roadmap for Semiconductors* (ITRS) [26], but at least two ADC-related aspects are not controlled by the ITRS scaling roadmap:

- The rate at which mainstream ADC research activities will migrate to newer CMOS technology.
- To what extent the ADC research community will attempt to push the envelope with respect to ultra-low voltage operation.

Regarding the former, it was observed in section 2, "*CMOS node adoption*", that the number of early adopters for each node is very small. In any year, the absolute majority of experimental ADCs have so far been implemented in technology being 2–5 generations behind the scaling front. How the "mainstream" will behave in the future is next to impossible to predict, as it is influenced by future industrial needs, research grant policies, research community group dynamics, journal and conference publication targets, as well as many other hard and soft parameters of which we know very little today.

The latter depends on a handful of pioneers choosing to explore the outer limits of ultra-low voltage ADC operation. It was seen in the previous section that there have been rather few attempts to push in this direction, which reveals that only a few groups have historically chosen this focus. If no one decides to have a shot at the current world record – the 0.2 V, VCO-based ∑-Δ modulator presented by *Wismar* et al. in [28] – we may never see it nudged. It is therefore very difficult to predict the future VDD trends for analog-to-digital converters, both with respect to the ultra-low voltage state-of-the-art, and the mainstream supply voltage. What we *can* do, however, is to observe historical trends and use them as a reference.

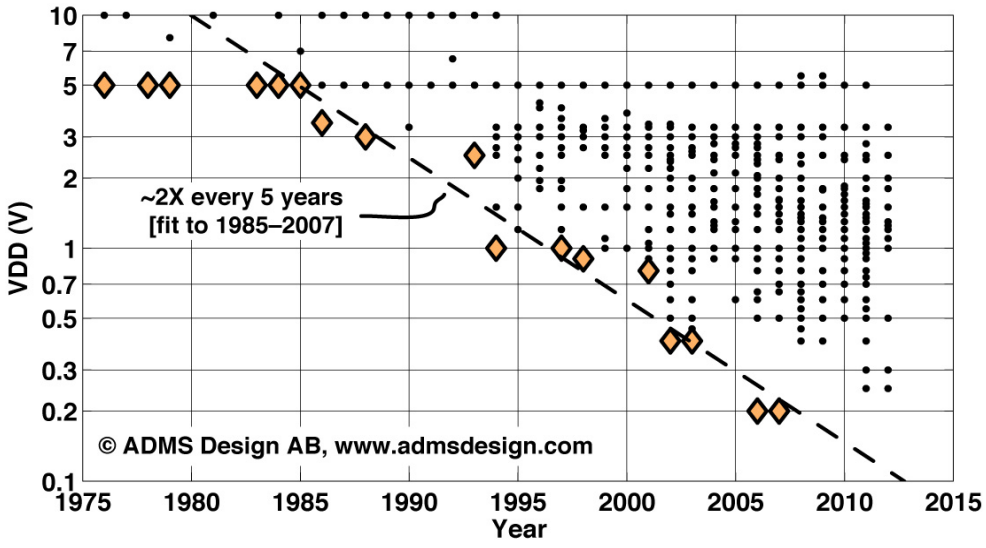

Figure 4.1. Supply voltages used for scientifically reported CMOS ADCs over time. Data points representing the evolution of low-voltage state-of-the-art have been highlighted (◇). Trend line fit to 1985-2007 data.

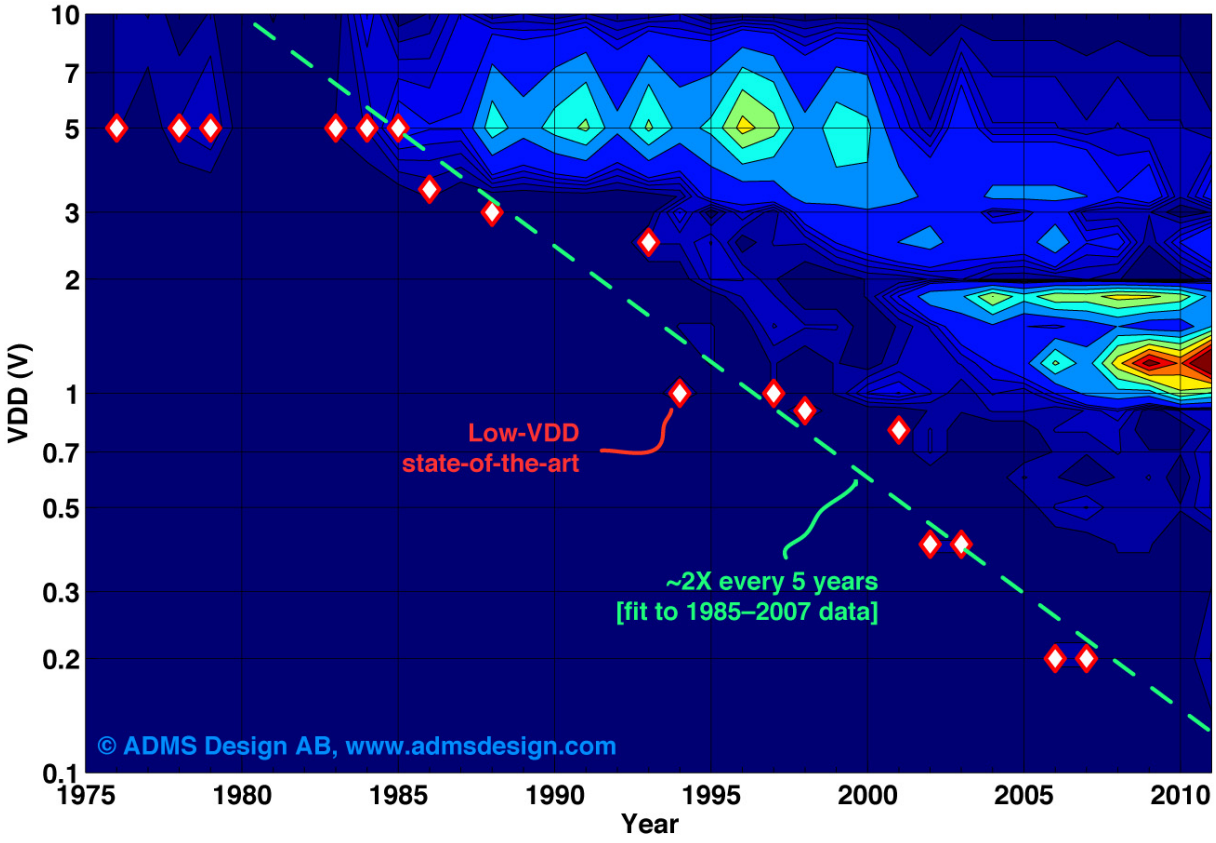

Figure 4.2. Voltage supplies used for scientific ADCs over time. Color represents number of publications. The low-voltage state-of-the-art data points (◇) are superimposed along with a scaling trend estimated from 1985–2007 data.

4.1 Observation of supply voltage trends

Figure 4.1 shows the voltage supplies reported for CMOS A/D-converters reported in scientific publications until Q1-2012. The graph shows the highest supply voltage applied to the circuit. It means that, if a circuit used several independent supplies, then VDD = max(VDD1, VDD2, …, VDDn), so that true low-voltage operation is promoted. The evolution of low-voltage state-of-the-art has been highlighted.

A similar graph in [18] shows data for *all* ADCs (CMOS, as well as bipolar and BiCMOS) but with data only to Q1-2010. Focusing on CMOS, and adding two more years of empirical data yields a different scatter. Nevertheless, the low-voltage state-of-the-art sequence here is nearly identical to that in [18] because the global state-of-the-art almost completely coincides with CMOS ADCs, and also did not improve since 2006. As observed in [18], the lowest reported VDD remained unchanged at 5 V until 1985, after which it started to follow a noisy but distinct scaling trend for 20 years. Fitting to the state-of-the-art data from 1985–2007, yields that the lowest reported VDD was scaled by **~2× every five years** during this period.

Figure 4.2 shows the distribution of scientific ADC implementations over supply voltage and publication year as a contour plot. The state-of-the-art data points and trend fit from Figure 4.1 have been superimposed for reference. Just as in the previous section, manually selected bin centers and non-linear contour levels have been used in order to render a meaningful and readable (but simplified) plot. The main purpose is to illustrate *the difference*

between mainstream VDD and state-of-the-art low-voltage operation each year. It is observed that:

- Mainstream focus remained at 5 V for over 20 years.
- The state-of-the-art VDD scaling front started to go below 5 V around 12–15 years before any noticeable change in the mainstream focus.
- The low-voltage scaling front appears to be approximately 5–6× below, and 10–15 years ahead of the mainstream VDD for each year.
- Supply voltages from 5 V and down seem to have an extremely long lifetime in publications.

4.2 Future VDD scaling for ADCs

I'm *very* aware that there are good reasons why ultra-low VDD scaling may not be able go much further, so please note that I'm not saying here that it *will*. Perhaps it is physically impossible, or functionally meaningless to go significantly further than the 200 mV operation achieved by *Wismar*, et al. On the other hand, I'm old enough to have heard one "hard" limit after another being suggested for MOST scaling, and we're still scaling them. So, let's just see where we would end up if it should turn out to be possible also for the voltage supply:

If the current trend for ultra-low voltage ADCs should be maintained, the **low-voltage pioneers** would have to publish ADCs according to the following approximate schedule:

Year	VDD
2015	73 mV
2020	36 mV
2025	18 mV

Table 4.1. Projections of lowest supply voltage at current rate of scaling.

Again, I'm not saying that it will happen. But I still found it interesting to see what kind of supplies the historical trend is projecting towards. Does anyone dare to predict a hard limit for A/D-converter supply voltage? Will we ever see an ADC operating at 73 mV? Is 36 mV impossible? What are the possibilities in context of the impossibilities? In case anyone wish to make their own projections, the trend fit expression is:

Eq 4.1 $$VDD = 10^{-0.061011 \times year + 121.7998}$$

5. Thermal noise

In the quest for high-resolution, high-bandwidth ADCs it's got to hurt when you hit the noise floor. Looking at the evolution trends for several noise-related parameters in *scientific ADC publications*, it certainly looks as if we've already touched that floor, or at least are very close to smack into it. We'll look at performance trends for *absolute* and *relative* noise, as well as *sampling jitter* over a series of three sections so you can make your own assessment. But before diving into the thermal noise evolution below, I will first serve up some background theory:

5.1 Separation of noise sources

Quantization errors, thermal noise, flicker noise, switching transients and sampling jitter all contribute to the noise power observed at the ADC output. In order to enable a deeper analysis of noise performance, it is desirable to separate different contributors from each other. A first step is to separate ideal quantization errors inherent to the algorithm from actual circuit noise.

5.1.1 Quantization vs. circuit noise

The ideal quantization error is often treated as a noise component, although in reality it is a deterministic signal with a large number of harmonics yielding a noise-like spectrum for ADCs with a sufficient number of bits *N* [29]-[30]. The ideal *signal-to-quantization-noise ratio* (SNR_Q) in dBFS is defined as

Eq 5.1
$$SNR_Q = 6.02 \times N + 1.76$$

Any noise above the ideal quantization noise level is due to the actual circuit implementation, and is referred to as "*circuit noise*" in this treatment. It originates from both analog and digital circuits, and can be further subdivided as described in the following subsections and sections. Using Eq 5.1 for SNR_Q, the *signal-to-circuit-noise ratio* (SNR_C) is estimated by noise power subtraction based on reported *signal-to-noise ratio* (SNR) and resolution *N* as described by Eq 5.2.

Eq 5.2
$$SNR_C = -10\log_{10}\left(10^{-\frac{SNR}{10}} - 10^{-\frac{6.02+1.76}{N}}\right)$$

Since SNR and *N* were simultaneously reported in only 22% of all publications, SNR was used instead of SNR_C when *N* was not available, and the *signal-to-noise-and-distortion ratio* (SNDR) was used as a conservative estimate of SNR when the latter was not explicitly reported. Using these conservative approximations, 86% of all reported ADCs could be assigned an SNR_C estimate.

Although Eq 5.2 gives a better view of circuit noise, some caution needs to be exercised when the reported SNR is very close to the ideal quantization noise. In such a case the subtraction in Eq 5.2 becomes numerically unsound, and small rounding errors can have great impact on the estimated SNR_C.

5.1.2 Noise sources excluded from this survey

Some well-known noise sources are unfeasible to analyze from ADC survey data, and will not be included in this treatment. These are:

- Flicker noise
- Switching noise
- Noise due to non-zero DNL errors

The effects of ***flicker noise*** [31] can be suppressed by *chopper techniques* [32] or the use of *correlated double sampling* [33]-[34]. Very few ADC publications report flicker noise performance explicitly, and in many wide-band applications the flicker noise is not of particular concern. Thus it is not included in this study.

On-chip ***switching noise*** result from transients caused by the switching of analog and digital circuits during normal operation. Such interference propagates as transients on power supply lines, through the substrate, or by inductive and capacitive coupling. Its effect can be suppressed using supply separation, decoupling, guard-rings, electromagnetic shielding, and by controlling the slope or timing of dominating transients such as those generated by digital output pins. In a large circuit with many simultaneously switching elements, any deterministic relations between input signal and on-chip state changes are likely to be obscured by the large amount of switching nodes. Switching interference is then observed as an increased total noise level, possibly indistinguishable from thermal noise. Switching noise is usually not reported explicitly in ADC publications, and is difficult to estimate from other data typically reported. Therefore it is not treated in this survey.

Due to component mismatch, real ADC implementations have finite (non-zero) ***differential non-linearity*** (DNL) errors. Such DNL errors cause a non-uniform distribution of decision levels that increase the quantization error power above that of the ideal quantizer described by Eq 5.1. The exact amount of additional noise depends non-trivially on the magnitude and distribution of DNL errors across the converter code range, and not only on the worst-case DNL error normally reported. Furthermore, because of the difficulty to define a “typical” DNL error magnitude that would correctly represent the entire gamut of ADC implementations surveyed, “DNL-noise” is not included in the treatment.

5.1.3 Thermal Noise

Unlike non-linear distortion, memory effects, etc., *thermal white noise* is entirely random from sample to sample and thus impossible to predict or compensate for by calibration. It is therefore a more fundamental limitation of ADC performance than non-linear distortion or memory effects. A simplified model of ADC noise refers the noise to a noisy input source resistance R_n while assuming the rest of the signal path to be noiseless [10]. The mean squared thermal noise over R_n in a bandwidth $BW = f_s/2$ is [10], [31]

Eq 5.3
$$v_n^2 = 4kTR_n f_s/2$$

where k is Boltzmann’s constant (1.38×10-23 J/K) and T is the absolute temperature in Kelvin. Assuming a full-scale peak-to-peak input swing of V_{FS}, the equivalent single-tone SNR is

Eq 5.4
$$SNR = 10 \times \log_{10}\left(\frac{V_{FS}^2}{16kTR_n f_s}\right)$$

Eq 5.4 defines a theoretical limit on the achievable SNR for an ADC with a given input swing and impedance level. Assuming fixed values of T and V_{FS}, the theoretical limits for different R_n can be added to SNR or *effective number-of-bits* (ENOB) vs. f_s plots of reported experimental data as a visual guide, e.g., as done by *Walden* [10]. In section 9, "*Sampling rate and resolution*", something similar will be done with the ADMS Design survey data set. It should be noted, however, that V_{FS} vary between 5 mV [35] and 20 V [32] in the data set. This corresponds to a **72 dB variation** in Eq 5.4, and simply assuming a fixed V_{FS} therefore gives a very coarse approximation of R_n. Hence the reported values of V_{FS} were also accounted for in this treatment. It was additionally compensated for the fact that ideal quantization noise is not contributing to the thermal noise, and R_n was therefore estimated from the estimated *circuit* noise rather than the *total* noise, as described by Eq 5.5. As far as the author is aware, this compensation for quantization noise and full-scale range (FSR) was not applied in any other ADC survey when estimating thermal noise performance from empirical data.

Eq 5.5
$$R_n = \frac{V_{FS}^2}{16kTf_s} \times 10^{-\frac{SNR_C}{10}} = \frac{V_{FS}^2}{16kTf_s} \times \left(10^{-\frac{SNR}{10}} - 10^{-\frac{6.02 \times N + 1.76}{10}} \right)$$

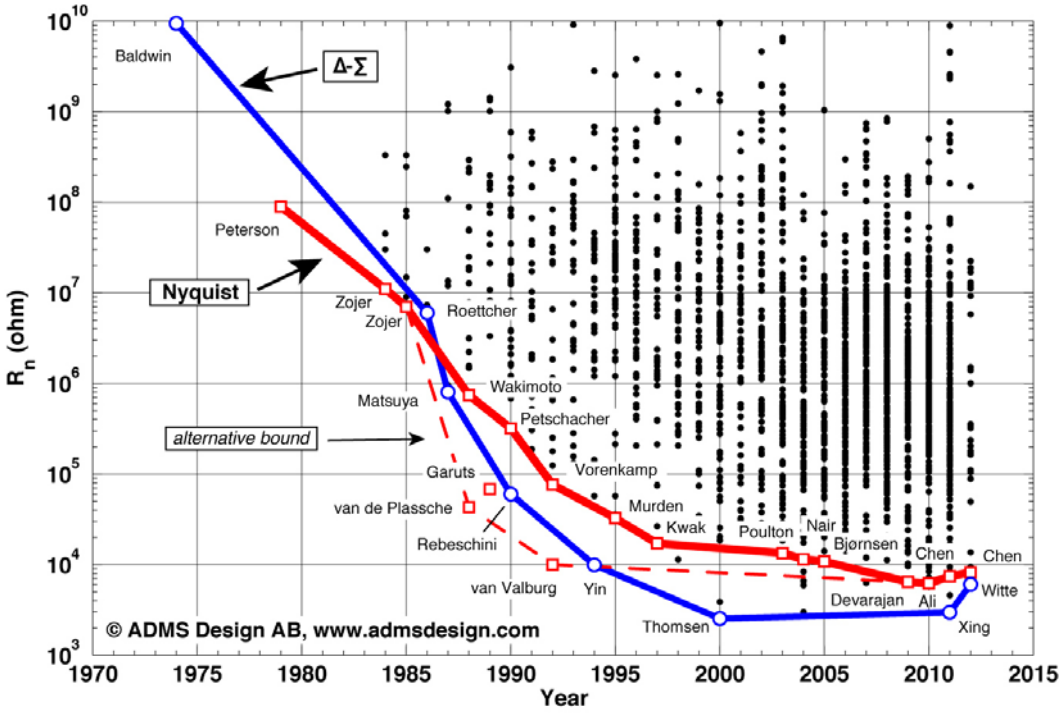

Figure 5.1. Evolution of absolute thermal noise levels as represented by equivalent input-referred noise resistance R_n in scientific publications. The state-of-the-art envelopes for Δ-∑ modulator (○) and Nyquist (□) ADCs have been highlighted.

5.2 Observation of ADC thermal noise trends

The evolution of R_n over time is shown in Figure 5.1 for an assumed T = 300K, and the state-of-the-art envelope for *delta-sigma modulator* (DSM) and Nyquist ADCs are highlighted. As explained above, the value of R_n reflects the *absolute* noise present in the circuit with ideal quantization errors removed. The lowest absolute noise level reported for Δ-∑ modulators to this date is equivalent to $\boldsymbol{R_n \approx 2.5\ k\Omega}$ for the DC measurement ADC in [36] by *Thomsen* et

al., and the best Nyquist converter is the 250 MS/s, IF-sampling, 16-b, pipeline ADC by *Ali* et al. [37] with $\boldsymbol{R_n \approx 6.2\ k\Omega}$.

Note that the estimation of SNR_C in Eq 5.2 and Eq 5.5 becomes sensitive to numerical errors when the reported SNR is close to the ideal quantization noise limit. In order to not overestimate the flatness of the saturation/slowdown region, three designs [38]-[40] were treated as potential (but not certain) outliers, possibly created by numerical problems in the estimation of SNR_C, and therefore the main state-of-the-art envelope was drawn inside of these designs. As an example, a peak ENOB of ~7.97-b is reported for the 8-b folding ADC design by *van Valburg* et al. [38]. It corresponds to an ADC with almost nothing but ideal quantization noise, and hence a design that has extremely low analog noise levels in comparison with its overall spec. The same applies to the designs by *van de Plassche* et al. [39] and *Garuts* et al. [40]. For completeness, the R_n values estimated for these designs have been added to Figure 5.1, together with the alternative state-of-the-art bound resulting from their inclusion.

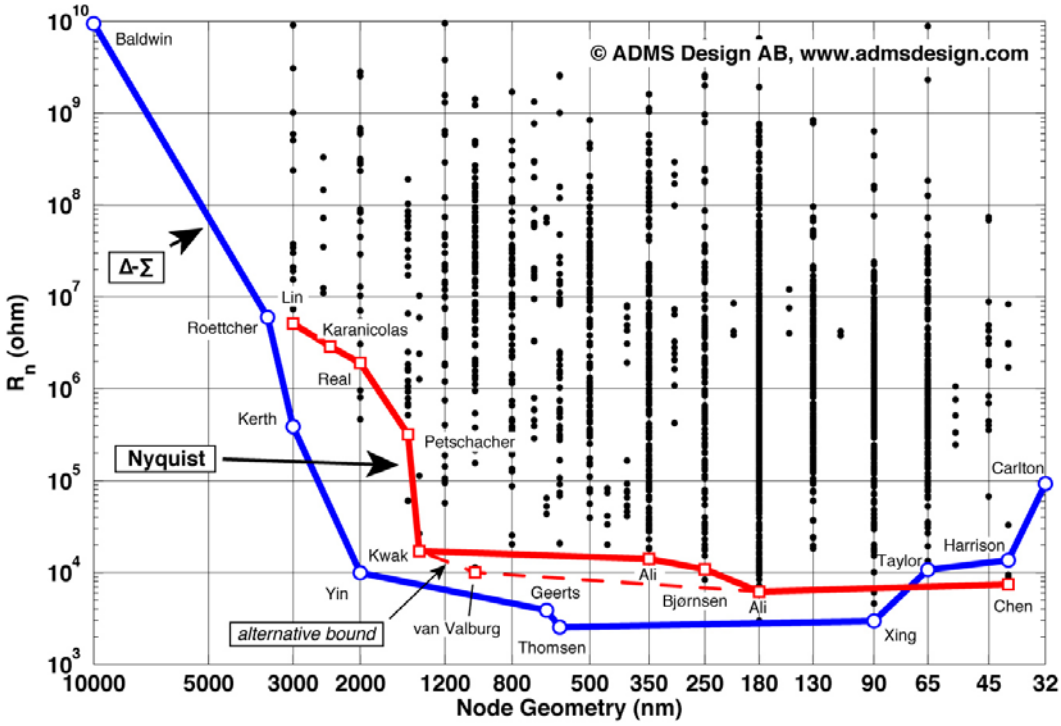

Figure 5.2. Evolution of equivalent input-referred noise resistance R_n over node geometry (any technology) according to scientific publications. The state-of-the-art envelopes for Δ-∑ modulator (○) and Nyquist (□) ADCs have been highlighted.

It is further seen from Figure 5.1 that absolute noise levels have *not improved after year 2000 for delta-sigma modulators*, and only improved by less than 3× since the 17 kΩ achieved in 1997 [41] for Nyquist ADCs. From what is seen in Figure 5.1, it is highly likely that absolute noise power performance is currently in a state of saturation, or significant slowdown, where only moderate improvements are reported over time. Observing how R_n evolves over technology scaling, as shown in Figure 5.2, gives further reason to believe in such saturation. It seems that the state-of-the-art R_n is almost independent of technology node, which is in line with [42], where it is suggested that the absolute noise in devices may

remain constant over scaling. Similar trends of saturation will be shown for other noise-related performance parameters in two subsequent sections.

In the next section we will examine the evolution of A/D-converter *jitter* performance.

5.3 Additional remarks

- Please note that what we believe to be Australia's first[1] – the 20 MHz bandwidth, LC bandpass delta-sigma ADC in 40 nm CMOS by *Harrison* et al. [43] – is also one of the ADCs currently defining the state-of-the-art envelope for R_n vs. node scaling. Now, doesn't that seem like a good way to introduce yourself? I think so.
- The apparent increase in R_n below 90 nm is not necessarily due to physics: There are too few designs below 65 nm to assess these nodes. In particular, the effort by *Carlton* et al. [27] is the only 32 nm ADC so far. The higher R_n in 65 nm can be a temporary setback (as also observed for 130 nm).
- It should be understood that the R_n values in this section are *estimates* of thermal noise under the assumption that the remaining noise is dominated by thermal noise (at low input frequencies) when quantization noise is removed. For the designs with state-of-the-art noise performance, this is a reasonable assumption.
- Survey data used in this work has a near-exhaustive coverage of all scientifically published ADCs, but did not include commercial parts. It is sometimes suggested that the inclusion of commercial parts would paint a completely different picture. At least when it comes to noise-related performance, they don't seem to do that. Although commercial parts may settle to a slightly different (usually better) noise-level, they show a similar slow-down and saturation behavior over time.
- The term "*Node Geometry*" used in Figure 5.2 indicates a more liberal inclusion of all survey data points for which a process technology "feature size" was reported. In contrast, the more discriminating "*Channel Length*" is sometimes used in other sections.

[1] See http://converterpassion.wordpress.com/2012/05/14/australias-first-scientific-adc/

6. Jitter

In the previous section we observed the evolution of absolute thermal noise levels in ADCs. In this second noise-related section out of three, we will look at *sampling time uncertainty*, commonly referred to as *jitter*. The future evolution of A/D-converter jitter performance will have great impact on the development of advanced communications infrastructure, or any other application where you wish to sample at *radio frequencies* (RF) or beyond. A thorough assessment of past and present jitter evolution trends is therefore a highly valuable reference for system-level strategists, as it gives an indication of what kind of ADC performance to expect in the future.

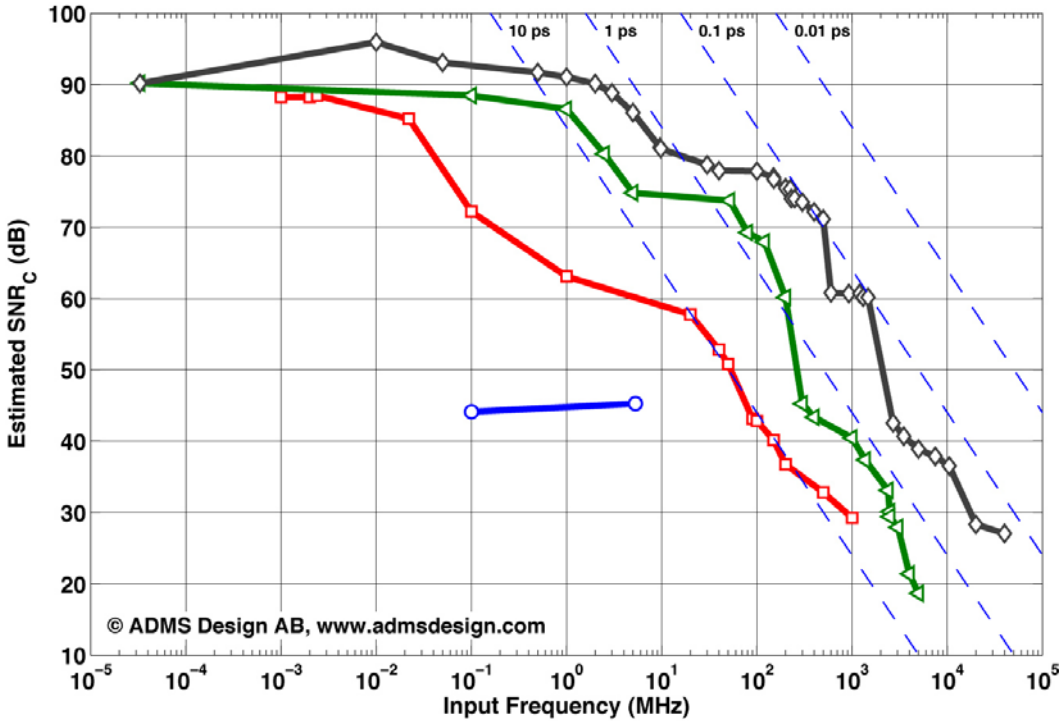

Figure 6.1. Evolution of SNR-vs.-f_{in} envelope for Nyquist ADCs: Current state-of-the art (◇) is compared to state-of-the-art envelopes at 1980 (○), 1990 (□) and 2000 (◁). Theoretical SNR limits for jitter are indicated.

6.1 Observation of ADC jitter trends

Sampling a single-tone input with frequency f_{in} and rms sampling time uncertainty σ_t in an otherwise ideal system yields a jitter-defined SNR in dBc defined by Eq 6.1.

Eq 6.1
$$SNR_J = 20 \times \log_{10}\left(\frac{1}{2\pi f_{in}\sigma_t}\right)$$

An ADC where the circuit noise is dominated by sampling jitter has a *circuit* signal-to-noise ratio $SNR_C \approx SNR_J$. (*SNR_C as defined by* Eq 5.2.) Observing the SNR achieved at a particular input frequency therefore leads to a worst-case estimate of sampling jitter, and looking at the SNR_C vs. f_{in} progress for the entire body of scientific Nyquist ADC data thus renders a conservative estimate of *jitter performance evolution over time*. Figure 6.1 shows the state-of-the-art envelopes for SNR_C vs. f_{in} at 1980, 1990, and 2000 compared to present

day (~Q1 2012). By using SNR_C instead of SNR, the ideal quantization noise component that would falsely add to the jitter estimate is removed. For data points where the effect of jitter has fully kicked in, this actually makes little difference. Since not all papers report performance at such high input frequencies, the use of SNR_C instead of SNR still gives a better jitter estimation.

An SNR_J vs. f_{in} jitter raster according to Eq 6.1 has been included as a visual guide in Figure 6.1 for jitter values of 0.01, 0.1, 1, and 10 ps. The state-of-the-art in 1980 is defined by only a few designs and therefore difficult to interpret. By 1990, a much larger number of attempts have been made, and the high-frequency roll-off is almost perfectly aligned with the theoretical SNR limit for 8–10 ps jitter. *Garuts* et al. achieves 5.5 ps at 1 GHz input with a 4-b, 1 GS/s ADC (9.8 ps without quantization error subtraction) [44]. During the following decade, the lowest jitter estimate was reduced by almost exactly ten times to the 0.53 ps reported by *Singer* et al. for an IF-sampling 12-b, 65 MS/s ADC in [45]. A slight slowdown is then observed for the last 11 years, when the jitter evolution front progress by less than a factor of ten across most of the frequency range. Current state-of-the-art is defined by the 14□b, 125 MS/s, RF/IF-sampling ADC reported by *Ali* [46]. The 88 fs rms jitter estimated at 500 MHz f_{in} is a six times improvement over [45]. Looking at Figure 6.1, it is evident where the main effort was focused. Whereas the state-of-the-art boundary is almost straight at 1990, there is an obvious “bump” in the 50–200 MHz range by 2000, which has migrated to 100–500 MHz to this date, and a secondary bump has emerged recently at 1–2GHz, defined by data reported in [47]. Although there are other applications matching these frequency ranges, the progress of the evolution front aligns very well with the evolution of requirements for wideband communications infrastructure. The current state-of-the art boundary is largely defined by publications stating wireless or digital communication as the target application. Hence it is concluded that communications applications have been a main driver for jitter performance over the last two decades. **A point of concern for the communications industry** could be that the current state-of-the-art [46] was achieved in 2005, using a 0.35 µm BiCMOS process. Observing the state-of-the-art estimate of σ_t (at any f_{in}) over time, as shown in Figure 6.2, jitter performance appears to be in a state of saturation. A closer inspection of the underlying data set gives further reason to assume saturation of jitter performance: Figure 6.3 shows the evolution of jitter performance vs. *technology scaling*. Although it’s possible that jitter eventually starts to improve in deeply scaled nodes, the current trend clearly supports the assumption of **jitter performance saturation**.

In the third and final section on noise performance evolution, we will study the evolution of A/D-converter *relative noise floor*.

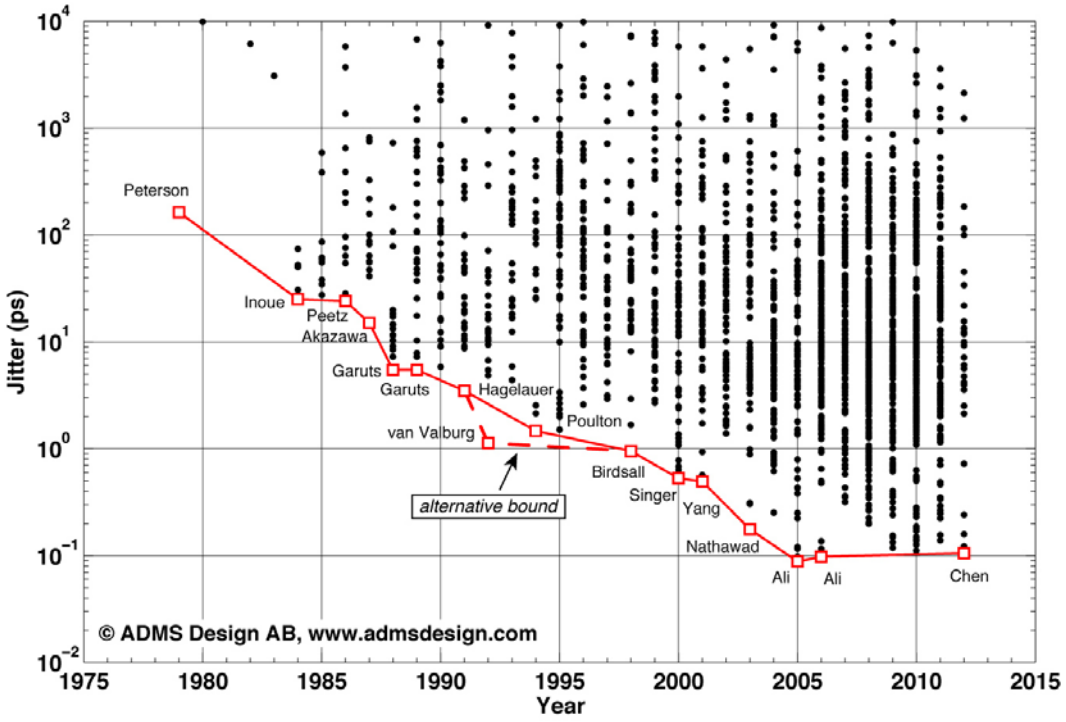

Figure 6.2. Evolution of worst-case jitter estimate for Nyquist ADCs in scientific publications. The state-of-the-art envelope has been highlighted.

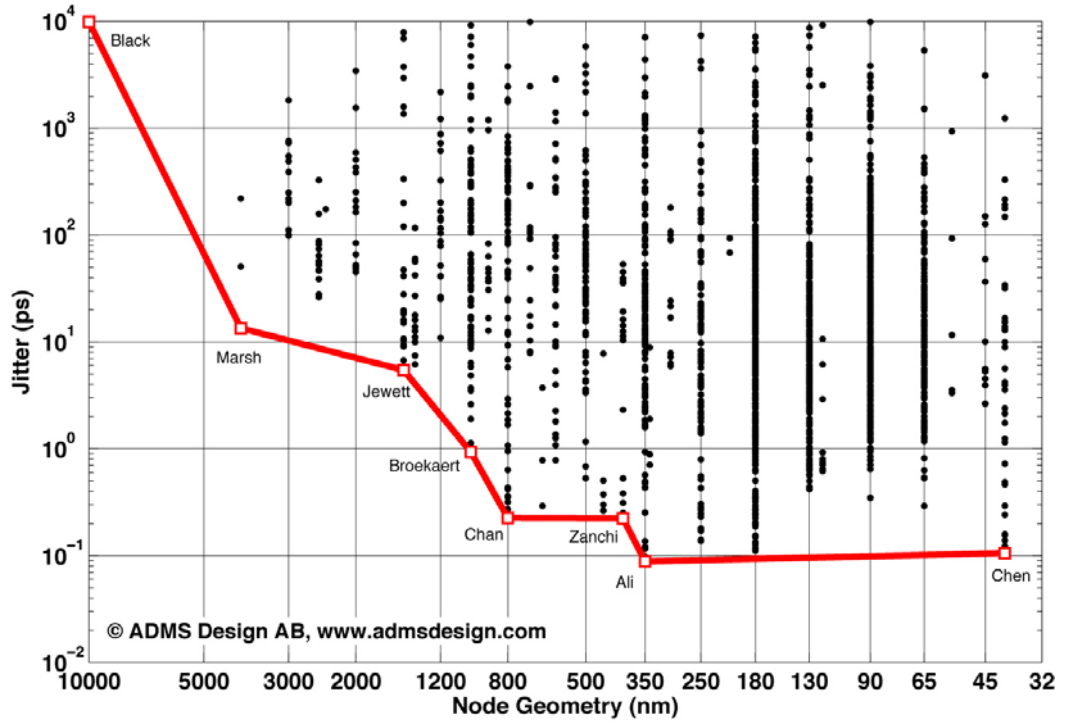

Figure 6.3. Evolution of worst-case jitter estimate for scientific Nyquist ADCs over node geometry (any technology). The state-of-the-art envelope has been highlighted.

6.2 Additional remarks

- Note that jitter was analyzed by observing Nyquist ADCs. Bandpass delta-sigma modulators can achieve much better combinations of SNR and input frequency. The current state-of-the-art is a continuous-time bandpass delta-sigma modulator by *Luh* [48]. A significant part of the sampling jitter in a delta-sigma modulator is suppressed either by the noise transfer function or when selecting a band-of-interest in the subsequent signal processing. An arbitrarily small jitter estimate could be achieved simply by choosing to measure the SNR over a lower bandwidth. Delta-sigma, and other narrowband ADCs were therefore excluded from the jitter observation, but are included in both the previous and the subsequent noise-evolution sections.
- Designs that could influence the state-of-the-art envelope, but possibly suffered from numerical problems in the SNR_C estimation (for reasons described in section 5), were handled as follows: (a) All data points for which $N - \text{ENOB} < 0.05$ were filtered out before generating Figure 6.1, and (b) The design by *van Valburg* [38] was left outside the main estimation of state-of-the-art envelope in Figure 6.2. No special handling was necessary for Figure 6.3.
- ADCs using some form of optical or optoelectronic solution may be able to sample with considerably lower timing uncertainty. As of yet, such ADCs are quite rare, and mostly implemented in unusual or purely experimental technology. Optoelectronic or all-optical solutions may very well be the way forward if classic electronic sampling saturates at unacceptable jitter levels. The survey here, however, did not cover optically sampled A/D-converters.

7. Relative noise floor

We have in the previous sections studied the evolution of absolute thermal noise levels, and sampling jitter for analog-to-digital converters (ADC). Finally, the *overall noise performance evolution* is observed with all noise contributors included. Whereas the two previous sections analyzed two fundamental noise components in isolation, this section looks at the actual noise performance achieved with everything included. This is the ADC performance you actually get.

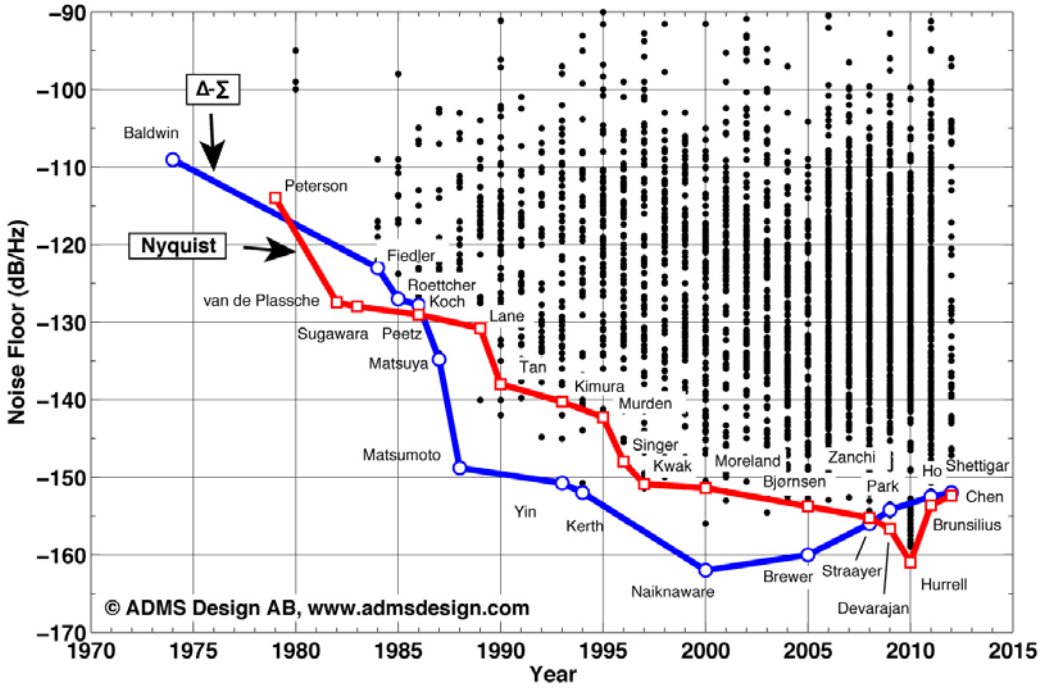

Figure 7.1. Evolution of relative noise-floor for DSM (○) and Nyquist (□) ADCs over time.

7.1 Observation of ADC noise floor trends

The ADC survey data spans a very wide range of converter specifications. An SNR of *x* dB in 20 kHz bandwidth is not as impressive as achieving the same in a 1GHz band. Using the relative noise-floor n_r in dB/Hz derived by Eq 7.1 allow ADCs with widely different Nyquist bandwidths (BW) to be compared with respect to noise performance.

Eq 7.1
$$n_r = -\left(SNR + 10 \times \log_{10} BW\right)$$

Figure 7.1 shows the evolution of n_r for *delta-sigma modulators* (DSM) and Nyquist ADCs over time. A similar plot, based on less data, and not differentiating between DSM and Nyquist ADCs is found in [18]. From a linear fit of the state-of-the-art data points, it is seen here that ADC noise-floor for has evolved at an average rate of ~2.2 dB/year until year 2000 for DSM, after which it has remained in saturation. Nyquist ADCs have developed at a slower rate of ~1.3 dB/year until 2010. The current state-of-the-art is approximately the same for both: **–162 dB/Hz** for **DSM** [34], and **–161 dB/Hz** for **Nyquist** [49]. Since the state-of-the-art for Nyquist converters was so recently reported, it cannot be concluded only from Figure 7.1 that the noise floor for Nyquist ADCs is in saturation. A likely explanation for the DSM trend is, however, the lower signal swing, and thus higher relative noise-floor, implied by the

continuous scaling of semiconductor technology [50], [17]. This is also seen in Figure 7.2. Although the *absolute* noise-floor may remain constant in devices [42], the *relative* noise-floor is raised when signal swing is reduced. New technologies may allow higher bandwidths, but the simultaneous combination of SNR and bandwidth has not improved for a decade due to this inherent dynamic-range limitation of nanometer technology [17]. It is likely to assume that Nyquist converters will suffer from this limit *at least* as much as delta-sigma modulators do. Noise-performance normalized to signal bandwidth therefore seems to have reached the physical limits of process technology defined by the available signal swing. Expecting a further reduction in signal swing [26], future ADCs could very well fail to maintain the current state-of-the-art in noise performance.

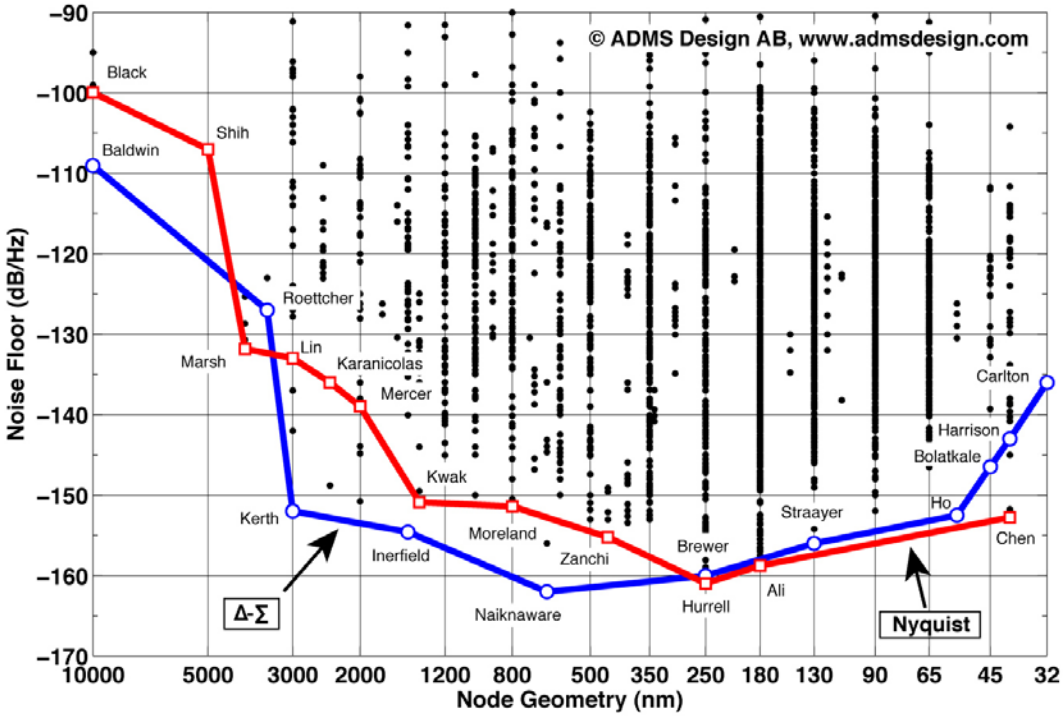

Figure 7.2. Evolution of relative noise-floor for DSM (○) and Nyquist (□) ADCs vs. node geometry (any technology).

7.2 Conclusion: ADC noise performance trends

Over the last three sections, it was seen that the overall state-of-the-art with respect to *absolute noise power*, *sampling jitter*, **and** *relative total noise floor* has not improved during the last 5–10 years. It is therefore concluded that *all significant aspects of ADC noise performance appear to have reached saturation*. This is an expected, yet significant result of the study as it clearly confirms the commonly raised concerns regarding analog design and dynamic range in scaled technologies, e.g., in [15], [17], [20], [42], [50].

My conclusion is that A/D-converters have already hit the noise floor – at least its softer upper coating.

7.3 Additional remarks

- As commented in section 5, the trends of degradation observed below 65 nm in Figure 7.2 may be due to lack of reported attempts, and not necessarily due to physics.
- Using a circuit design that allows for large input signal swing can help a lot in improving relative noise floor performance. As an example, the state-of-the-art design by *Hurrell* et al. [49] reports an **8.2 V** peak-to-peak input full-scale range.

8. Linearity (SFDR)

After observing trends for technology scaling, voltage scaling, and noise, we have arrived at *linearity*. As far as the author is aware, there has been no large survey on A/D-converter linearity evolution published to this date, except for a recent work by *Walden* [12], where an "*SFDR-bits*" vs. f_{in} scatter plot illustrates the state-of-the-art movement between 1999 and 2007 for the outermost corner/edge of the data, based on approximately **175** ADCs. In this work we will instead observe the migration of the entire *envelope* (state-of-the-art) for SFDR vs. input frequency (and sampling rate) in order to observe how SFDR evolved across multiple frequency ranges between 1990, 2000 and present (~Q1-2012). We will also slice through the data set, and specifically observe SFDR evolution at four very different speed grades of minimum sampling rate. The underlying data is from a survey of **1708** scientific ADC papers published between 1974 and Q1-2012.

8.1 ADC linearity trends: SFDR-vs-frequency envelope

While SNR, or an aggregate noise-and-distortion measure such as SNDR, is a sufficient measure of ADC performance for some applications, there are other applications where non-linear distortion is independently specified. Such applications include high-end audio and many wireless communication systems. Wireless communication systems often need to cope with the presence of a strong interferer in the form of a neighboring channel or carrier, while correctly interpreting a weak signal of interest. Without a sufficiently linear signal path, the interfering signal will generate harmonics or intermodulation products that may completely block the in-channel signal. The evolution of ADC linearity is therefore as important as the noise and ENOB evolution.

State-of-the-art envelopes for single-tone *spurious-free dynamic range* (SFDR) vs. input frequency f_{in} and Nyquist sampling rate f_s have been plotted in Figure 8.1 and Figure 8.2 respectively. Current state-of-the-art at ~Q1-2012 is compared to that of 1990 and 2000 in order to illustrate the evolution over all frequencies. Starting with the envelopes at 1990, the linearity vs. f_{in} and f_s is evenly distributed across almost straight lines representing the increasing difficulty to achieve high linearity as the input frequency and sampling rate is increased. The first of the two noticeable performance peaks in Figure 8.1 coincide with the 20 kHz audio bandwidth, and the second peak is defined by video and instrumentation ADCs in the frequency range 10-100 MHz. It is evident from both plots that most of the progress from 1990 to the current state-of-the-art was achieved in the first decade 1990-2000 when SFDR vs. f_{in} was improved by 20-40 dB across all input frequencies in the 100 kHz to 1GHz range, and SFDR vs. f_s increased by 5-30 dB for the same range of sampling frequencies. Although state-of-the-art linearity has been increased by 5-10 dB over many segments of the frequency range during the last 11 years, Figure 8.1 and Figure 8.2 clearly show that there has been a slowdown in the evolution of linearity over a broad range of frequencies and speed grades. One noticeable exception is the 30 dB performance lift in the 100-250 MS/s speed range, which reflects the specifications of the more recent wideband *radio base-stations* (RBS). It was concluded in section 6, "*Jitter*", that the evolution of communications standard requirements has been a strong driver for ADC jitter performance. Observing that the strongest push of the *linearity* envelopes also occurred at frequencies and sampling rates matching the specifications for wideband RBS, e.g., [37], [51]-[53], **it is concluded that communications applications have been a key driver for ADC linearity as well**. This is also the conclusion of *Walden* in [12]. Another significant achievement during the last decade

has been to improve performance at the high-frequency end of the spectrum, with sampling rates above 4 GS/s and input frequencies beyond 1 GHz [54]-[62]. Again, a similar observation was made in [12].

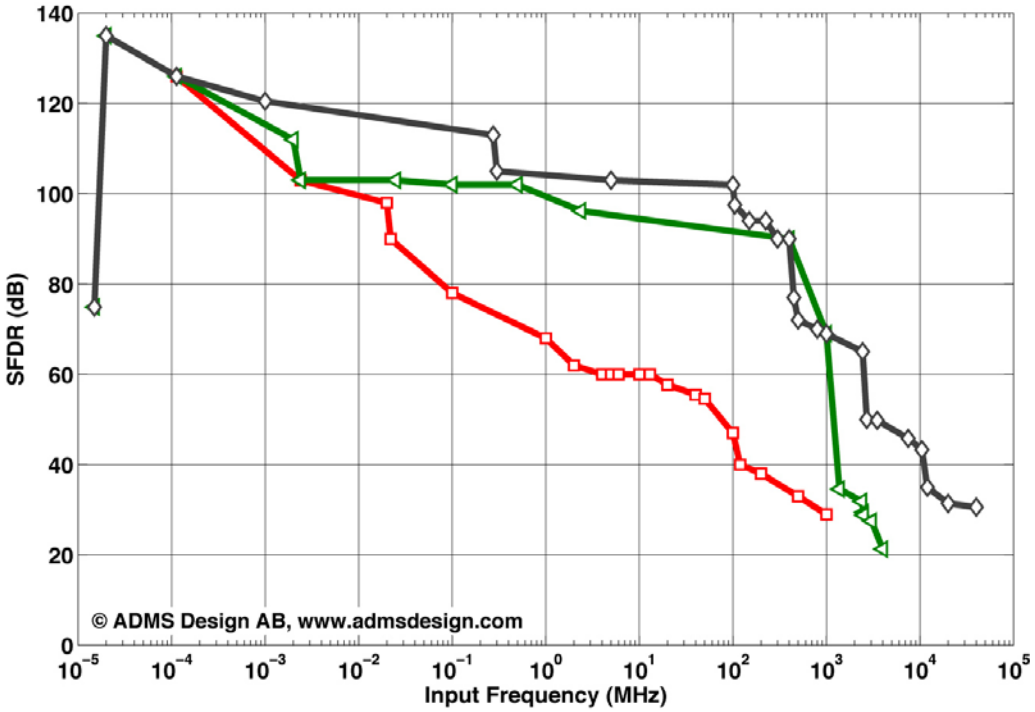

Figure 8.1. Evolution of SFDR-vs.-f_{in} envelope for scientifically reported ADCs: Current state-of-the art (◇) is compared to state-of-the-art envelopes at 1990 (□) and 2000 (◁).

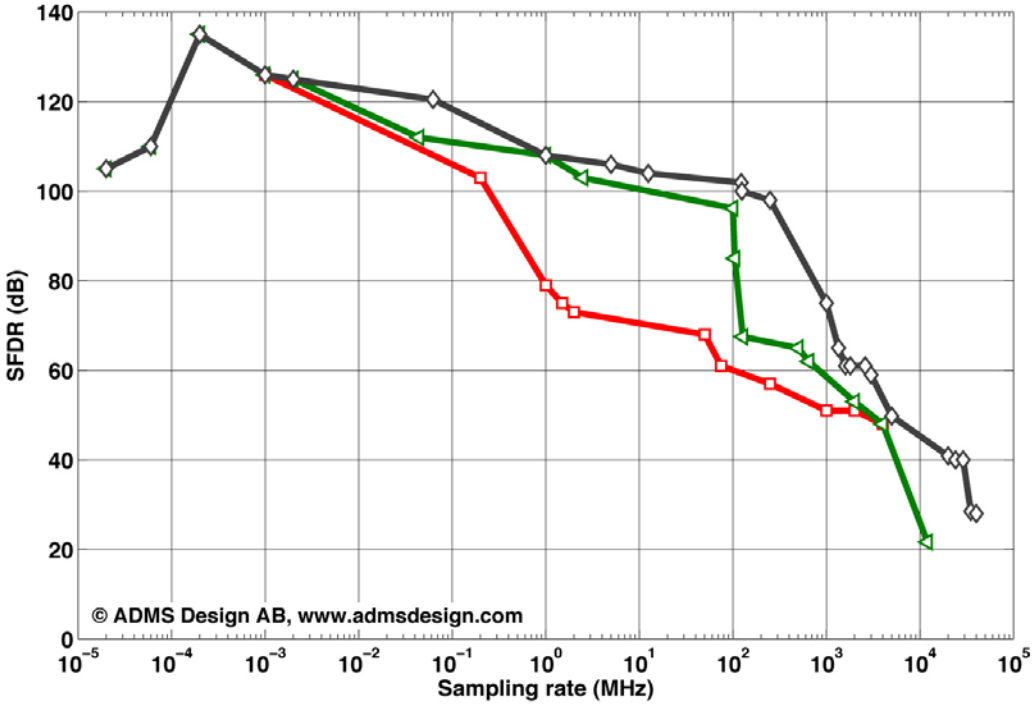

Figure 8.2. Evolution of SFDR-vs.-fs envelope for scientifically reported ADCs: Current state-of-the art (◇) is compared to state-of-the-art envelopes at 1990 (□) and 2000 (◁).

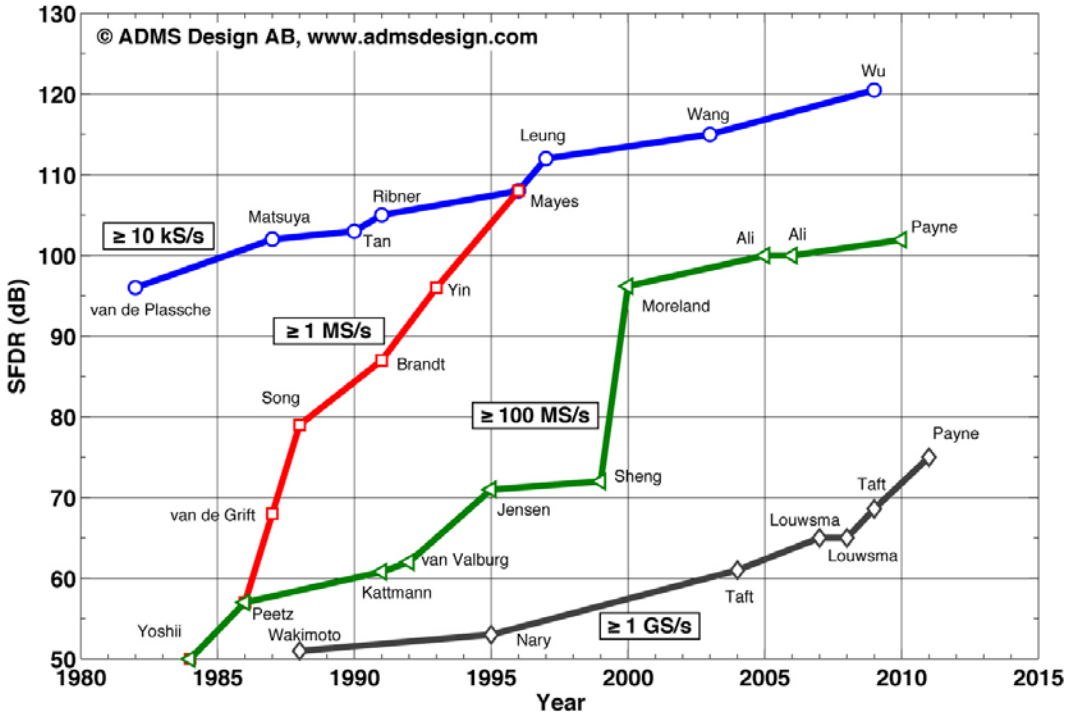

Figure 8.3. Scientifically reported ADC implementations: Peak SFDR evolution over time for minimum sampling rates of 10k (○), 1M (□), 100M (◁), and 1GS/s (◇).

8.2 ADC linearity trends: SFDR by speed grade

Figure 8.3 shows the evolution of peak SFDR at minimum speed grades of $f_s \geq \{10k, 1M, 100M, 1G\}$ samples/s. The curves show *the monotonically improving upper edge* for each subset of survey data. What is included in each subset is defined by the four minimum sampling rate constraints. As in Figure 8.1 and Figure 8.2, the overall scatter is removed for readability.

My interpretation of Figure 8.3 is as follows:

- 1MS/s ADCs appear to have saturated at an SFDR of 108 dB [63], and have not improved since 1996.
- 100 MS/s ADCs has slowed-down evolution beyond 100 dB to ~6 dB/decade or less. Current state-of-the-art is 102 dB SFDR [53].
- At very low, and very high sampling rates there are no signs of saturation yet. For ADCs with $f_s \geq 10$ kS/s, there has been an almost constant progress of ~9dB/decade from 1992 [64] to 2009 [65].
- ADCs with $f_s \geq 1$ GS/s are currently evolving at an *accelerated* rate of ~3 dB/year. **If** this rate is maintained, gigasample ADCs could go from 75 dB [66] to upwards of 100 dB SFDR by 2020.
- The accelerated rate of evolution seen at different times for different speed grades may reflect how research activities migrate to higher and higher sampling rates depending on what applications are in focus. Previously more of a niche product, gigasample ADCs are now becoming a mainstream necessity.

8.3 Commercial ADC parts

Although not shown here, the results were also compared with the data from **595** commercially released ADC parts. Current state-of-the-art envelope for both sets align well across most of the speed range, with one significant exception: There are already commercial parts with significantly better SFDR than their scientific counterparts at 2 MSPS and below, e.g., AD7766 [67] and AD7986 [68]. Commercial ADCs appear to have evolved beyond their experimental siblings in later years within this speed segment.

Another difference is in the paths each subset has followed towards today's (mostly similar) state-of-the-art. In the GSPS range there is for example MAX 104 [69], specifying 69 dB SFDR at f_{in} = 125 MHz almost a decade before the scientific publication by *Taft* [70], while scientific efforts seem to have been ahead in other frequency ranges (e.g., below 2MS/s) during earlier years[2].

Linearity trends are therefore more difficult to interpret, and to some degree depend on what products were reported scientifically and not. Such dependency could not be observed for any noise-related parameter analyzed in this work.

[2] Please note that these statements on commercial ADCs are only my best guesstimates, as the commercial data set (although large) is not as exhaustive as that for scientific ADCs.

9. Sampling rate and resolution

Previous sections analyzed noise and linearity separately. Another common approach is to review the overall ADC performance in terms of sampling rate and effective resolution ENOB. In Figure 9.1, the current state-of-the-art at ~Q1-2012 is compared to the envelopes for 1990 and 2000 in order to show the simultaneous evolution of the two parameters throughout the entire parameter space. *SNR-only results have been excluded* from this plot because ENOB is not fully defined by SNR. Hence, there is no experimental data available before 1980. By 1990 the curve has assumed the expected shape. Between 1990 and 2000 there is a 1-4 bits improvement across the full range of sampling rates. The main advances were in the 200kS/s – 100MS/s speed range. This corresponds to typical telecommunications specifications – from single-carrier GSM to multi-carrier WCDMA receivers. From year 2000 to present day, the more significant advances were at 12.5 MS/s [49], from 100–250 MS/s, [37], [71], at 3 GS/s [72], and above 10 GS/s [54], [60], [62].

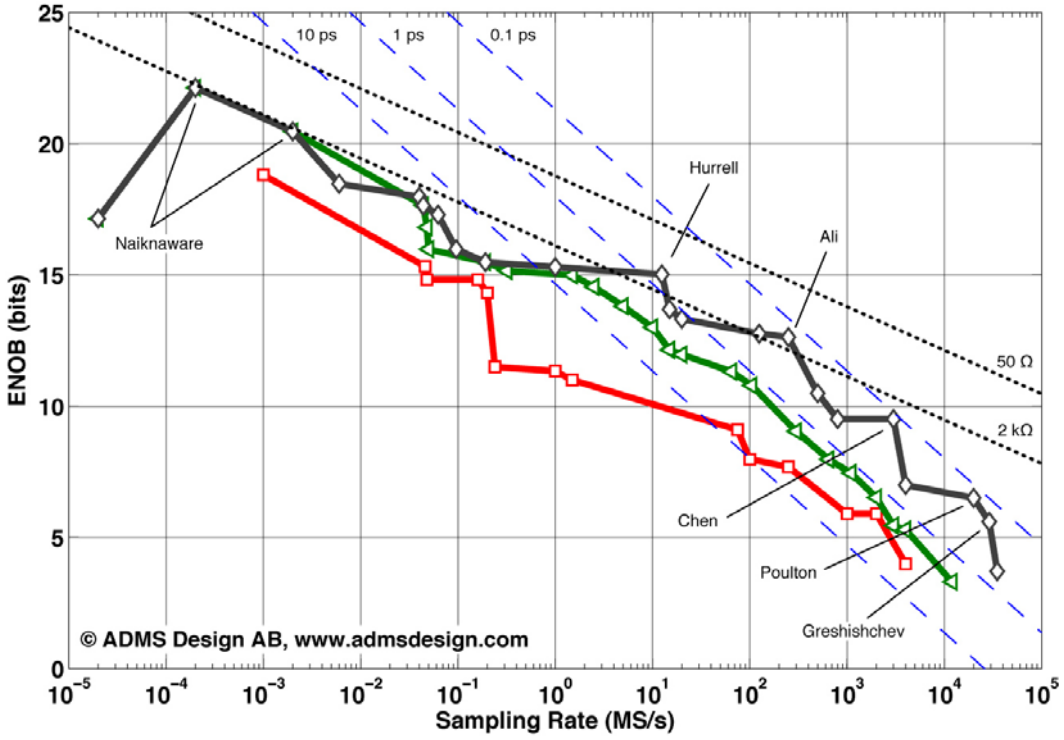

Figure 9.1. Evolution of ENOB vs. f_s envelope for scientifically reported ADCs. Current state-of-the art (◇) is compared to state-of-the-art envelopes at 1990 (□) and 2000 (◁). Theoretical limits for thermal noise @ V_{FS} = 1V (dotted) and jitter (dashed) are indicated.

The thermal noise limits according to Eq 5.4 have been included as a visual guide, using V_{FS} = 1V, T = 300 K, and R_n = {50, 2000} Ω. Similarly, the theoretical jitter-limited ENOB at $f_{in} = f_s/2$ according to Eq 6.1 has been added for σ_t = {0.1, 1, 10} ps. The R_n and σ_t values were deliberately chosen to simplify comparison with a similar plot in *Walden's* survey [12] (see also *Additional remarks* below). Although the jitter limits should preferably be observed from SNR vs. f_{in} (as done in section 6, "*Jitter*"), the shape of the state-of-the-art envelopes in Figure 9.1 clearly indicate the regions where ADC performance is limited by thermal noise and jitter respectively. The design by *Naiknaware* et al. [34] is limited by thermal noise, while

those by *Poulton* et al. [54] and *Greshishchev* et al. [62] are limited by sampling jitter (and/or metastability [10]). At the boundary between thermal noise and jitter limited designs are the ADCs that suffer from both noise sources in equal amount, such as the design by *Ali* et al. [37]. Designs in this corner put strict demands on the simultaneous design for jitter *and* thermal noise.

9.1 Additional remarks

- It may seem that the state-of-the-art thermal noise according to Figure 9.1 is equivalent to less than 2 kΩ for some designs. This would obviously be in contradiction to the **2.5** and **6.2 kΩ** state-of-the-art reported for *delta-sigma modulator* and *Nyquist* ADCs, respectively. The thermal noise limits in Figure 9.1 are only valid for V_{FS} = 1 V_{pp}, and the apparently better results here are because of a larger full-scale range, e.g., 2.5 V for [37]. The correct noise-resistance estimations are found in section 5, "*Thermal noise*".
- The corresponding jitter limits in [12] have a 0.5-bit offset because it appears that *Walden* derives the rms-signal to *peak*-noise ratio by assuming that the signal is always sampled where the slope is greatest, i.e., in the zero-crossings [10]. In reality, the signal is sampled anywhere along the waveform for all but pathological cases, and therefore the rms slope should be used instead, as was done in this treatment.
- In [18], the evolution trends for peak sampling rate at fixed minimum ENOB grades {4, 8, 12, 14} bits, and the complementary peak ENOB at fixed minimum sampling rates {10k, 100k, 1M, 100M, 1G} S/s are shown in a style similar to Figure 8.3.

10. Walden figure-of-merit (FOM)

A work on A/D-converter performance trends would not be complete without an analysis of *figure-of-merit* (FOM) trends, would it? We will therefore take a look at the two most commonly used FOM, starting with the by far most popular:

Eq 10.1
$$F_{A1} = \frac{P}{2^{ENOB} \times f_s}$$

where P is the *power dissipation*, f_s is Nyquist *sampling rate*, and ENOB is the *effective number of bits* defined by the *signal-to-noise and-distortion ratio* (SNDR) as:

Eq 10.2
$$ENOB = \frac{SNDR - 1.76}{6.02}$$

F_{A1} is sometimes referred to as the *Walden* or ISSCC FOM and relates the ADC power dissipation to its performance, represented by sampling rate and conversion error *amplitude*. The best reported F_{A1} value each year has been plotted for *delta-sigma modulators* (DSM) and Nyquist ADCs in Figure 10.1. Trajectories for state-of-the-art have been indicated, and trends have been fitted to these state-of-the-art data points. The average improvement trend for *all* ADCs (**2×/2.6 years**) is included for comparison.

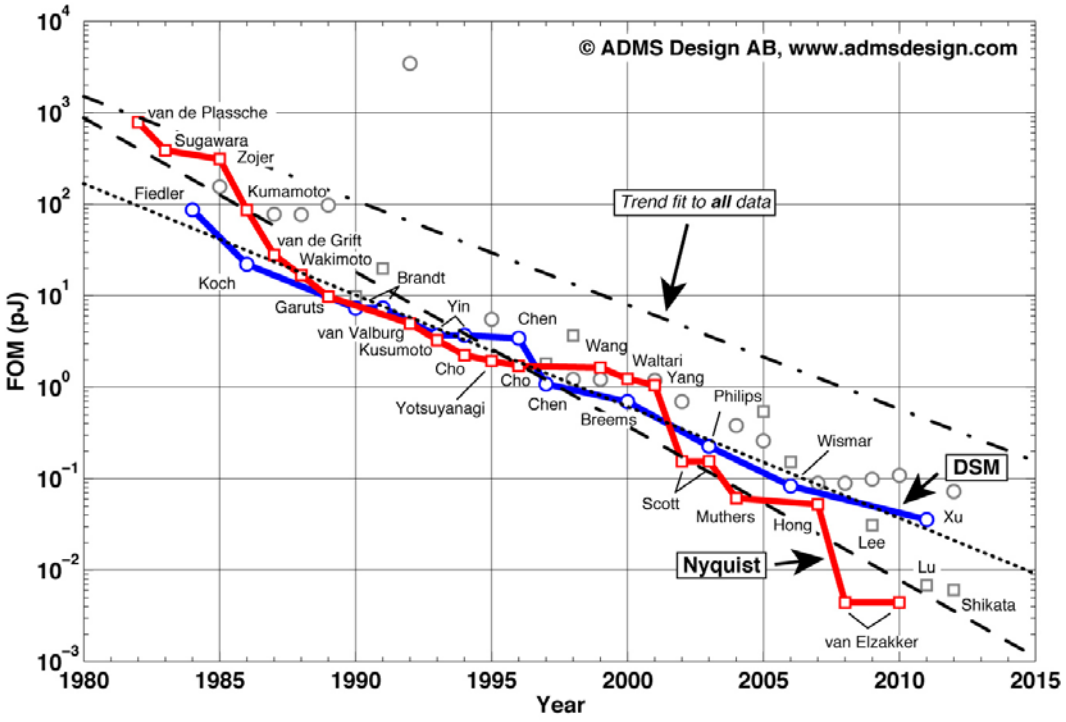

Figure 10.1. Evolution of best reported Walden FOM for delta-sigma modulators (○) and Nyquist ADCs (□). Monotonic state-of-the-art improvement trajectories have been highlighted. Trend fit to DSM (dotted), and Nyquist (dashed) state-of-the-art. Average trend for all designs (dash-dotted) included for comparison.

By dividing the data into DSM and Nyquist subsets, it is seen that delta-sigma modulators have improved their state-of-the-art FOM at an almost constant rate of **2×/2.5 years** throughout the existence of the field – just slightly faster than the overall average. State-of-the-art Nyquist ADCs have followed a steeper and more S-shaped evolution path. Their overall trend fits to a **2× improvement every 1.8 years**, although it is obvious that evolution rates have changed significantly over time. A more accurate analysis of Nyquist ADC trends should probably make individual fits of the early days glory, the intermediate slowdown, and the recent acceleration phase. This was done in [18] where evolution was analyzed with DSM and Nyquist data merged. However, for simplicity I'll just stick to the more conservative overall Nyquist trend. [I wouldn't want anyone to suggest that I'm producing "*subjective*" or "*highly speculative*" trend estimates, would I? ☺]

Still, if anyone *is* curious to know, the state-of-the-art data points fit to a **2×/14 months** trend between 2000 and 2010. That's actually ***faster* than *Moore's Law***, which is traditionally attributed a 2×/18 months rate [8]-[9]. A new twist on "*More than Moore*", perhaps? Even the more conservative overall 2×/21 months trend is close enough to conclude that the state-of-the-art FOM for Nyquist ADCs has developed exponentially in a fashion closely resembling *Moore's Law*. And that's *got* to be an impressive trend for any analog/mixed circuit performance parameter.

Irrespective of what's the best fit to data, it should be evident from Figure 10.1 that Nyquist ADCs broke away from the overall trend around year 2000, and has since followed a steeper descent in their figures-of-merit. They have also reached further (**4.4 fJ**) [73] than DSM (**35.6 fJ**) [74]. The overall trend projects to a **0.2 fJ** ADC FOM in **2020**. Whether or not that's possible, we'll leave for another time. A deeper look at the data also reveals that:

- The acceleration in state-of-the-art is almost completely defined by *successive-approximation* (SAR) ADCs [73], [75]-[80], accompanied by a single *cyclic* ADC [81]. The superior energy efficiency of the SAR architecture was empirically shown in [24].
- A significant part of the acceleration can be explained by the increased tendency to leave out, for example I/O power dissipation when reporting experimental results – a trend also observed by *Bult* [82]. The FOM in the graph was intentionally calculated from the *on-chip* rather than *total* power dissipation because: (a) ADCs are increasingly used as a system-on-chip (SoC) building block, which makes the stand-alone I/O power for a prototype irrelevant, and (b) Many authors don't even report the I/O power anymore.
- F_{A1} has a bias towards low-power, medium resolution designs rather than high-resolution, and thus benefits from CMOS technology scaling as shown in [17], [25]. An analysis of the underlying data shows that, for the best F_{A1} every year, the trajectories for ENOB and P follow distinct paths towards consistently lower power and medium resolution [83]. You simply gain more in F_{A1} by lowering power dissipation than by increasing resolution because Eq 10.1 does not correctly describe the empirically observed power-resolution tradeoff for ADCs [17], [25].

In order to compare high-resolution ADCs limited by thermal noise, it has therefore been proposed to use a slightly different FOM, sometimes labeled the "*Thermal FOM*" [52], [84]. This figure-of-merit will be the topic of the next section.

11. Thermal figure-of-merit (FOM)

As mentioned in the previous section, a slightly different FOM, sometimes labeled the "*Thermal FOM*" [52], [84], has been proposed in order to better compare high-resolution ADCs limited by thermal noise. The thermal FOM, F_{B1}, is expressed as

Eq 11.1
$$F_{B1} = \frac{P}{2^{2\times ENOB} \times f_s}$$

The thermal FOM considers error *power* rather than *amplitude* (as in the Walden FOM), and therefore the value of F_{B1} improves by 4× (rather than 2×) for every additional bit of resolution. This matches the theoretical 4× minimum increase in power if ENOB is limited by *kT*/*C*-noise [15] and the architecture remains unchanged [82]. It was shown in [25] that the thermal FOM represents a better description of the state-of-the-art power-resolution tradeoffs according to empirical data than the Walden FOM for ENOB ≥ 9. As seen in Figure 11.1, there is a significant difference between DSM and Nyquist ADCs with respect to F_{B1}. With the exception of two early 14-b designs [64], [85], *the global state-of-the-art is defined entirely by delta-sigma modulator implementations* while Nyquist ADCs lag distinctly behind. A possible explanation could be that the thermal FOM favors converters whose power dissipation is truly limited by thermal noise, and that high-resolution Δ-∑ ADCs are more distinctly driven into the thermal noise limit than their Nyquist counterparts. Another point is that many scientific DSM implementations use an off-chip (i.e., zero power) decimation filter implemented in software. This will give DSM an unfair advantage over Nyquist, although it can hardly be the only explanation for a one order of magnitude FOM difference.

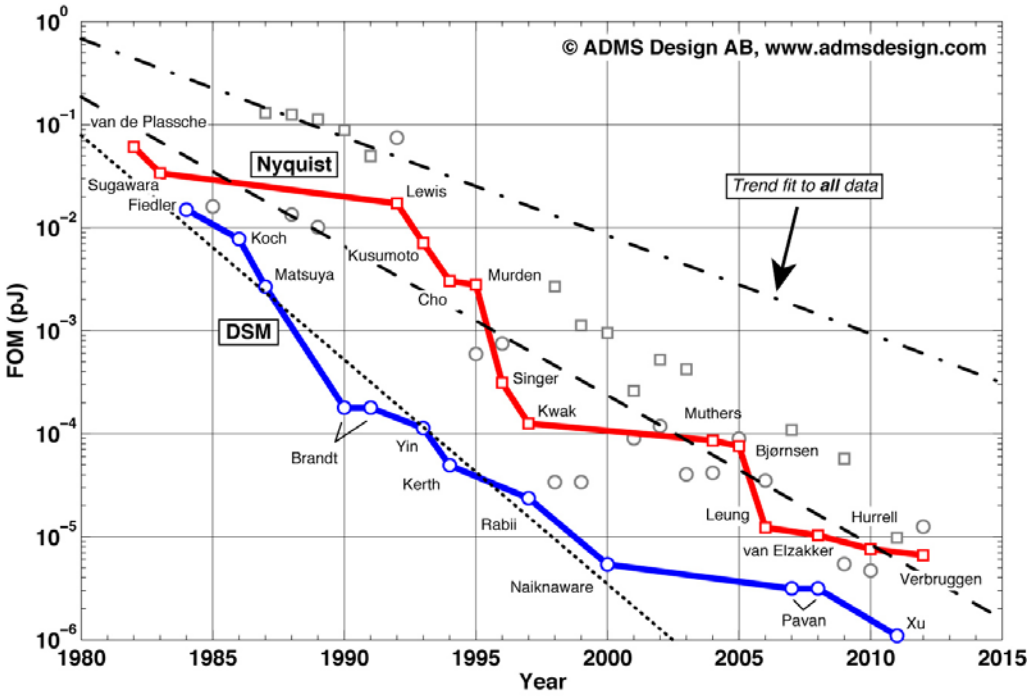

Figure 11.1. Evolution of best reported thermal FOM for delta-sigma modulators (○) and Nyquist ADCs (□). Monotonic state-of-the-art improvement trajectories have been highlighted. Trend fit to state-of-the-art points for DSM [1984–2000] (dotted), and Nyquist [1982–2012] (dashed). Average trend for all designs (dash-dotted) included for comparison.

Since the thermal FOM for Nyquist converters has evolved over a rather uneven path, I'll not make any elaborate interpretations of its shape. The trend (dashed) is simply fitted to all the state-of-the-art points from 1982–2012, revealing an average improvement rate of 2**×** every two years. The DSM envelope appears to have three main segments with breakpoints at 1990 and 2000, respectively. For simplicity, a single trend was estimated for the envelope until *Naiknaware* [34], after which the thermal FOM has evolved significantly slower. From *Fiedler* [86] to *Naiknaware*, the average improvement rate is **2**× **every 17 months** (1.4 years) – again faster than *Moore's Law* [8]-[9] – whereas from year 2000 to present day [74], the state-of-the-art points fit to a more modest **2**×**/5.5 years** slope. Even if the latter is from a fit of only four data points, and the exact slopes can be discussed, it is clear from Figure 11.1 that the thermal FOM for DSM experienced a distinct slowdown after year 2000. This coincides with the breakpoint where the relative noise floor – approximately the denominator in Eq 11.1 – *also* goes into saturation[3]. It can further be noticed that it coincides with the accelerated evolution of F_{A1} as well[4]. A possible, but perhaps speculative interpretation is that the ADC community first focused on thermal noise performance and related design optimization, and after hitting the noise floor around year 2000 moved on to focus on power efficiency.

As a small postlude, a list of known prior art ADC surveys are included in the appendix for those of you that (like myself) have an insatiable appetite for technology trend estimations and empirical data dots.

[3] See section 7, "*Relative noise floor*".

[4] See section 10, "*Walden figure-of-merit (FOM)*".

12. Appendix: A survey of ADC surveys

If the A/D-converter survey just concluded did not fully satisfy your desire for scatter plots and tech trends, then this appendix will provide a list of prior ADC survey works as suggested "further reading". In fact, I'd recommend everyone serious about data-converter technology trends to get hold of these documents. The list will also serve as a brief history of the ADC survey field. But first some thoughts on surveys:

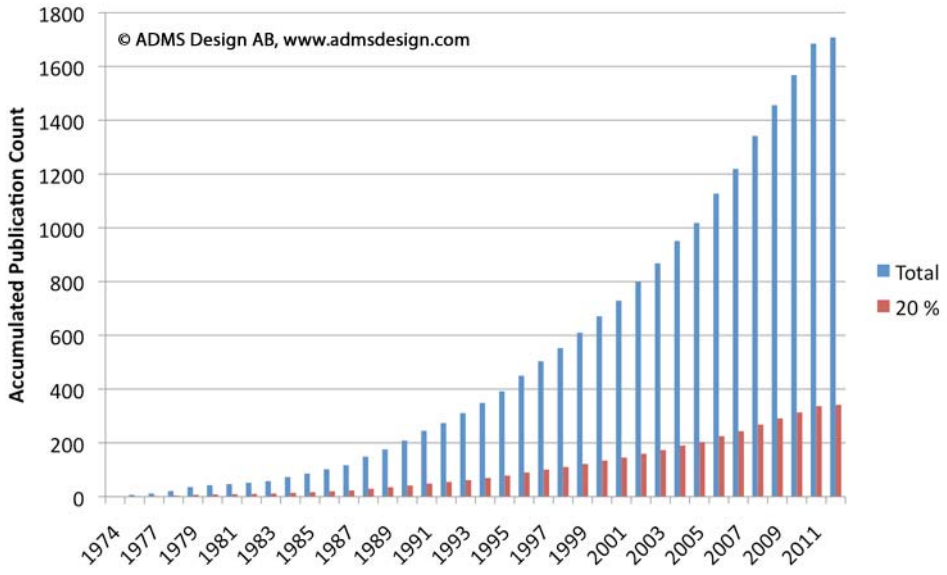

Figure 12.1. Accumulated publication count for scientifically reported ADC implementations in mainstream IEEE sources. The number of publications equivalent to 20% of total is indicated for reference.

12.1 Survey characteristics

When is it a "survey"? — I'm not going to spend too much energy on a stringent definition of a "survey". My guideline is that a survey should be based on a significant amount of data, and that the visualization, discussion and interpretation of the data *is* the main work. Many scientific papers nowadays include a scatter plot that compares a particular design with 5–20 relevant prior efforts. While it's a good idea to do so, these papers are not considered surveys in this context. Others use large amounts of empirical data to validate or derive a model, but the focus is more on the model.

What is "a significant amount of data"? — The size of the survey should be related to the total amount of data available at the time of the survey. A survey of 200 papers would have been exhaustive in 1990. Today it represents less than 12% of all scientific publications. The accumulated amount of scientific papers over time is shown in Figure 12.1. While this is not the *absolute* total number of ADC publications, it covers the ADC implementations reported in nearly all journals and conferences central to the A/D-converter field, and shall for simplicity be referred to as the "total" amount here. The number of sources equivalent to 20% of the accumulated total at any given time is also shown in Figure 12.1.

So, how much of the total do I need? Well, it depends on what you're trying to do. When it comes to survey data, I'm a firm believer in "*the more the merrier*", but there are tasks which

can be done with a fairly small subset. For example, if you want to get an idea of the overall trend for one parameter vs. another or just make a quick sanity check.

Small subsets *do* have some limitations though. For the subset to function as a reasonably generic approximation of the exhaustive set, its data must span roughly the same chunk of parameter space, and have similar distribution of values in all dimensions. This is difficult to achieve unless you make a random sampling of the exhaustive set. A smaller set also risks running out of data, for example when dividing it further according to some parameter such as resolution or architecture.

How quickly will a survey become dated? — I really don't know. I guess it depends on what you wish to study. But we can observe, as in Figure 12.2, how the accumulated total at any given time relates to the overall total (here 1708 papers), and what percentage of currently available works were yet unpublished at any given time (e.g., at the end year of a particular survey). It is seen that approximately **50%** of all currently available papers (~Q1-2012) were published in the last 8–8.5 years, i.e., ***after* 2003**, and almost **30%** were not yet published in 2007. By the end of 1997, **70%** of today's body of empirical data was still unpublished.

You can use Figure 12.2 to assess how old a survey can be before it's no longer useful for your purpose. Can you make business decisions based on trend estimates where the most recent half of the data set is missing? Probably not. Most recent 30/20/10%? If so, you need a survey that's less than approximately 4/3/1 years old.

It is clear that continuously updated surveys, such as Murmann's, or the one used by Converter Passion/ADMS Design AB are preferable over single-shot attempts, since the former allows for continuously updated trend estimations.

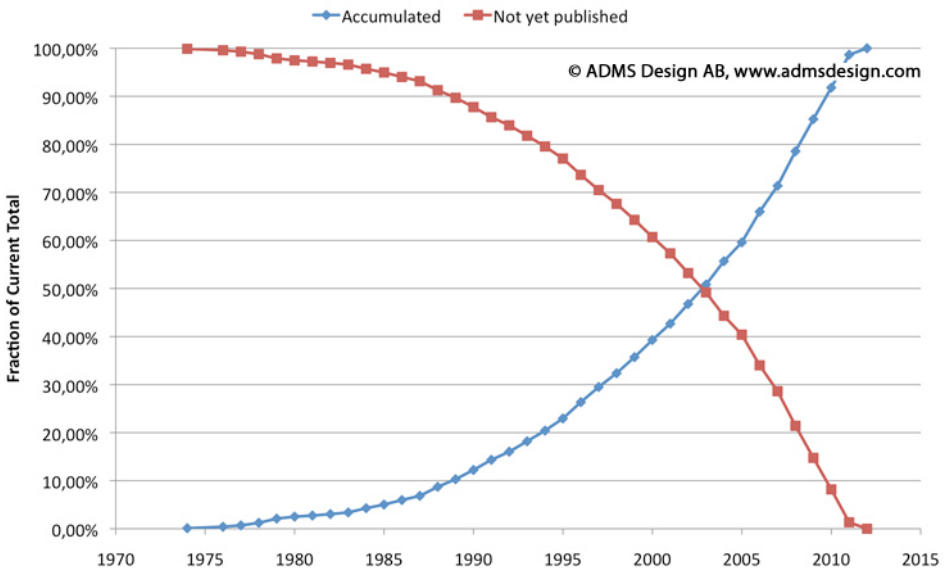

Figure 12.2. Paper "yield". The fraction of current total already published (blue) and yet to be published (red) at the end of any given year.

12.2 Known ADC surveys

Author	Size	Years	Type	What	Ref
Walden	100	≤ 1994	Both	Performance limits, FOM	[11]
Walden	150	1978-1997	Both	Performance limits, FOM, jitter, evolution	[10]
Merkel	150	1993-2002	Both	Performance limits, SFDR, power, VDD, scaling, device, architecture	[16]
Le	1000	1983-2004	Parts	Performance limits, jitter, cost, architecture, no. channels, N	[14]
Walden	175	1978-2007	Both	Update++	[12]
Walden	n/a	1978-2008	Both	Update	[13]
Murmann	~260	1997-2008	Sci	Performance limits, VDD, scaling, FOM, evolution	[15]
Jonsson	1400	1974-2010	Sci	Performance, VDD, scaling, FOM, evolution	[18]
Jonsson	1100	1976-2010	Sci	Performance & FOM vs. CMOS scaling, evolution	[17]
Fuiano	5540	1970-2010	Sci	Data-converters, research/patent correlation	[87]
Jonsson	1400	1974-2010	Sci	Energy/sample by architecture	[24]
Jonsson	1500	1974-2011	Sci	Area efficiency by architecture	[88]
Murmann	~350	1997-2012	Sci	Online survey data	[89]
Jonsson	1700	1974-2012	Sci	Performance, VDD, scaling, jitter, SFDR, FOM, evolution	[90]

Table 12.1. Known ADC surveys.

12.3 About the surveys

12.3.1 Walden

The “mother of all ADC surveys”, and the most frequently cited of all, is the pioneering work by *Walden* [10] where 150 scientific and commercial ADCs were analyzed, and performance trends were extracted. An earlier version was published already in 1994 [11], but this extended work became “*The Walden Survey*” to most of us. Although the 150 source documents originated from a mix of commercial and experimental designs, the *Walden* survey had a size equivalent to 30% of all scientific publications available at the time. The methods introduced in [10] are still useful, but Figure 12.1 and Figure 12.2 suggest that the trends extracted in [10] are unlikely to be valid and applicable today. At least they would have to be confirmed using more recent data. Two updated versions of the survey were published in 2008 – one covering 175 ADCs and data until 2007 [12], and one with an unspecified survey size and data until 2008 [13]. It is unclear how the 175 converters included in [12] were selected. During the time from *Walden’s* classic survey to 2007, the academic output alone generated another 715 new sources – commercial parts not counted. The +25

increase in source data therefore seems surprisingly incremental. Still, some of the results in [12] align very well with Converter Passion data, so apparently it was a carefully chosen subset.

12.3.2 Merkel & Wilson

Merkel and *Wilson* surveyed *150 commercial and scientific ADCs* with specifications suitable for defense space applications [16]. Their data appear to span from 1993–2002, and the selection criteria for inclusion in the survey was a sampling rate $fs \geq 1$ MS/s, and nominal resolution $N \geq 12$ bits. The paper does not reveal the mix between scientific papers and commercial parts, but *gathering 150 sources must have been quite an effort by the authors*. The total scientific output matching these specs and the time period is no more than 81 papers, and only 59 in the two sources (ISSCC, JSSC) the authors mention as primary. An additional minimum of 69–91 commercial parts must have been included to reach 150 sources. It is therefore assumed that the *Merkel* & *Wilson* data set was close to exhaustive for the spec range surveyed, and exhaustive data sets are always applauded here at Converter Passion.

The analysis and discussion itself is geared towards the stated application and focused on linearity (SFDR) to the extent that noise parameters are not treated at all. Power dissipation, supply voltage, speed, device type, scaling and architecture were observed.

12.3.3 Le, Rondeau, Reed & Bostian

An enormous data set, covering nearly 1000 commercial ADC parts from 1983–2004 was used in the survey by *Le*, *Rondeau*, *Reed* and *Bostian* [14]. As a comparison, the scientific output from the same years (not included in their survey) is 900 papers. The work is firmly rooted in the *Walden* tradition, but also considers parameters such as the number of channels per package and cost vs. performance. Additionally, the treatment separates the data by architecture, which adds an interesting extra dimension. Because of the larger volume and time span of the data set, part of the focus is to establish differences between this work and the classic *Walden* paper. Unfortunately, some exponentially improving parameters were plotted along linear axes, which makes many results from the survey difficult to see or interpret. Nevertheless, the contribution by *Le* et al. is a gigantic work and a key reference.

12.3.4 Murmann

The survey by *Murmann* [15] is a significant recent contribution to the analysis of empirical performance data. It covers approximately *260 scientific ADCs* reported 1997–2008 at the two conferences *VLSI Circuit Symposium* and *ISSCC*. The work analyzes ADC performance trends with a focus on *energy per sample* and *signal-to-noise-and-distortion ratio* (SNDR). The impact of process and voltage scaling is considered. If you don't have this paper already, you should definitely head over to IEEE Xplore and get it right now.

Murmann's survey has further benefits in that it is continuously updated and the data set is available online [89]. The latter opens up a lot of possibilities for anyone wishing to analyze the data in their own way, and makes the survey a very important contribution to the field. It currently includes around 350 sources.

12.3.5 Fuiano, Cagnazzo & Carbone

A rather different angle is taken in [87], where *Fuiano*, *Cagnazzo* and *Carbone* use survey data to analyze the correlation between scientific literature and patent activity. Compared to more "Waldenesque" surveys, this is a rather different animal. It nevertheless appeals to me as it illustrates an attempt to mine large amounts of survey data for something more unusual than ENOB, f_s and FOM.

12.3.6 Jonsson

The ADMS Design data set used at Converter Passion has also been used in five scientific papers, of which four are "surveys":

- ADC trends and performance evolution over time was analyzed in [18].
- The impact of CMOS scaling on ADC performance was empirically analyzed in [17].
- ADC architectures were compared with respect to energy efficiency in [24].
- Area-efficiency of ADC architectures was surveyed in [88].

The largest survey for which this data set has been used so far is the recently published Converter Passion article [90] of which this document is a slightly edited copy.

12.4 Other survey-related literature

A few other prior publications that are "survey-ish", or otherwise use a large set of empirical data for their analysis are listed here:

- *Vogels* and *Gielen* used a multidimensional regression fit to derive an ADC power dissipation model/FOM based on ≥ 70 empirical data points divided by architecture [91]. A similar approach was recently used by *Verhelst* and *Murmann* to analyze power dissipation and area vs. scaling based on Murmann's data set [92].
- *Sundström*, *Murmann*, and *Svensson* derived theoretical power dissipation bounds in [23], and used the Murmann set to compare theory with empirical reality.
- In [25], the author illustrated how the quality of a figure-of-merit (FOM) can be assessed by testing it against a large set of empirical data.

13. References

[1] G. L. Baldwin, “A linear delta modulator/demodulator with 10 Mbit/s sampling rate,” *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, Philadelphia, USA, pp. 192–193, Feb., 1974.

[2] G. L. Baldwin, and S. K. Tewksbury, “Linear delta modulator integrated circuit with 17-Mbit/s sampling rate,” *IEEE Trans. Circuit and Systems*, Vol. CAS-21, pp. 553–561, June, 1974.

[3] D. R. Breuer, and J. D. Hyde, “10-Bit, 5-Megasample/second monolithic A/D converter,” *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, Philadelphia, USA, pp. 152–153, Feb., 1976.

[4] A. G. E. Dingwall, and B. D. Rosenthal, “Low-power monolithic COS/MOS dual-slope 11-Bit A/D Converter,” *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, Philadelphia, USA, pp. 146–147, Feb., 1976.

[5] A. R. Hamadé, and E. Campbell, “A single-chip 8-bit A/D converter,” *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, Philadelphia, USA, pp. 154–155, Feb., 1976.

[6] F. H. Musa, and R. C. Taft, “A CMOS monolithic 31/2-digit A/D converter,” *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, Philadelphia, USA, pp. 144–145, Feb., 1976.

[7] R. A. Nordstrom, “High-speed integrated A/D converter,” *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, Philadelphia, USA, pp. 150–151, Feb., 1976.

[8] G.E. Moore, “Cramming more components onto integrated circuits,” *Electronics*, Vol. 38, No. 8, Apr. 1965.

[9] G. E. Moore, “No exponential is forever: but “forever” can be delayed!,” IEEE *ISSCC, Dig. Tech. Papers*, San Francisco, CA, Feb. 2003, pp. 20–23.

[10] R. H. Walden, “Analog-to-digital converter survey and analysis,” *IEEE J. Selected Areas in Communications*, no. 4, pp. 539–550, Apr. 1999.

[11] R. H. Walden, “Analog-to-digital converter technology comparison,” in *Proc. of GaAs IC Symp.*, pp. 228–231, Oct., 1994.

[12] R. Walden, “Analog-to-digital conversion in the early twenty-first century,” *Wiley Encyclopedia of Computer Science and Engineering*, pp. 126–138, Wiley, 2008.

[13] R. H. Walden, “Analog-to-digital converters and associated IC technologies,” in *Proc. Compound Semiconductor Integrated Circuits Symp.*, Monterey, Oct. 2008, pp. 1–2.

[14] B. Le, T. W. Rondeau, J. H. Reed, and C. W. Bostian, “Analog-to-digital converters [A review of the past, present, and future],” *IEEE Signal Processing Magazine*, pp. 69–77, Nov. 2005.

[15] B. Murmann, “A/D converter trends: Power dissipation, scaling and digitally assisted architectures,” *Proc. of IEEE Custom Integrated Circ. Conf. (CICC)*, San Jose, California, USA, pp. 105–112, Sept., 2008.

[16] K. G. Merkel, and A. L. Wilson, “A survey of high performance analog-to-digital converters for defense space applications,” in *Proc. IEEE Aerospace Conf.*, Big Sky, Montana, Mar. 2003, vol. 5, pp. 2415–2427.

[17] B. E. Jonsson, “On CMOS scaling and A/D-converter performance,” *Proc. of NORCHIP*, Tampere, Finland, pp. 1–4, Nov. 2010.

[18] B. E. Jonsson, “A survey of A/D-converter performance evolution,” *Proc. of IEEE Int. Conf. Electronics Circ. Syst. (ICECS)*, Athens, Greece, pp. 768–771, Dec., 2010.

[19] F. Maloberti, F. Francesconi, P. Malcovati, and O. J. A. P. Nys, “Design considerations on low-voltage low-power data converters,” *IEEE Trans. Circuits and Systems, pt. I*, vol. 42, no. 11, pp. 853–863, Nov. 1995.

[20] Y. Chiu, B. Nicolic, and P. R. Gray, “Scaling of analog-to-digital converters into ultra-deep-submicron CMOS,” in *Proc. Custom Integrated Circuits Conf.*, San Jose, Sept. 2005, pp. 375–382.

[21] S. Kawahito, “Low-power design of pipeline A/D converters,” in *Proc. Custom Integrated Circuits Conf.*, San Jose, Sept. 2006, pp. 505–512.

[22] C. Svensson, S. Andersson, and P. Bogner, “On the power consumption of analog to digital converters,” *Proc. of NORCHIP*, Linköping, Sweden, pp. 49–52, Nov., 2006.

[23] T. Sundström, B. Murmann, and C. Svensson, “Power dissipation bounds for high-speed Nyquist analog-to-digital converters,” *IEEE Trans. Circuits and Systems, pt. I*, vol. 56, no. 3, pp. 509–518, Mar. 2009.

[24] B. E. Jonsson, “An empirical approach to finding energy efficient ADC architectures,” *Proc. of 2011 IMEKO IWADC & IEEE ADC Forum*, Orvieto, Italy, pp. 1–6, June 2011. http://www.imeko.org/publications/iwadc-2011/IMEKO-IWADC-2011-28.pdf

[25] B. E. Jonsson, “Using Figures-of-Merit to Evaluate Measured A/D-Converter Performance,” *Proc. of 2011 IMEKO IWADC & IEEE ADC Forum*, Orvieto, Italy, pp. 1–6, June 2011. http://www.imeko.org/publications/iwadc-2011/IMEKO-IWADC-2011-43.pdf

[26] *International Technology Roadmap for Semiconductors (ITRS),* 2011 Edition [Online]. Available: http://www.itrs.net

[27] B. R. Carlton, H. Lakdawala, E. Alpman, J. Rizk, Y. W. Li, B. Perez-Esparza, V. Rivera, C. F. Nieva, E. Gordon, P. Hackney, C.-H. Jan, I. A. Young, and K. Soumyanath, “A 32nm, 1.05V, BIST enabled, 10-40MHz, 11-9 bit, 0.13mm2 digitized integrator MASH ΔΣ ADC,” *Symp. VLSI Circ. Digest of Technical Papers*, Kyoto, Japan, pp. 36–37, June, 2011.

[28] U. Wismar, D. Wisland, and P. Andreani, “A 0.2V 0.44µW 20 kHz analog to digital ∑Δ modulator with 57 fJ/conversion FoM”, *Proc. of Eur. Solid-State Circ. Conf. (ESSCIRC)*, Montreux, Switzerland, pp. 187-190, Sept., 2006.

[29] H. Pan and A. Abidi, “Spectral spurs due to quantization in Nyquist ADCs,” *IEEE Trans. Circuits and Systems, pt. I*, vol. 51, no. 8, pp. 1422–1439, Aug. 2004.

[30] M. S. Oude Alink, A. B. J. Kokkeler, E. A. M. Klumperink, K. C. Rovers, G. J. M. Smit and B. Nauta, “Spurious-free dynamic range of a uniform quantizer,” *IEEE Trans. Circuits and Systems, pt. II*, vol. 56, no. 6, pp. 434–438, June 2009.

[31] C. D. Motchenbacher, and J. A. Conelly, *Low noise electronic system design*, Wiley, 1993.

[32] B. P. Del Signore, D. A. Kerth, N. S. Sooch, and E. J. Swanson, “A monolithic 20-b delta–sigma A/D converter,” *IEEE J. Solid-State Circuits*, Vol. 25, pp. 1311–1317, Dec., 1990.

[33] D. Reynolds, and S. Ho, “An integrated 12 bit analog front end for CCD based image processing applications,” *Symp. VLSI Circ. Digest of Technical Papers*, Honolulu, USA, pp. 96–97, June, 1996.

[34] R. Naiknaware, and T. Fiez, “142dB ΔΣ ADC with a 100nV LSB in a 3V CMOS process,” *Proc. of IEEE Custom Integrated Circ. Conf. (CICC)*, Orlando, USA, pp. 5–8, May, 2000

[35] K. Philips, P. A. C. M. Nuijten, R. Roovers, F. Munoz, M. Tejero, and A. Torralba, “A 2mW 89dB DR continuous-time ΣΔ ADC with increased immunity to wide-band interferers,” *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, San Francisco, California, pp. 86–87, Feb., 2004.

[36] A. Thomsen, E. de Angel, S. X. Wu, A. Amar, L. Wang, and W. Lee, “A DC measurement IC with 130nVpp noise in 10Hz,” *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, San Francisco, California, pp. 334–335, Feb., 2000.

[37] A. M. A. Ali, A. Morgan, C. Dillon, G. Patterson, S. Puckett, P. Bhorashkar, H. Dinc, M. Hensley, R. Stop, S. Bardsley, D. Lattimore, J. Bray, C. Speir, and R. Sneed, “A 16-bit 250-MS/s IF Sampling Pipelined ADC With Background Calibration,” *IEEE J. Solid-State Circuits*, Vol. 45, pp. 2602-2612, Dec., 2010.

[38] J. van Valburg, and R. J. van de Plassche, “An 8-b 650-MHz Folding ADC,” *IEEE J. Solid-State Circuits*, Vol. 27, pp. 1662-1666, Dec., 1992.

[39] R. J. van de Plassche, and P. Baltus, “An 8-bit 100-MHz Full-Nyquist Analog-to-Digital Converter,” *IEEE J. Solid-State Circuits*, Vol. 23, pp. 1334-1344, Dec., 1988.

[40] V. E. Garuts, Y.-C. S. Yu, E. O. Traa, and T. Yamaguchi, "A Dual 4-bit 2-Gs/s Full Nyquist Analog-to-Digital Converter Using a 70-ps Silicon Bipolar Technology with Borosenic–Poly Process and Coupling-Base Implant," *IEEE J. Solid-State Circuits*, Vol. 24, pp. 216-222, Apr., 1989.

[41] S.-U. Kwak, B.-S. Song, and K. Bacrania, "A 15-b, 5-Msample/s Low-Spurious CMOS ADC," *IEEE J. Solid-State Circuits*, Vol. 32, pp. 1866-1875, Dec., 1997.

[42] W. Sansen, "Analog design challenges in nanometer CMOS technologies," *Proc. of IEEE Asian Solid-State Circ. Conf. (ASSCC)*, Jeju, Korea, pp. 5–9, Nov., 2007.

[43] J. Harrison, M. Nesselroth, R. Mamuad, A. Behzad, A. Adams, and S. Avery, "An LC Bandpass ΔΣ ADC with 70dB SNDR Over 20MHz Bandwidth Using CMOS DACs," *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, San Francisco, California, pp. 146–147, Feb., 2012.

[44] V. E. Garuts, E. O. Traa, Y.-C. S. Yu, and T. Yamaguchi, "A dual 4-bit, 1.5 Gs/s analog to digital converter," *IEEE Bipolar Circuits and Technology Meeting*, pp. 141–145, Sept., 1988.

[45] L. Singer, S. Ho, M. Timko, and D. Kelly, "A 12b 65MSample/s CMOS ADC with 82dB SFDR at 120MHz," *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, San Francisco, California, pp. 38–39, Feb., 2000.

[46] A. M. A. Ali, C. Dillon, R. Sneed, A. Morgan, J. Kornblum, L. Wu, and S. Bardsley, "A 14-bit 125 MS/s IF/RF sampling pipelined A/D converter," *Proc. of IEEE Custom Integrated Circ. Conf. (CICC)*, San Jose, USA, pp. 391–394, Sept., 2005.

[47] C.-Y. Chen, J. Wu, J.-J. Hung, T. Li, W. Liu, and W.-T. Shih, "A 12-Bit 3 GS/s Pipeline ADC With 0.4 mm2 and 500 mW in 40 nm Digital CMOS," *IEEE J. Solid-State Circuits*, Vol. 47, pp. 1013–1021, Apr., 2012.

[48] L. Luh, J. F. Jensen, C.-M. Lin, C.-T. Tsen, D. Le, A. E. Cosand, S. Thomas, and C. Fields, "A 4GHz 4th-order passive LC bandpass delta-sigma modulator with IF at 1.4GHz," *Symp. VLSI Circ. Digest of Technical Papers*, Honolulu, USA, pp. 168–169, June, 2006, IEEE.

[49] C. P. Hurrell, C. Lyden, D. Laing, D. Hummerston, and M. Vickery, "An 18 b 12.5 MS/s ADC With 93 dB SNR," *IEEE J. Solid-State Circuits*, Vol. 45, pp. 2647-2654, Dec., 2010.

[50] K. Bult, "Analog design in deep sub-micron CMOS," *Proc. of Eur. Solid-State Circ. Conf. (ESSCIRC)*, Stockholm, Sweden, pp. 126–132, Sept., 2000.

[51] A. Namdar, and B. H. Leung, "A 400MHz 12b 18mW IF digitizer with mixer inside a ΣΔ modulator loop," *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, San Francisco, California, pp. 62–63, Feb., 1999.

[52] A. M. A. Ali, C. Dillon, R. Sneed, A. S. Morgan, S. Bardsley, J. Kornblum, and L. Wu, "A 14-bit 125 MS/s IF/RF sampling pipelined ADC with 100 dB SFDR and 50 fs jitter," *IEEE J. Solid-State Circuits*, Vol. 41, pp. 1846–1855, Aug, 2006.

[53] R. Payne, M. Corsi, D. Smith, T.-L. Hsieh, and S. Kaylor, "A 16-Bit 100 to 160 MS/s SiGe BiCMOS Pipelined ADC With 100 dBFS SFDR," *IEEE J. Solid-State Circuits*, Vol. 45, pp. 2613-2622, Dec., 2010.

[54] K. Poulton, R. Neff, B. Setterberg, B. Wuppermann, T. Kopley, R. Jewett, J. Pernilo, C. Tan, and A. Montijo, "A 20GS/s 8b ADC with a 1MB memory in 0.18μm CMOS," *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, San Francisco, California, pp. 318–319, Feb., 2003.

[55] W. Cheng, W. Ali, M.-J. Choi, K. Liu, T. Tat, D. Devendorf, L. Linder, and R. Stevens, "A 3b 40GS/s ADC-DAC in 0.12μm SiGe," *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, San Francisco, California, pp. 262–263, Feb., 2004.

[56] L. Y. Nathawad, R. Urata, B. A. Wooley, and D. A. B. Miller, "A 20GHz bandwidth, 4b photoconductive-sampling time-interleaved CMOS ADC," *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, San Francisco, California, pp. 320–321, Feb., 2003.

[57] S. M. Louwsma, A. J. M. van Tuijl, M. Vertregt, and B. Nauta, "A 1.35 GS/s, 10 b, 175 mW Time-Interleaved AD Converter in 0.13 μm CMOS," *IEEE J. Solid-State Circuits*, Vol. 43, pp. 778–786, Apr., 2008.

[58] B. Chan, B. Oyama, C. Monier, and A. Guiterrez-Aiken, “An ultra-wideband 7-Bit 5-Gsps ADC implemented in submicron InP HBT technology,” *IEEE J. Solid-State Circuits*, Vol. 43, pp. 2187–2193, Oct., 2008.

[59] P. Schvan, J. Bach, C. Falt, P. Flemke, R. Gibbins, Y. Greshischev, N. Ben-Hamida, D. Pollex, J. Sitch, S.-C. Wang, and J. Woczanski, “A 24GS/s 6b ADC in 90nm CMOS,” *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, San Francisco, California, pp. 544–545, Feb., 2008.

[60] S. Shahramian, S. P. Voinigescu, and A. C. Carusone, “A 35-GS/s, 4-bit flash ADC with active data and clock distribution trees,” *IEEE J. Solid-State Circuits*, Vol. 44, pp. 1709–1720, June, 2009.

[61] J. Ryckaert, A. Geis, L. Bos, G. Vand Der Plas, and J. Craninckx, “A 6.1 GS/s 52.8 mW 43 dB DR 80 MHz Bandwidth 2.4 GHz RF Bandpass $\Delta\sum$ ADC in 40 nm CMOS,” *Proc. of IEEE Radio Frequency Integrated Circ. Symp*., Anaheim, CA, USA, pp. 443-446, May, 2010.

[62] Y. M. Greshishchev, J. Aguirre, M. Besson, R. Gibbins, C. Falt, P. Flemke, N. Ben-Hamida, D. Pollex, P. Schvan, and S.-C. Wang, “A 40GS/s 6b ADC in 65nm CMOS,” *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, San Francisco, California, pp. 390–391, Feb., 2010.

[63] M. K. Mayes, and S. W. Chin, “A 200 mW, 1 Msample/s, 16-b pipelined A/D converter with on-chip 32-b microcontroller,” *IEEE J. Solid-State Circuits*, Vol. 31, pp. 1862-1872, Dec., 1996.

[64] R. J. van de Plassche, and H. J. Schouwenaars, “A monolithic 14 bit A/D converter,” *IEEE J. Solid-State Circuits*, Vol. SC-17, pp. 1112-1117, Dec., 1982.

[65] J.-Y. Wu, Z. Zhang, R. Subramoniam, and F. Maloberti, “A 107.4 dB SNR multi-bit sigma delta ADC with 1-ppm THD at 0.12 dB from full scale input,” *IEEE J. Solid-State Circuits*, Vol. 44, pp. 3060-3066, Nov., 2009.

[66] R. Payne, C. Sestok, W. Bright, M. El-Chammas, M. Corsi, D. Smith, and N. Tal, “A 12b 1GS/s SiGe BiCMOS Two-Way Time-Interleaved Pipeline ADC,” *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, San Francisco, USA, pp. 182-184, Feb., 2011.

[67] AD7766, “24-Bit, 8.5 mW, 109 dB, 128 kSPS/64 kSPS/32 kSPS ADCs” *Data Sheet*, Analog Devices Inc., Aug, 2007.

[68] AD7986, “18-Bit, 2 MSPS PulSAR 15 mW ADC in LFCSP (QFN),” *Data Sheet*, Analog Devices Inc., Apr., 2009.

[69] MAX104, “+/-5V, 1Gsps, 8-bit ADC with on-chip 2.2GHz track/hold amplifier,” *Data Sheet*, Maxim Integrated Products Inc., Sept., 1999.

[70] R. C. Taft, P. A. Francese, M. R. Tursi, O. Hidri, A. MacKenzie, T. Hoehn, P. Schmitz, H. Werker, and A. Glenny, “A 1.8V 1.0GS/s 10b self-calibrating unified-folding-interpolating ADC with 9.1 ENOB at Nyquist frequency,” *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, San Francisco, California, pp. 78-79, Feb., 2009.

[71] S. Devarajan, L. Singer, D. Kelly, S. Decker, A. Kamath, and P. Wilkins, “A 16b 125MS/s 385mW 78.7dB SNR CMOS pipeline ADC,” *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, San Francisco, California, pp. 86–87, Feb., 2009.

[72] C.-Y. Chen, and J. Wu, “A 12b 3GS/s Pipeline ADC with 500mW and 0.4 mm2 in 40nm Digital CMOS,” *Symp. VLSI Circ. Digest of Technical Papers*, Kyoto, Japan, pp. 120–121, June, 2011.

[73] M. van Elzakker, E. van Tuijl, P. Geraedts, D. Schinkel, E. Klumperink, and B. Nauta, “A 1.9μW 4.4fJ/Conversion-step 10b 1MS/s Charge-Redistribution ADC,” *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, San Francisco, California, pp. 244–245, Feb., 2008.

[74] J. Xu, X. Wu, M. Zhao, R. Fan, H. Wang, X. Ma, and B. Liu, “Ultra Low-FOM High-Precision $\Delta\Sigma$ Modulators with Fully-Clocked SO and Zero Static Power Quantizers,” *Proc. of IEEE Custom Integrated Circ. Conf. (CICC)*, San Jose, California, USA, pp. 1–4, Sept., 2011.

[75] A. Shikata, R. Sekimoto, T. Kuroda, and H. Ishikuro, “A 0.5 V 1.1 MS/sec 6.3 fJ/Conversion-Step SAR-ADC With Tri-Level Comparator in 40 nm CMOS,” *IEEE J. Solid-State Circuits*, Vol. 47, pp. 1022–1030, Apr., 2012.

[76] T.-C. Lu, L.-D. Van, C.-S. Lin, C.-M. Huang, "A 0.5V 1KS/s 2.5nW 8.52-ENOB 6.8fJ/Conversion-Step SAR ADC for Biomedical Applications," *Proc. of IEEE Custom Integrated Circ. Conf. (CICC)*, San Jose, California, USA, pp. 1–4, Sept., 2011.

[77] S.-K. Lee, S.-J. Park, Y. Suh, H.-J. Park, and J.-Y. Sim, "A 1.3μW 0.6V 8.7-ENOB Successive Approximation ADC in a 0.18μm CMOS," *Symp. VLSI Circ. Digest of Technical Papers*, Honolulu, USA, pp. 242–243, June, 2009.

[78] H.-C. Hong, and G.-M. Lee, "A 65-fJ/Conversion-Step 0.9-V 200-kS/s Rail-to-Rail 8-bit Successive Approximation ADC," *IEEE J. Solid-State Circuits*, Vol. 42, pp. 2161–2168, Oct., 2007.

[79] M. D. Scott, B. E. Boser, and K. S. J. Pister, "An Ultra-Low Power ADC for Distributed Sensor Networks," *Proc. of Eur. Solid-State Circ. Conf. (ESSCIRC)*, Firenze, Italy, pp. 255–258, Sept., 2002.

[80] M. D. Scott, B. E. Boser, and K. S. J. Pister, "An Ultralow-Energy ADC for Smart Dust," *IEEE J. Solid-State Circuits*, Vol. 38, pp. 1123–1129, July, 2003.

[81] D. Muthers, and R. Tiekert, "A 0.11mm2 low-power A/D-converter cell for 10b 10MS/s operation," *Proc. of Eur. Solid-State Circ. Conf. (ESSCIRC)*, Leuven, Belgium, pp. 251–254, Sept., 2004.

[82] K. Bult, "Embedded analog-to-digital converters," *Proc. of Eur. Solid-State Circ. Conf. (ESSCIRC)*, Athens, Greece, pp. 52–60, Sept., 2009.

[83] B. E. Jonsson, "The path to a good A/D-converter FOM", *Converter Passion*, Aug., 2012, Available: http://converterpassion.wordpress.com/2011/07/24/the-path-to-a-good-ad-converter-fom/

[84] C. Wulff, and T. Ytterdal, "Design of a 7-bit, 200MS/s, 2mW pipelined ADC with switched open-loop amplifiers in a 65nm CMOS technology," *Proc. of NORCHIP*, Aalborg, Denmark, Nov., 2007.

[85] T. Sugawara, M. Ishibe, H. Yamada, S.-I. Majima, T. Tanji, and S. Komatsu, "A Monolithic 14 Bit/20 μs Dual Channel A/D Converter," *IEEE J. Solid-State Circuits*, Vol. SC-18, pp. 723-729, Dec., 1983.

[86] H. L. Fiedler, and B. Hoefflinger, "A CMOS Pulse Density Modulator for High-Resolution A/D Converters," *IEEE J. Solid-State Circuits*, Vol. SC-19, pp. 995-996, Dec., 1984.

[87] F. Fuiano, L. Cagnazzo, and P. Carbone, "Data Converters: an Empirical Research on the Correlation between Scientific Literature and Patenting Activity," *Proc. of 2011 IMEKO IWADC & IEEE ADC Forum*, Orvieto, Italy, pp. 1–6, June, 2011.

[88] B. E. Jonsson, "Area Efficiency of ADC Architectures," *Proc. of Eur. Conf. Circuit Theory and Design (ECCTD)*, Linköping, Sweden, pp. 560–563, Aug., 2011.

[89] B. Murmann, "ADC Performance Survey 1997-2012," [Online]. Available: http://www.stanford.edu/~murmann/adcsurvey.html.

[90] B. E. Jonsson, "A/D-converter Performance Evolution," *Converter Passion*, Aug., 2012, Available: http://converterpassion.wordpress.com/articles/ad-converter-performance-evolution/.

[91] M. Vogels, and G. Gielen, "Architectural Selection of A/D Converters," *Proc. of Des. Aut. Conf. (DAC)*, Anaheim, California, USA, pp. 974–977, June, 2003.

[92] M. Verhelst, and B. Murmann, "Area scaling analysis of CMOS ADCs," *El. Letters*, Vol. 48, No. 6, pp. 315–315, Mar., 2012, IEE.

14. Revision history

14.1 Version 1.0

The first version of the document.

14.2 Version 1.1

Inserted correct Fig. 6.3 (instead of a duplicate of 6.2).