

Chapter 2

ADC Architectures

2.1 Overview

In this chapter an overview of factors which typically determine ADC resolution is outlined. A brief discussion of popular Nyquist-rate ADC topologies is given, where the topologies most relevant to the focus of this book (besides pipelined ADCs) are discussed with the associated tradeoffs of each topology noted. The goal of this section is to provide a general overview and motivation of ADCs in electronic systems, as well as the system constraints which motivate a particular architecture over the other.

2.2 Factors Which Determine ADC Resolution and Linearity

Digital transmission is the most common form of data communication due to its superior signal integrity in the presence of noise. Digital signals can for example be transmitted at baseband as shown in Fig. 2.1 or modulated by a carrier signal, LO, to a higher frequency as shown in Fig. 2.2.

Different applications have different resolution requirements for the ADC in the receiver. For example, in many wireless receiver systems a receiver is required to resolve a received input with a minimum Signal to Noise and Distortion Ratio (SNDR), which is the ratio in dB of signal power to the power of all harmonics and total noise. The Effective Number Of Bits (ENOB) resolved by an ADC is determined according to the following formula [10]:

$$ENOB = \frac{SNDR - 1.76}{6.02} \tag{2.1}$$

In many applications the power of the received signal can vary significantly (e.g. the distance changing between transmitter and receiver in mobile applications), yet

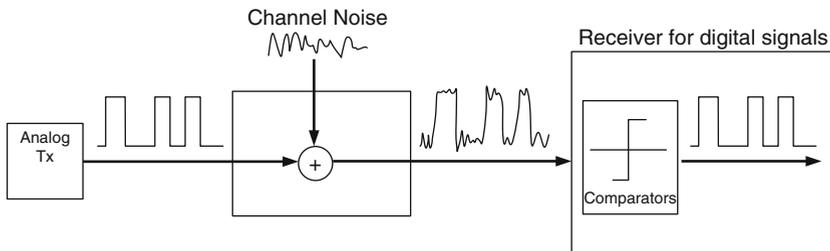


Fig. 2.1 Baseband binary digital transmission

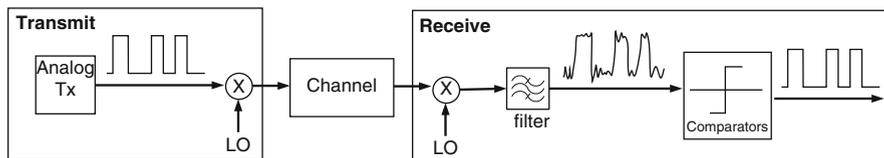


Fig. 2.2 Digital AM transmission

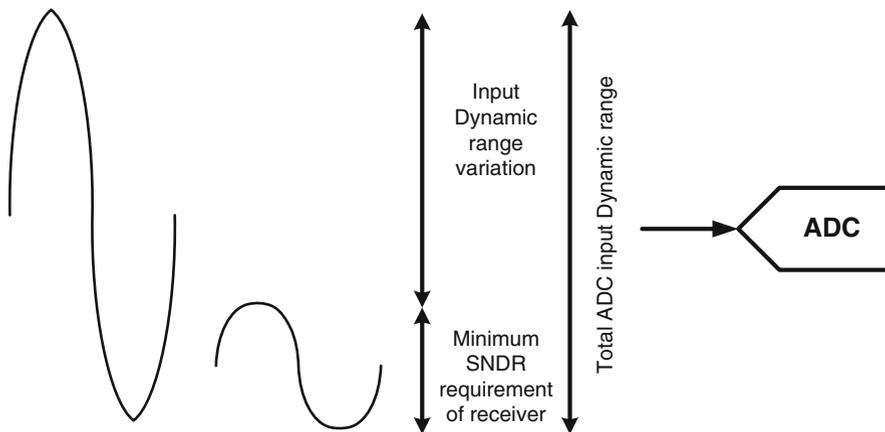


Fig. 2.3 Requirements of ADC input dynamic range

the ADC is still required to linearly digitize the full range of analog inputs with a minimum resolution. In many communication standards, the receive signal is required to maintain a minimum SNDR, thus the total dynamic range of the input (i.e. desired SNDR of the ADC) becomes the product of the minimum SNDR and variation in input signal strength as shown in Fig. 2.3.

In some applications it may be possible to include an Automatic Gain Control (AGC) circuit before the ADC to compensate for the effect of input dynamic range variation, thus reduce the dynamic range hence resolution of the ADC as shown in Fig. 2.4.

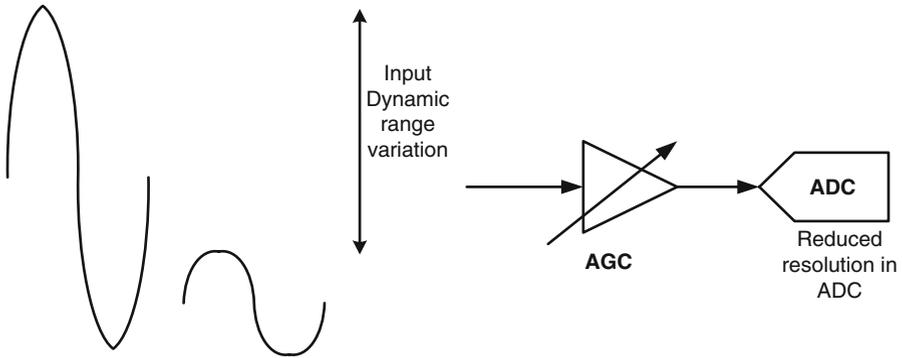


Fig. 2.4 AGC before ADC input to relax ADC input dynamic range requirements

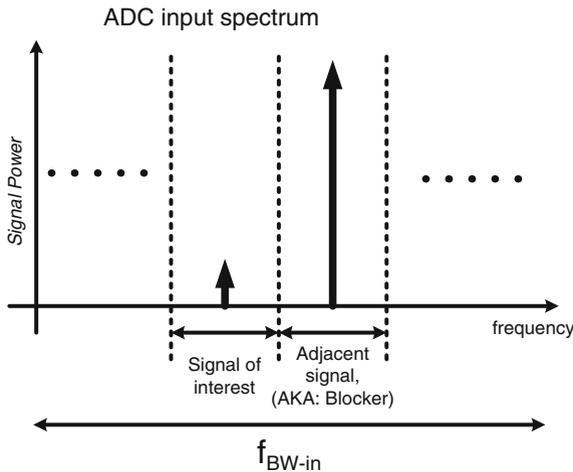


Fig. 2.5 ADC input spectrum illustrating case when desired signal is spectrally next to a much more powerful signal

The required resolution of an ADC can also be determined by the fact that in some applications the signal of interest which is to be digitized is adjacent to another signal in the frequency domain which is orders of magnitude more powerful, as shown in Fig. 2.5.

In such applications it may not be feasible to economically or adequately suppress the adjacent signal (referred to as a blocker) using analog techniques before the ADC input. As a result in many applications the entire signal bandwidth f_{BW-in} is digitized, and out of band signals efficiently eliminated using digital filters. In such applications to avoid introducing nonlinearity before digitally filtering out the blocker, the ADC must be able to linearly digitize both the large blocker and the signal of interest. Thus the dynamic range requirement of the ADC is set by the sum

in dB of the ratio of largest to smallest signal within f_{BW-in} , in addition to the minimum SNDR required to resolve the desired signal of interest.

The required linearity [10] requirement of an ADC can be determined by the fact that in many applications a large number of signals over a large bandwidth are digitized by a single ADC, such as for example OFDM systems, an example spectrum of which is shown in Fig. 2.6.

In an OFDM system the bandwidth of a single signal f_{sig-BW} is only a small fraction of the overall bandwidth f_{BW-in} digitized by the ADC. Thus the in-band thermal noise floor for a single signal is very small and given by $P_{NF}f_{BW-in}/N_{ch}$ (where N_{ch} is the total number of channels), and hence the ADC does not need to be designed with a very low input referred thermal noise floor. If the ADC has a nonlinear transfer characteristic however as shown in Fig. 2.7, each single tone input to the ADC produces an output with the same input tone plus additional harmonics, reducing the linearity in the digital output spectrum. Linearity is commonly assessed in ADCs by taking the ratio in dB of the largest and smallest harmonic tone in-band – referred to as the Spurious Free Dynamic Range (SFDR).

Thus in a system in which the input consists of several frequency adjacent tones such as OFDM, harmonics generated by each unique OFDM signal can fall into the

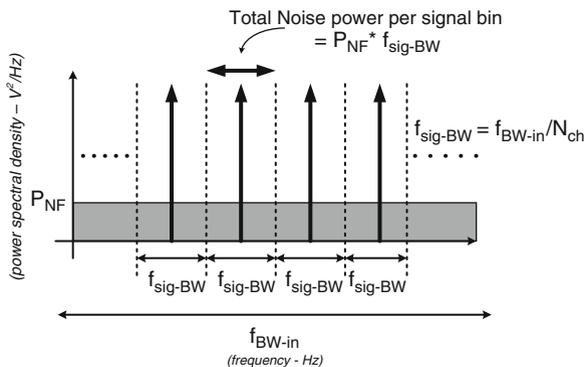


Fig. 2.6 Example OFDM spectrum

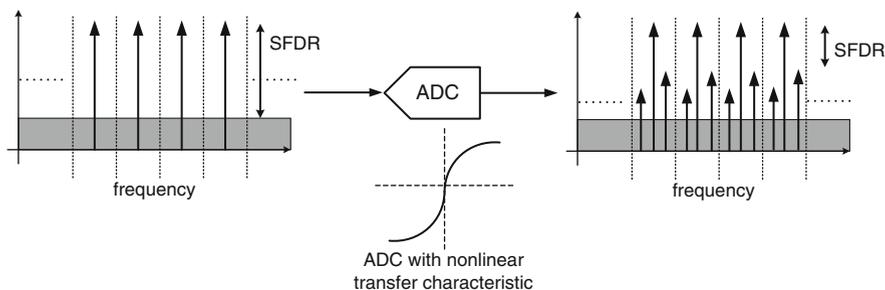


Fig. 2.7 Illustration of harmonic distortion in ADC output spectrum

bandwidth of another OFDM tone, thereby reducing the minimum SNDR within each signal bandwidth f_{sig-BW} . As a result in many systems the linearity of an ADC is the key design parameter, rather than the SNDR. In fact in some publications the ADC ENOB is calculated using the SFDR rather than the SNDR.

From the discussion in this section it is clear that the required resolution and linearity of an ADC can be determined by a number of factors. Furthermore by modifying the receive path by using an AGC and/or a combination of analog filtering techniques, the requirements of an ADC can be relaxed. The optimal configuration for a receiver depends on the specific constraints of a system. It is noted that the general trend in industry is to try and perform as much as possible in the digital domain, thereby pushing the ADC closer to the receiver input, thus demanding ADCs with more resolution and bandwidth – making ADCs essential enablers for future technologies.

2.3 ADC Architectures

As ADCs can consume a large percentage of power in a receiver, it is of vital interest to minimize ADC power consumption. Over the years several architectures have been developed which achieve optimal power consumption for different sampling rates, and resolutions as shown in Fig. 2.8.

The ADC topologies of Flash, SAR, and Pipelined are reviewed in subsequent sections as they are essential to understand within the context of this book. A detailed discussion of topologies not discussed in this book (Delta Sigma, and

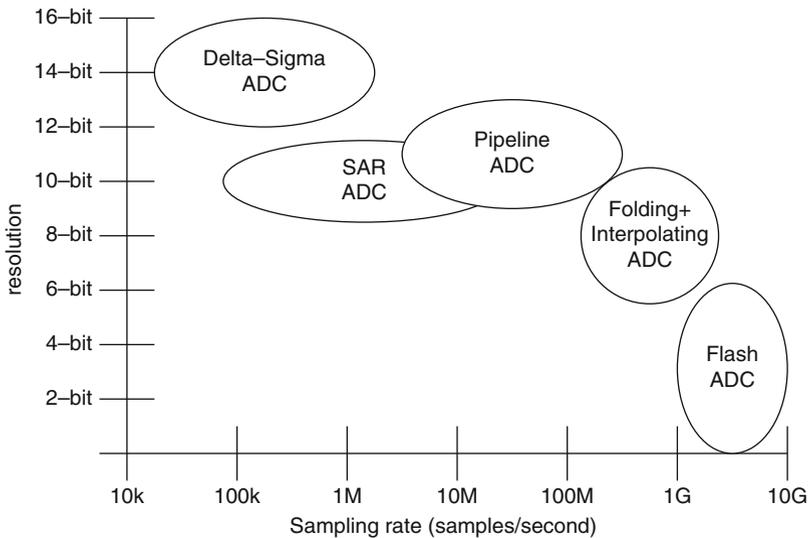


Fig. 2.8 ADC architecture comparison

Table 2.1 Comparison of ADC architectures

Architecture	Latency	Speed	Accuracy	Area
Flash	Low	High	Low	High
SAR	Low	Low-medium	Medium-high	Low
Folding + interpolating	Low	Medium-high	Medium	High
Delta-sigma	High	Low	High	Medium
Pipeline	High	Medium-high	Medium-high	Medium

Folding + Interpolating) can be found in [10]. Table 2.1 summarizes the key tradeoffs of the different ADC topologies shown in Fig. 2.8.

2.4 ADC Figure-of-Merit

A popular Figure-of-Merit (FOM) used to compare different ADCs is

$$FOM = \frac{Power}{(2^{ENOB})(f_s)} \text{ (pJ/step)} \quad (2.2)$$

where f_s is the sampling rate in Nyquist-rate ADCs. This figure of merit is commonly used to compare published reports as the accuracy term is based on easily measured quantities, and calculates a value that has meaningful units (i.e. energy required per conversion step – thus lower FOM means a better ADC).

A slight variation of equation (2.2) is the following FOM:

$$FOM = \frac{Power}{(2^{2 \times ENOB})(f_s)} \text{ (pJ/step)} \quad (2.3)$$

In equation (2.3) the ENOB is multiplied by 2 to account for the fact that due to thermal noise limitations, to achieve twice the resolution $\times 4$ the power is required. In general similar FOMs can be achieved with different ADC topologies, however it is noted that ADCs with lower resolutions tend to be able to achieve better FOMs using equation (2.2).

2.5 Flash ADC

Many ADC architectures have been developed over the years, each with different tradeoffs with respect to power, speed, and accuracy. Most ADC architectures however are in some form a variant of the Flash ADC or use a Flash ADC in their implementation.

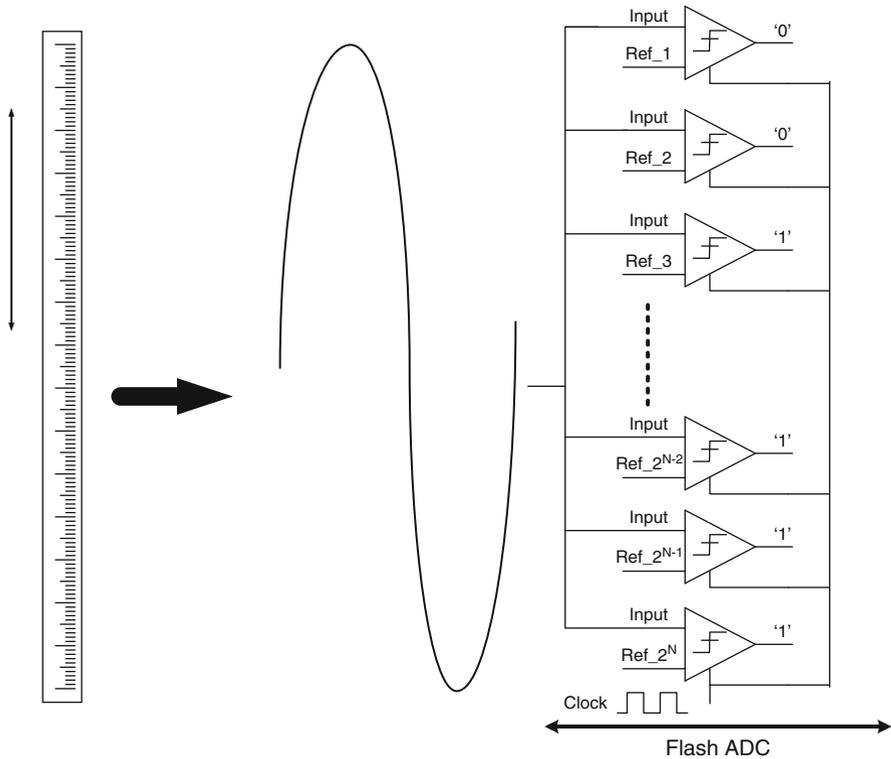


Fig. 2.9 Analogy between a ruler and a flash ADC

Much like a ruler with a fixed resolution maps an infinite precision length to a finite accuracy (e.g. measure a length in millimeters); Flash ADCs *measure* an analog signal into a digital signal by comparing an analog input with fixed reference values as shown in Fig. 2.9. The number of fixed references used determines the accuracy of the digital output. For example, 4-bit accuracy is obtained by comparing against $2^4 - 1 = 15$ reference values, 10-bit accuracy by comparing against $2^{10} - 1 = 1,023$ reference values. Determining which reference values the input is in-between forms a length 2^N bit (where N is the accuracy of the ADC) thermometer code representation of the analog input. Mapping the unique thermometer code to its binary equivalent forms a length N binary representation of the analog input [10].

In a Flash ADC, the number of comparators required is thus exponentially related to the desired resolution (in bits). It is noted that to achieve increased resolution in Flash ADCs, large devices are required to suppress process variation effects. Hence in the interest of minimizing area, Flash ADCs are most commonly used in applications where only low resolutions are required.

One of the key advantages of the Flash topology is that it has a potential latency of only one clock cycle – that is the digital output is available within one clock cycle

of the input being sampled. In certain systems where an ADC is required in a control loop (e.g. quantizer in a Delta Sigma [10]), it is critical to implement the ADC with as low a latency as possible to maximize closed-loop stability. As low latency is an attractive feature in some systems, many techniques have been developed to enable increased resolution while reducing area consumption, using essentially the Flash topology (e.g.: folding, interpolating, and averaging [10]).

2.6 SAR ADC

The algorithm which forms the basis of the Successive Approximation Register (SAR) ADC has been known since the 1500s, however it was first patented as an algorithm for use in ADCs in 1958 by Bernard M. Gordon [11]. SAR became a popular topology to implement ADCs in the 1970s with the availability of several logic ICs from companies such as AMD. The SAR ADC is relevant to an understanding of pipelined ADCs as from Fig. 2.8 the SAR and pipelined ADC have overlapping areas of use – roughly for low-medium speed applications with 8–10 bits of desired resolution. Thus a clear understanding of each topology and their associated tradeoffs is beneficial in understanding under which circumstances a designer would choose one or the other.

The algorithm used in Successive Approximation is based on a binary search algorithm, and thus is more component efficient than Flash ADCs which use a brute force approach to perform data conversion. Fig. 2.10 illustrates the topology of a standard SAR ADC.

In a SAR ADC the analog input is sampled by a sample-and-hold circuit which operates at the effective Nyquist sampling rate of the ADC, f_s . A sequential binary search is performed on the sampled input by initializing the N-bit register to midscale,

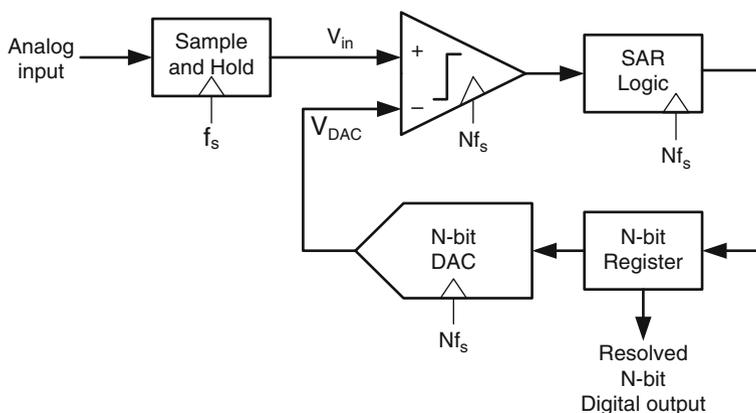


Fig. 2.10 SAR ADC topology

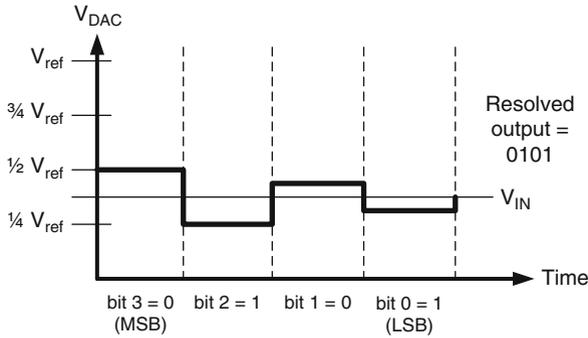


Fig. 2.11 Example SAR conversion – DAC voltage

which forces the decision threshold of the comparator to be $V_{\text{ref}}/2$ (where V_{ref} is the full scale input voltage). As a result if the sampled input is greater than $V_{\text{ref}}/2$, the MSB of the N-bit register remains at ‘1’, whereas if the input is below $V_{\text{ref}}/2$, the MSB of the N-bit register is changed to ‘0’ [12]. By successively repeating the same algorithm and initializing the next bit in the N-bit register to ‘1’, the digital representation of the analog input can be determined to N bits, where N is the resolution of the SAR ADC and the number of times the algorithm is repeated. Figure 2.11 illustrates an example of 4-bit conversion.

The significant advantage of the SAR ADC is that it uses only a few analog components (notably only a single comparator) to implement N-bit data conversion, resulting in a compact area and simple design. Furthermore since the topology produces a new digital output every $1/f_s$, the latency of the ADC with respect to the effective sampling rate is only one clock cycle of the Nyquist-rate clock f_s . As a result the SAR topology can be useful in systems which require ADCs in feedback.

For the SAR ADC to operate with an effective sampling rate of f_s however, the comparator, DAC, and SAR logic are required to operate at Nf_s . For example, if the desired sampling rate and resolution are 100 MHz and 10-bits respectively, the DAC, comparator, and SAR logic are required to operate at 1 GHz. Thus although the SAR ADC allows for a significant reduction in the number of analog components it comes at the cost of restricting the maximum sampling rate to only a fraction of the maximum speed available by a given technology. From a system designer’s point of view a SAR ADC may not be feasible in some systems where the available clock is only at the Nyquist sampling rate. As a result SAR ADCs have traditionally been restricted to low to medium speed, and medium to high accuracy applications.

From a design perspective, it is noted that while the SAR topology determines 1-bit of the final digital output every clock cycle, the DAC is required to settle to the full accuracy of the ADC every clock cycle. Also, while any static offset in the comparator appears as an input referred offset to the ADC, the comparator is required to be able to resolve inputs as small as the LSB of the ADC. As a result

much effort is required to optimize the DAC and comparator blocks for high speed and high accuracy. It should be noted however that recently published reports [13,14] have begun to emerge which show that low power passive DACs based on charge sharing between capacitors can be used to reduce power consumption and increase the operating speed of SAR ADCs to the point that in 90 nm operating speeds of 50 MS/s can be achieved with resolutions on the order of 8–9 bits while only consuming 0.7 mW [13].

2.7 Sub-sampling

The majority of Nyquist-rate ADCs use a sampling rate which is twice the highest frequency component of the input - this guarantees that the sampled input can be perfectly reconstructed (ignoring the ADC's resolution). In some applications however it is desirable to sample at a fraction of the input frequency – this type of sampling is referred to as sub-sampling. Sub-sampling is commonly used to alias a high frequency input down to a lower frequency. To understand how this happens consider a sinusoidal input $y(t)$ of frequency $(f_o + f_{LO})$, which is sampled by an ADC at a sampling rate of f_s . The discrete time values of the input digitized by the ADC are given by:

$$y[m] = \sin\left(2\pi\frac{f_o + f_{LO}}{f_s}m\right) \quad (2.4)$$

If $f_o + f_{LO}$ is greater than f_s , the periodicity of the sine function results in $y[m]$ appearing as if it were sampled at a much lower frequency: modulo $(f_o + f_{LO}, f_s)$. Thus rather than use a mixer to translate an input bandwidth to a lower frequency as shown in Fig. 2.12, sub-sampling can be used to efficiently perform frequency translation as shown in Fig. 2.13. The tradeoff in using a sub-sampled ADC is the maximum input frequency of the ADC is significantly increased.

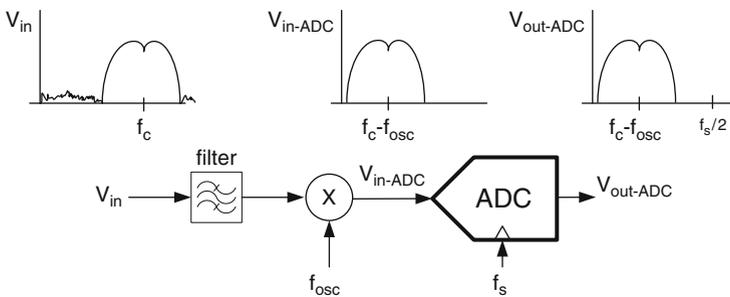


Fig. 2.12 Frequency translation using a front-end mixer

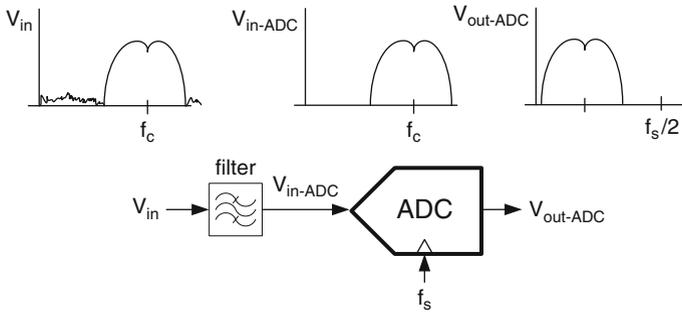


Fig. 2.13 Frequency translation using sub-sampling

2.8 Summary

This chapter an overview of ADC topologies was given. A brief discussion of the system level tradeoffs which determine ADC resolution and speed was given. Nyquist-rate ADC architectures were introduced, with an emphasis of the discussion on the Flash and SAR ADC topologies. The use of ADCs in sub-sampled systems was also described.