

Understanding Flash ADCs

Flash analog-to-digital converters, also known as parallel ADCs, are the fastest way to convert an analog signal to a digital signal. Flash ADCs are ideal for applications requiring very large bandwidth, however, they typically consume more power than other ADC architectures and are generally limited to 8-bits resolution. Flash ADCs are made by cascading high-speed comparators. Each comparator represents 1 LSB, and the output code can be determined in one compare cycle. This tutorial will also talk about flash converters vs. other converter types.

Flash analog-to-digital converters, also known as parallel ADCs, are the fastest way to convert an analog signal to a digital signal. They are suitable for applications requiring very large bandwidths. However, flash converters consume a lot of power, have relatively low resolution, and can be quite expensive. This limits them to high frequency applications that typically cannot be addressed any other way. Examples include data acquisition, satellite communication, radar processing, sampling oscilloscopes, and high-density disk drives.

Architecture Detail

Figure 1 shows a typical flash ADC block diagram. For an "N" bit converter, the circuit employs $2^N - 1$ comparators. A resistive divider with 2^N resistors provides the reference voltage. The reference voltage for each comparator is one least significant bit (LSB) greater than the reference voltage for the comparator immediately below it. Each comparator produces a "1" when its analog input voltage is higher than the reference voltage applied to it. Otherwise, the comparator output is "0". Thus, if the analog input is between v_{x4} and v_{x5} , comparators x_1 through x_4 produce "1"s and the remaining comparators produce "0"s. The point where the code changes from ones to zeros is the point where the input signal becomes smaller than the respective comparator reference voltage levels.

This is known as thermometer code encoding, so named because it is similar to a mercury thermometer, where the mercury column always rises to the appropriate temperature and no mercury is present above that temperature. The thermometer code is then decoded to the appropriate digital output code.

The comparators are typically a cascade of wideband low gain stages. They are low gain because at high frequencies it's difficult to obtain both wide bandwidth and high gain. They are designed for low voltage offset, such that the input offset of each comparator is smaller than a LSB of the ADC. Otherwise, the comparator's offset could falsely trip the comparator, resulting in a digital output code not representative of a thermometer code. A regenerative latch at each comparator output stores the result. The latch has positive feedback, so that the end state is forced to either a "1" or a "0".




Figure 1. Flash ADC architecture

Some reality checks now need to be added to optimize the flash converter architecture.

Sparkle Codes

Normally, the comparator outputs will be a thermometer code, such as 00011111. Errors may cause an output like 00010111 (i.e., there is a spurious zero in the result). This out of sequence "0" is called a sparkle. This may be caused by imperfect input settling or comparator timing mismatch. The magnitude of the error can be quite large. Modern converters like the MAX104 employ an input track-and-hold in front of the ADC along with an encoding technique that suppresses sparkle codes.

Metastability

When a digital output of a comparator is ambiguous (neither a one nor a zero), the output is defined as metastable. Metastability can be reduced by allowing more time for regeneration. Gray-code encoding can also greatly improve metastability. Gray-code encoding allows only one bit in the output to change at a time. The comparator outputs are first converted to gray-code encoding and then later decoded to binary if desired.

Another problem occurs when a metastable output drives two distinct circuits. It is possible for one circuit to declare the input a "1" while the other circuit thinks it's a "0". This can create major errors. To avoid this, only one circuit should sense a potentially metastable output.

Input Signal Frequency Dependence

When the input signal changes before all the comparators have completed their decision, the ADC performance is adversely impacted. The most serious impact is a drop-off in signal-to-noise ratio plus distortion (SINAD) as the frequency of the analog input frequency increases.

Measuring spurious free dynamic range (SFDR) is another good way to observe converter performance. The "effective bits" achieved is a function of input frequency. This can be improved by adding a track-and-hold (T/H) circuit in front of the ADC. This allows dramatic improvement, especially when input frequencies approach the Nyquist frequency, as shown in Figure 2 (taken from the MAX104 data sheet). Parts without the track-and-hold show a significant drop-off in SFDR.




Figure 2. Spurious free dynamic range as a function of input frequency

Clock Jitter

Signal-to-noise ratio (SNR) is degraded when there is jitter in the sampling clock. This becomes noticeable for high analog input frequencies. To achieve accurate results, it is critical to provide the ADC with a low-jitter, sampling clock source.

Architecture Tradeoffs

ADCs can be implemented by employing a variety of architectures. The principal tradeoffs between these alternatives are:

- *The time it takes to complete a conversion (conversion time).* For flash converters, the conversion time does not change materially with increased resolution. The conversion time for Successive Approximation Register (SAR) or Pipelined converters increases approximately linearly with an increase in resolution (Figure 3a). For integrating ADCs, the conversion time doubles with every bit increase in resolution.
- *Component matching requirements in the circuit.* Flash ADC component matching typically limits resolution to around 8-bits. Calibration and trimming are sometimes used to improve the matching available on chip. Component matching requirements double with every bit increase in resolution.

This applies to flash, successive approximation or pipelined converters, but not integrating converters. For integrating converters, component matching does not materially increase with an increase in resolution (Figure 3b).

- *Die size, cost and power.* For flash converters, every bit increase in resolution almost doubles the size of the ADC core circuitry. The power also doubles. In contrast, a SAR, Pipelined, or sigma-delta ADC die size will increase linearly with an increase in resolution, and an integrating converter core die size will not materially change with an increase in resolution (Figure 3c). An increase in die size increases cost.

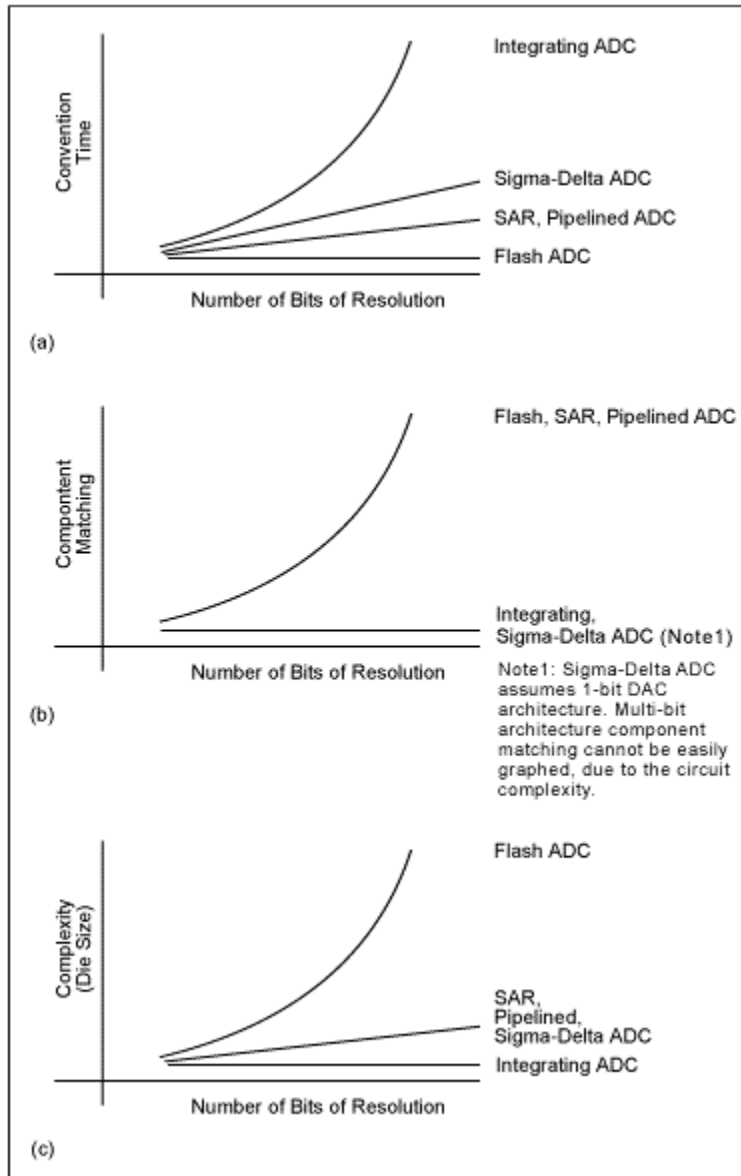


Figure 3. Architecture tradeoffs

FLASH ADC vs. Other ADC Architectures

Flash vs. Successive Approximation Register (SAR) ADCs

In a SAR converter, the bits are decided by a single high-speed, high-accuracy comparator one bit at a time (from the MSB down to the LSB), by comparing the analog input with a DAC whose output is updated by previously decided bits and thus successively approximates the analog input. This serial nature of the SAR

limits its speed to no more than a few Msps, while flash ADCs exceed giga-sample per second (Gsps) conversion rates.

SAR converters are available in resolutions up to 16-bits. An example of such a device is the MAX1132. Flash ADCs are typically limited to around 8-bits. The slower speed also allows the SAR ADC to be much lower in power. For example, the MAX1106, an 8-bit SAR converter, uses 100 μ A at 3.3V with a conversion rate of 25ksps. The MAX104 dissipates 5.25W. This is about 16,000 times higher power consumption compared to the MAX1106, but also 40,000 times faster in terms of its maximum sampling rate.

The SAR architecture is also less expensive. The MAX1106 at 1k volumes sells at approximately \$1.51, while the MAX104 sells at roughly \$398. Package sizes are larger for flash converters. In addition to a larger die size requiring a larger package, the package needs to dissipate a lot of power and needs many pins for power and ground signal integrity. The package size of the MAX104 is more than 50 times larger than the MAX1106.

Flash vs. Pipelined ADCs

A pipelined ADC employs a parallel structure in which each stage works on one to a few bits of successive samples concurrently. This improves speed at the expense of power and latency. However, each pipelined stage is much slower than a flash section. The pipelined ADC requires accurate amplification in the DACs and interstage amplifiers, and these stages have to settle to the desired linearity level. By contrast, in a flash ADC, the comparator only needs to be low offset and be able to resolve its inputs to a digital level (i.e., there is no linear settling time involved). However, some flash converters require preamplifiers to drive the comparators. Gain linearity needs to be carefully specified.

Pipelined converters are capable of conversion speeds of around 100Msps at 8 to 14-bit resolutions. An example of a pipelined converter is the MAX1449, a 105MHz, 10-bit ADC. For a given resolution, pipelined ADCs are around 10 times slower compared to flash converters of similar resolution. Pipelined converters are possibly the optimal architecture for ADCs that need to sample at rates up to around 100Msps with resolution at 10-bits and above. At resolutions of up to 10-bits, and conversion rates above a few hundred Msps, flash ADCs dominate.

Interestingly, there are some situations where flash ADCs are hidden inside a converter employing another architecture to increase its speed. This is the case, for example, in the MAX1200; a 16-bit pipelined ADC that includes an internal 5-bit flash ADC.

Flash vs. Integrating ADCs

Single, dual and multi-slope ADCs can achieve high resolutions of 16-bits or more are relatively inexpensive and dissipate materially less power. These devices support very low conversion rates, typically less than a few hundred samples per second. Most applications are for monitoring DC signals in the instrumentation and industrial markets. This architecture competes with sigma-delta converters.

Flash vs. Sigma-Delta ADCs

Flash ADCs do not compete with this architecture because currently the achievable conversion rates differ by up to two orders of magnitude. The sigma-delta architecture is suitable for applications with much lower bandwidth, typically less than 1MHz, with resolutions in the 12 to 16-bit range. These converters are capable of the highest resolution possible in ADCs. They require simpler anti-alias filters (if needed) to bandlimit the signal prior to conversion.

They trade speed for resolution by oversampling, followed by filtering to reduce noise. However, these devices are not always efficient for multi-channel applications. This architecture can be implemented by using sampled data filters (also known as modulators) or continuous time filters. For higher frequency conversion rates the continuous time architecture is potentially capable of reaching conversion rates in the hundreds of Msps range with low resolution of 6 to 8-bits. This approach is still in the early research and development stage and offers competition to flash alternatives in the lower conversion rate range.

Another interesting use of a flash ADC is as a building block inside a sigma-delta circuit to increase the conversion rate of the ADC.

Sub-Ranging ADCs

When higher resolution converters or smaller die size and power for a given resolution are needed, multi-stage conversion is employed. This architecture is known as a sub-ranging converter. This is also sometimes referred to as a multi-step or half-flash converter. This approach combines ideas from successive approximation and flash architectures.

Sub-ranging ADCs reduce the number of bits to be converted into smaller groups, which are then run through a lower resolution flash converter. This approach reduces the number of comparators and reduces the logic complexity, compared to a flash converter (Figure 4). The tradeoff results in slower conversion speed compared to flash.

The MAX153 is an 8-bit, 1Msps ADC implemented with a sub-ranging architecture. This circuit employs a two-step technique. Here a first conversion is completed with a 4-bit converter. A residue is created, where the result of the 4-bit conversion is converted back to an analog signal (with an 8-bit accurate DAC) and subtracted from the input signal. This residue is again converted by the 4-bit ADC and the results of the first and second pass are combined to provide the 8-bit digital output.

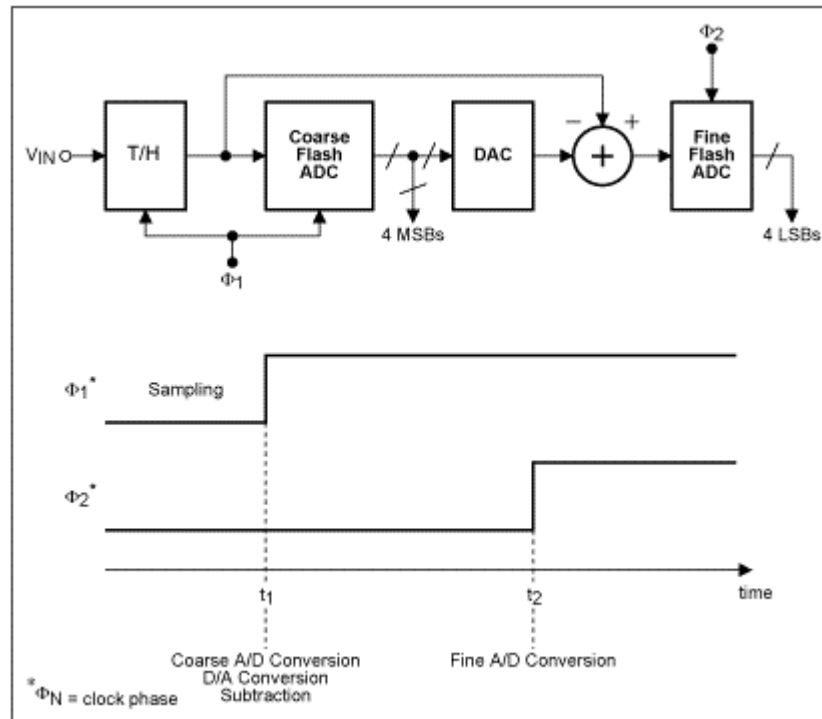


Figure 4. Sub-ranging ADC architecture

Process Technology

The fastest monolithic converters are built using bipolar technology. Flash converter speeds are currently in excess of 1Gsps. Examples are the MAX104/MAX106. CMOS flash converters are available at lower speed and resolutions compared to bipolar technology offerings and are typically intended for integration into a larger CMOS circuit. CMOS, BiCMOS and bipolar technologies will continue to improve, yielding increasingly higher conversion rates.

Conclusion

For applications requiring modest resolutions, typically up to 8-bits, at sampling frequencies in the high hundreds of MHz, the flash architecture may be the only viable alternative. The user must supply a low-jitter clock to ensure good ADC performance. For applications with high analog input frequencies, the ADC chosen should have an internal track-and-hold.

September 2001

MORE INFORMATION

MAX104:	QuickView	-- Full (PDF) Data Sheet (896k)	
MAX105:	QuickView	-- Full (PDF) Data Sheet (680k)	-- Free Sample
MAX1106:	QuickView	-- Full (PDF) Data Sheet (216k)	-- Free Sample
MAX153:	QuickView	-- Full (PDF) Data Sheet (456k)	-- Free Sample
MAX196:	QuickView	-- Full (PDF) Data Sheet (176k)	-- Free Sample