Data Converters: Getting to Know Dynamic Specs
ADC-AN-3

By Bob Leonard

The ongoing evolution in data converters requires that they be characterized beyond static accuracy and throughput.

A kaleidoscope of architectures is available for today's data converters. But ensuring that a data converter can obtain a desired transfer rate and accuracy at a particular input frequency is more important than any architectural nuances. Now that converter data sheets carry frequency-domain specifications, designers are finally gaining insight into how they'll perform on the job.

A designer's first consideration is to ensure that the converter can provide data at the required throughput. The second concern is whether the data converter can meet the system's static accuracy requirements. Stated simply, when the converter operates at the designated throughput, does it yield the desired transfer accuracy for a dc input? A third consideration is that the data converter, when operating at the designated throughput, must deliver the required dynamic accuracy. In other words, is the data converter accurately digitizing a high-frequency or transient input? This "dynamic accuracy" is often best assessed by using frequency-domain specifications.

Designers face many pitfalls in assessing throughput and static and dynamic accuracy, however. For one thing, the three performance considerations mentioned previously are often interrelated. For another, time and error budgets need to take into account the overall performance required for a system. And to assess what's needed in a data-acquisition system, designers must know the idiosyncrasies of particular analog-to-digital converter (ADC) architectures and sample-and-hold amplifiers (SHAs). For multichannel applications, an error analysis must be expanded to include the input multiplexer and instrumentation amplifier.

In the past, engineers tackled their data-converter requirements by buying a separate, stand-alone ADC. If the analog inputs were beyond basically dc signals, such as those found in measurements of temperature, pressure, and so forth, then a companion SHA was required. Some applications today still use only an ADC because of slow input signals. Others have a separate ADC and SHA. Sampling ADCs, which include an internal ADC and SHA in one package, are growing increasingly popular.

But designers must scrutinize other factors before deciding that a sampling ADC can deliver the needed accuracy under certain dynamic conditions. Clean input digitization when the ADC operates up to the Nyquist rate (signal frequency at one half the sampling rate) isn't a given. The flash ADC category is another example of an ADC used without a preceding SHA. The generic family of flash ADCs includes devices from 4 to 10 bits of accuracy with throughputs from 10MHz to 500MHz.

Some flash ADC architectures may or may not need a SHA in digitizing high-frequency signal inputs. Available flash ADC architectures include the "true flash" with a straight comparator design, an interpolation/folding design, and a two-pass design. Two factors determine if a SHA is needed. One is the aperture jitter specification, which often depends on architecture. The other deciding factor is the ability to drive the flash ADC's dynamic input impedance at the desired frequency.

As previously noted, there's a proliferation of ADC architectures today. For example, the integrating ADC with various slope architectures is still focused on slow, high-resolution applications – it doesn't require an external SHA.

The successive-approximation ADC and companion SHA have long been workhorses for medium speed and resolution. Now they share some of these applications with the deltasigma ADC.

The delta-sigma architecture has an inherent SHA function built in. To obtain increased speeds beyond the range of 20 kHz to 1 MHz, the successive approximation and deltasigma ADCs give way to the two-pass ADC. The various types of two-pass (also known as a two-step or subranging ADC) and companion SHAs also reach their speed/resolution limitations as next-generation limits are pushed.

The successive-approximation ADC and two-pass ADCs' speed (kilohertz to megahertz range) suit most applications beyond dc measurements. They use sample-and-hold amplifiers that can be combined with the ADC either externally or internally. An analysis of the constraints on the dataacquisition system's dynamic performance should start with the sample-and-hold amplifier to be used.

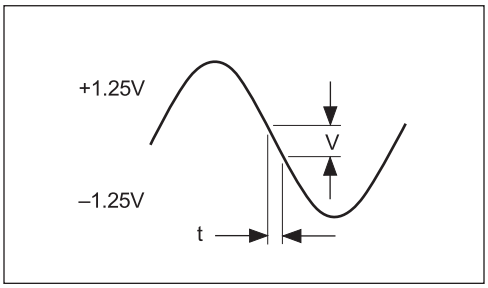

Figure 1. When a data converter digitizes a sinusoidal signal, Δt is the delay from the time that a designer wants to read the voltage to when the voltage was actually sampled. This period is called the aperture delay time. ΔV is the difference between the exact voltage that is digitized from the desired voltage.

(a)

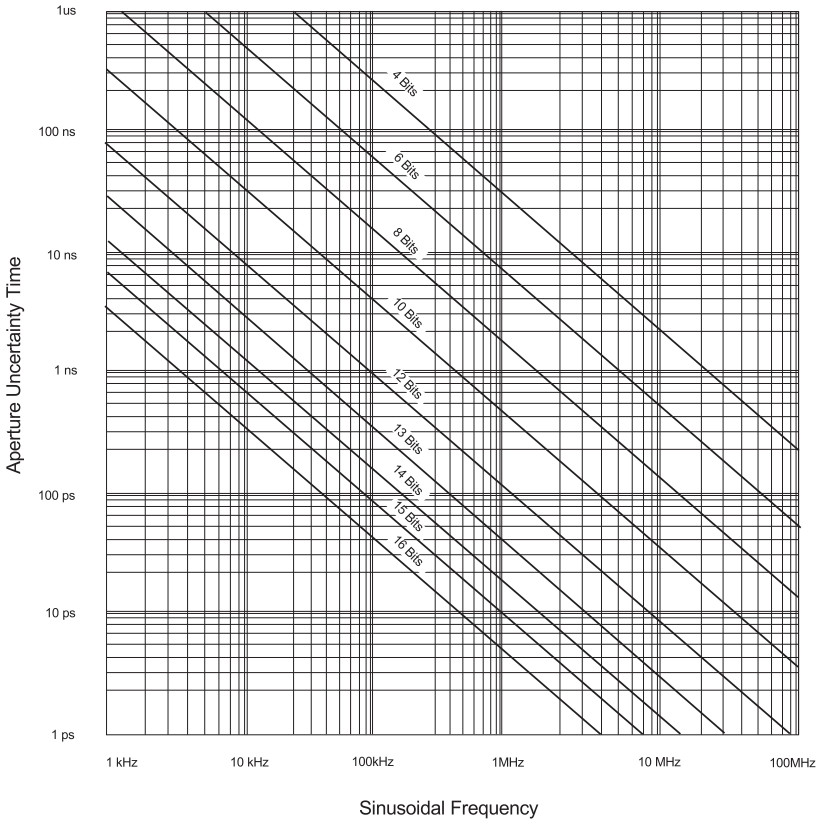

(b)

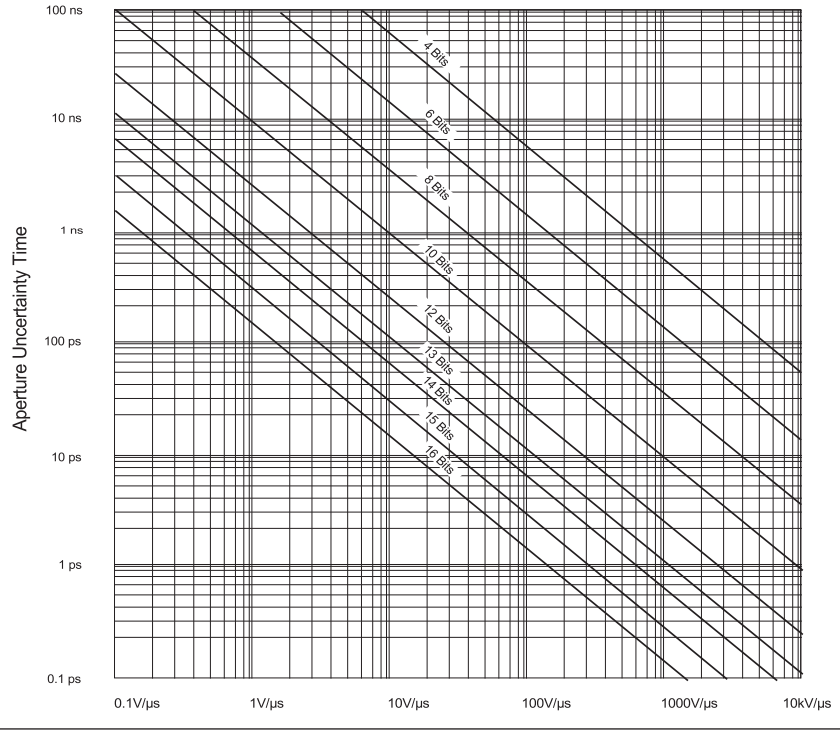

Figure 2. These graphs quickly identify what aperture-uncertainty specification is required to keep up with a particular input frequency or slew rate. Assuming a 10-V pk-pk full-scale signal, the first graph depicts the aperture time required to digitize a particular input frequency to an accuracy of 1 LSB (a). For the same set of conditions, the other graph shows the relationship between aperture time and slew rate (b).

Sinusoidal Signals

To ensure appropriate accuracy at a particular frequency, sinusoidal signals must be explored. Assume that the 2.5-V pkpk sinusoidal signal will be digitized by a data converter to 12- bit accuracy (Fig. 1). When designers want the ADC or sampleand- hold amplifier in front of the converter to take a sample, timing delays occur. As a result, the exact point that was originally intended isn't acquired.

Δt represents the delay from the time that designers want to read the voltage to when the voltage was actually sampled. This aperture-delay time appears on the data sheets for SHAs and for flash ADCs, where appropriate. The amplifier's aperture-delay time is affected by the logic threshold voltage on the pin that controls going into the sample or hold mode.

For TTL or other logic families, there's an internal delay in the digital logic going from a logic 1 (assume a sample-mode) to a logic 0 (assume a hold-mode) and the time the internal switch actually opens. The sample-and-hold amplifier also experiences an analog delay as the input goes through the input buffer to be stored on the hold capacitor. The difference between the digital and analog delays is known as the "effective" aperture delay time.

The aperture-delay time prevents designers from digitizing the sinusoid at exactly the right moment. The exact voltage that's digitized differs from the desired voltage by ΔV (Fig. 1, again). For a 12-bit converter with a ±1.25-V input (1.25-V full scale, 2.5-V full-scale range), 1 LSB equates to:

$2.5\ V/2^{12} = 2.5\ V/4096 = 1\ LSB = 610\ \mu V$

Maintaining 12-bit, ±1-LSB accuracy requires that during the aperture-delay time (Δt), the voltage change (ΔV) is held to ±1 LSB maximum. The change in voltage over the change in time (ΔV/Δt) is the maximum slew rate of the input signal allowed. Mathematically, formula 1 is derived as:

$$\text{Slew rate} = \Delta V/\Delta t = 2\pi fV$$

This yields the maximum input frequency of the sinusoid, which can be digitized as formula 2:

$$f = \Delta V/(\Delta t 2\pi V)$$

where f = maximum input frequency
V = full-scale voltage

In practice, substituting the aperture-delay specification into the formulas doesn't allow a very fast slew rate or input frequency before accuracy errors occur.

However, in digitizing sinusoidal signals, the aperture delay is repeatable from sample to sample. Upon performing a discrete Fourier transform (DFT) to define the frequency, the aperture delay doesn't affect identifying the input frequency. The aperture delay now just represents itself as a phase shift. However, designers interested in capturing single-shot events would be limited in accuracy by the "effective" aperture delay. In contrast, the maximum sinusoidal frequency that can be acquired isn't limited by the repeatable aperture delay times. Now, the input-frequency limiting factor is the uncertainty of the delay from sample to sample.

This uncertainty can be derived from the effect that various noise sources create when changing the exact threshold voltage of the sample-and-hold amplifier's digital control command. The aperture uncertainty (jitter) now becomes the limiting factor (Δt) in determining the maximum input frequency. A digital-signal-processing application on repetitive sinusoidal signals is one area whose ultimate performance and upper-frequency range depends upon the aperture-uncertainty specification. For example, the upper-frequency range of a spectrum analyzer becomes limited by the aperture-uncertainty specification.

Keep in mind that unpredictable events are digitized once they're detected. Such events include seismic or biological applications, which digitize an error in the signal based upon the aperture-delay specification. Graphing formulas 1 and 2 can quickly identify the aperture-uncertainty specification required to keep up with a particular input frequency or slew rate (Figs. 2a and 2b). To ensure that each sample has the desired accuracy, the aperture uncertainty should be a peak value vs. root mean square (rms).

Frequency Domain

It's essential to determine whether the SHA or the flash ADC can meet the input frequency and accuracy requirements based on aperture-uncertainty or slew-rate considerations. Although they're prerequisites, these specifications are insufficient in judging the performance to be obtained. Designers must also consider other dynamic specifications.

For example, because the ADS-130, 12-bit, 10MHz sampling ADC combines an ADC and SHA in one package, it allows dynamic specifications that are often reserved only for system instruments to be published (see the table). As a result, sampling ADCs are easing the manufacturer's ability to specify data converters in the frequency domain. If a standalone ADC had these specifications and a companion sampleand- hold amplifier were still required, there's no guarantee that an overall system would still meet these specifications.

Two common and important ADC specifications are integral and differential nonlinearity (see the table, again). What is uncommon, however, is that they're specified for Nyquist operation (5MHz input frequency with a 10MHz sampling rate). Errors in these specifications at a high input frequency appear as errors in the frequency-domain specifications.

A fast Fourier transform (FFT) plot is the basis for determining the frequency domain specifications, such as total harmonic distortion (THD), spurious-free dynamic range, signal-to-noise and distortion ratio (SINAD), and signal-tonoise ratio without distortion. These specifications are given at a particular input frequency and full-scale range (FSR) of –0.5 dB below full-scale (or about 95% of the input range). Beware of what these specifications might mean if the foregoing conditions aren't listed. The two-tone intermodulation distortion tests require two input frequencies (tones) to be generated for testing vs. the single fundamental for the other tests.

An FFT plot is an excellent way to identify unknown input frequencies (Fig. 3). Actually, a discrete Fourier transform (DFT) is performed on the discontinuous data from the data converters. Various windowing functions then ensure that the DFT's fast Fourier transform algorithm is for data representing complete cycles when noncoherent sampling is used. Examples of windowing functions include the Blackman-Harris or Hamming windows, whose functions are to prevent the signal energy from "leaking" into other frequencies.

If a continuous bandwidth-limited signal contains no frequency components higher than the analog input frequency (F_{IN}), then the original signal can be recovered without distortion if it's sampled (F_S) at more than twice the analog

input frequency (Nyquist rate). Because of imperfect filters, which go to zero asymptotically and introduce phase distortion, the sampling rate might be a minimum of four times the analog input in practice.

If the sampling frequency (F_S) isn't high enough or if harmonics are generated by the input or data converter, part of the spectrum folds over into the original signal spectrum (Fig. 4). This undesirable effect, called frequency folding, creates an alias frequency in the original signal spectrum as given in the following example:

The input frequency is analyzed as $F_{IN} = K(F_S/2) + \Delta F$, where:
(a) K is an odd or even integer that's a multiple of half the sampling period.
(b) $0 < \Delta F < Fs/2$

Case 1. Alias frequency = $(Fs/2) - \Delta F$ if K is odd.
Example: If Fin= 270 Hz and Fs = 100 Hz, then 270 Hz = 5 (100/2) + 20, where K = 5

K is odd so that 270 Hz has an alias appearing at (100/2) – 20 = 30 Hz in the original signal spectrum

Case 2. Alias frequency = ΔF if K is even
Example: if Fin = 820 Hz and Fs = 100 Hz, then 320 Hz = 6(100/2) + 20 where K = 6

K is even so that 320 Hz has an alias appearing at 20 Hz in the original spectrum.

Designers should also be careful when it comes to signal-tonoise ratios (SNRs). When inspecting the FFT plot, the classic definition defines the signal as the fundamental frequency. Noise is defined as all of the other unwanted errors (harmonics, spurious frequencies, and the noise floor) in the FFT plot. In practice, the term SNR may not always include the harmonics.

In specifying ADCs, manufacturers now use the term SINAD, for signal-to-noise and distortion ratio. Or they note if the signal-to-noise ratio includes or excludes distortion. After mathematical derivation through Gaussian quantization errors, the ideal SNR and distortion specification (rms signal to rms noise) is expressed in formula 3 as:

SNR = 6.02n dB + 1.76 dB, where n is the number of bits

For harmonics, the specification THD includes all harmonics. The in-band harmonic specification is meant to include the worst harmonic (usually the second one). The spurious-free dynamic range resembles the in-band harmonic specification. In practice, the spurious occurrences and noise present are lower than the harmonics. Total harmonic distortion, the ratio of the sum of all the harmonics to the fundamental signal as analyzed at the ADC output can be expressed in formula 4 as:

$$THD_{rms} = 20\log \sqrt{\left[10^{\frac{\text{2nd Harmonic}}{20}}\right]^2 + \left[10^{\frac{\text{3rd Harmonic}}{20}}\right]^2 + \ldots}$$

While total harmonic distortion is a more conservative specification, the in-band harmonic and spurious-free dynamic- range specifications can be of interest for particular applications. For example, insufficient usable dynamic range that enables noise and harmonics to intrude on the measurement could interfere in distinguishing a low-level radar signal. In the effective-bits specification, a mathematical representation of an ideal sine wave is compared with the sinusoid digitized by the data converter. Numerical method algorithms help compute the accuracy, which is specified in bits.

Basically, the effective-bit specification is a combination of all the other errors. Any errors from differential and integral nonlinearity, aperture uncertainty, and missing codes show up as part of the overall rms error. In what "no missing codes" means to ADCs and "monotonicity" to DACs, effective bits will likely be the leading comparative specification for the dynamic performance of sampling ADCs.

Effective bits can be derived from the SNR (and distortion) specification in formula 5 as:

Effective bits = (SNR –1.76)/6.02

If the signal-to-noise ratio without distortion is employed, then the THD specification can be used to derive the effective bits in formula 6 as:

$$\text{Effective bits} = \{[-20\log 10\ \overline{10^{(-SNR/10)} \ldots}$$

$$\overline{\ldots + 10^{(THD/10)}}] - 1.76\}/6.02$$

The specification of input bandwidth is also important. Most dynamic specifications already discussed are based upon sinusoids. It's certainly commendable that a data converter can operate at its Nyquist rate for sinusoidal inputs. The data converter would then digitize minus full-scale signals to plus full-scale signals on alternative samples.

Some applications, however, encounter a more difficult situation, such as when channels are switched from a multiplexed input. In this case, the sample-and-hold amplifier doesn't get to track the input all along. This is a tougher test of its input bandwidth than other, less-demanding applications.

Similarly, fast-slewing transient conditions need wide input bandwidth and short sample-and-hold acquisition time to acquire the signal. A wide-input-bandwidth data converter also minimizes phase shifts on the signal being digitized.

As mentioned, ADC specifications have different significance with respect to various applications. The input bandwidth, for example, becomes important in avoiding phase shifts between inphase and quadrature channels in a radar system. Imaging systems, whether based on a chargecoupled device or on infrared, step changes from one pixel to the next. This resembles multichannel signals being switched through a multiplexer. In an imaging application, the edge between a black object on a white background simulates this step condition.

Fast acquisition times and wide input bandwidth ensure performance for these applications. Some applications require that the harmonics, spurious frequencies, and noise not be as large as the lowest signal to be digitized. Spurious-free dynamic range becomes important to these users. In practice, the inband harmonics specification yields essentially the same information because the worst harmonic is usually greater than the noise floor and spurious frequencies.

Digital storage oscilloscopes (DSOs) take advantage of the aliasing of a data converter and its sampling rate to digitize repetitive frequencies outside the normal range of a data converter. Here, the effective-bits specification yields the DSO's overall accuracy at particular frequencies. Communication systems may have multiple frequencies (tones) and their interaction is of interest. In this case, specifications of two-tone (two-frequency) intermodulation distortion give the needed insight into ADC requirements.

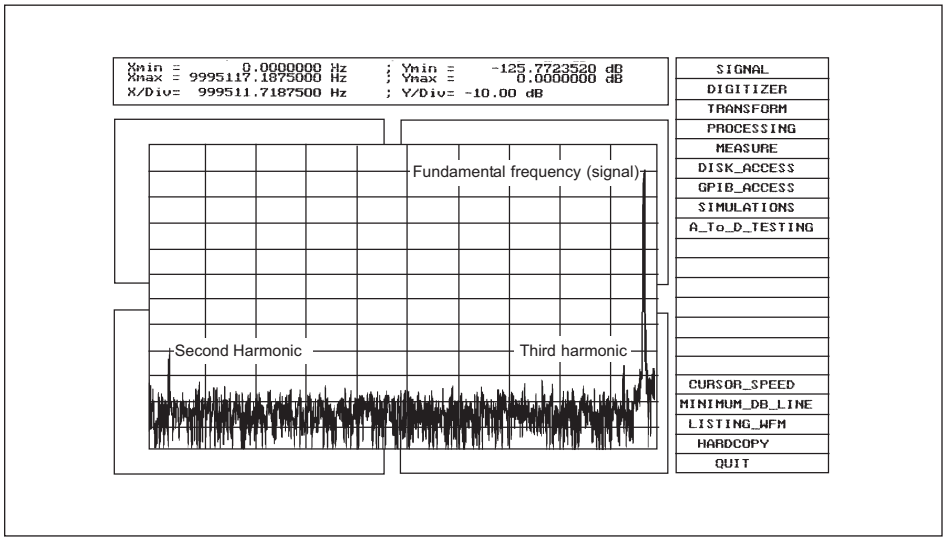

Figure 3. A Fast Fourier Transform plot is a handy method for viewing unknown input frequencies or distortions introduced by a data converter. This plot is based on data from an ADS-130 ADC.

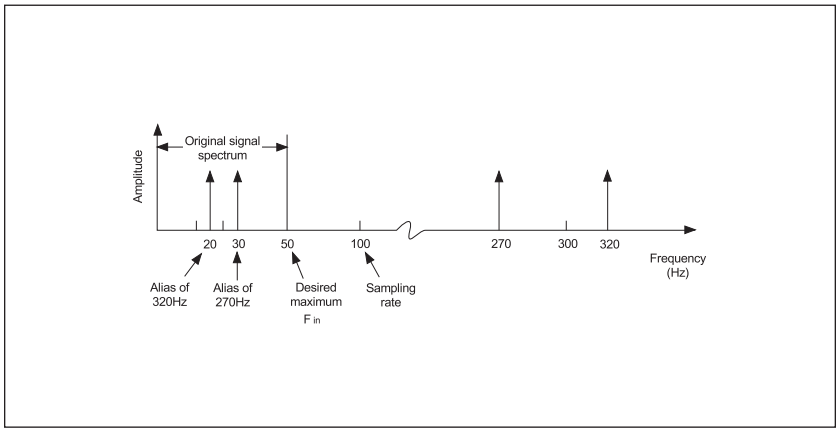

Figure 4. When aliasing occurs, frequency components exceeding the nyquist criteria are folded back into the desired frequency spectrum. In this example, 270 Hz and 320 Hz frequencies appear at 30 Hz and 20 Hz, respectively.

ADS-130 Dynamic Specifications

	Minimum	Type	Maximum	Units
Integral nonlinearity				
@fin=5MHz, +25°C			±1	LSB
0 to +70°C			±1	LSB
–55°C to +125°C			±2	LSB
Differential nonlinearity				
@fin=5MHz, +25°C			±1	LSB
0 to +70°C			±1	LSB
–55°C to +125°C			±2	LSB
Total harmonic distortion (–0.5 dB)				
Dc to 500kHz	–68	–70		dB below FS
500kHz to 2.5MHz	–65	–67		dB below FS
2.5MHz to 5MHz	–65	–67		dB below FS
In-band harmonics				
Dc to 500kHz	–69	–70		dB below FS
500kHz to 2.5MHz	–66	–67		dB below FS
2.5MHz to 5MHz	–66	–67		dB below FS
Spurious-free dynamic range				
Dc to 500kHz	–69	–70		dB below FS
500kHz to 2.5MHz	–66	–67		dB below FS
2.5MHz to 5MHz	–66	–67		dB below FS
Signal-to-noise ratio (without distortion, –0.5dB)				
Dc to 500kHz	–67	–70		dB below FS
500kHz to 2.5MHz	–65	–69		dB below FS
2.5MHz to 5MHz	–65	–69		dB below FS
Signal-to-noise ratio and distortion, –0.5dB				
Dc to 500kHz	–65	–66		dB below FS
500kHz to 2.5MHz	–63	–65		dB below FS
2.5MHz to 5MHz	–63	–65		dB below FS
Effective bits, –0.5dB				
Dc to 500kHz	10.6	11.0		bits
500kHz to 2.5MHz	10.2	10.5		bits
2.5MHz to 5MHz	10.0	10.2		bits
Two-tone intermodulation				
Distortion (fin=2.2MHz, 2.3MHZ, fs=8MHz, –0.5dB)	–72	–75		dB below FS
Input Bandwidth				
Small signal (–20dB input)	50	65		MHz
Large signal (0dB input)	30	40		MHz
Slew rate	175	200		v/µs
Aperture delay time		5	7	ns
Aperture uncertainty		5	7	ps
Sample and hold acquisition time to 0.01% (2.5V step)		30	50	ns
FS=full scale				