Siglent SDS1104X-E Review

The graphs above show input noise for all timebase settings from 1ms/div to 1ns/div at maximum vertical sensitivity of 500µV/div. We can see a continuous 1/f characteristic for the peak to peak noise amplitude. In contrast, the RMS reading is fairly constant and clearly doesn't tell the full story.

Trigger

The trigger is one of the most important features of any oscilloscope, since we can only analyze a signal that we are able to trigger in a reliable and stable manner. Consequently, this section deals with the trigger system and its various options in the Siglent SDS1104X-E, which has a modern fully digital trigger system. Expectations are high, because triggering has been excellent on previous Siglent X-series DSOs.

Trigger Coupling

This gives us the opportunity of filtering the sampled input signal data (by means of digital signal processing) before the actual trigger condition is checked. If this is done in a proper way that suits the application, we can often get a stable trigger even on complex and/or noisy signals. There are four choices of trigger signal coupling and consequently filtering:

- DC: The entire input frequency range – no filtering.
- AC: High-pass filter with a lower corner frequency of about 8Hz
- LFRJ (LF-Reject): High-pass filter with a corner frequency of about 2MHz
- HFRJ (HF-Reject): Low-pass filter with a corner frequency of about 2.2MHz

The corner frequency of AC coupling can be determined by the phase shift between trigger point and signal.

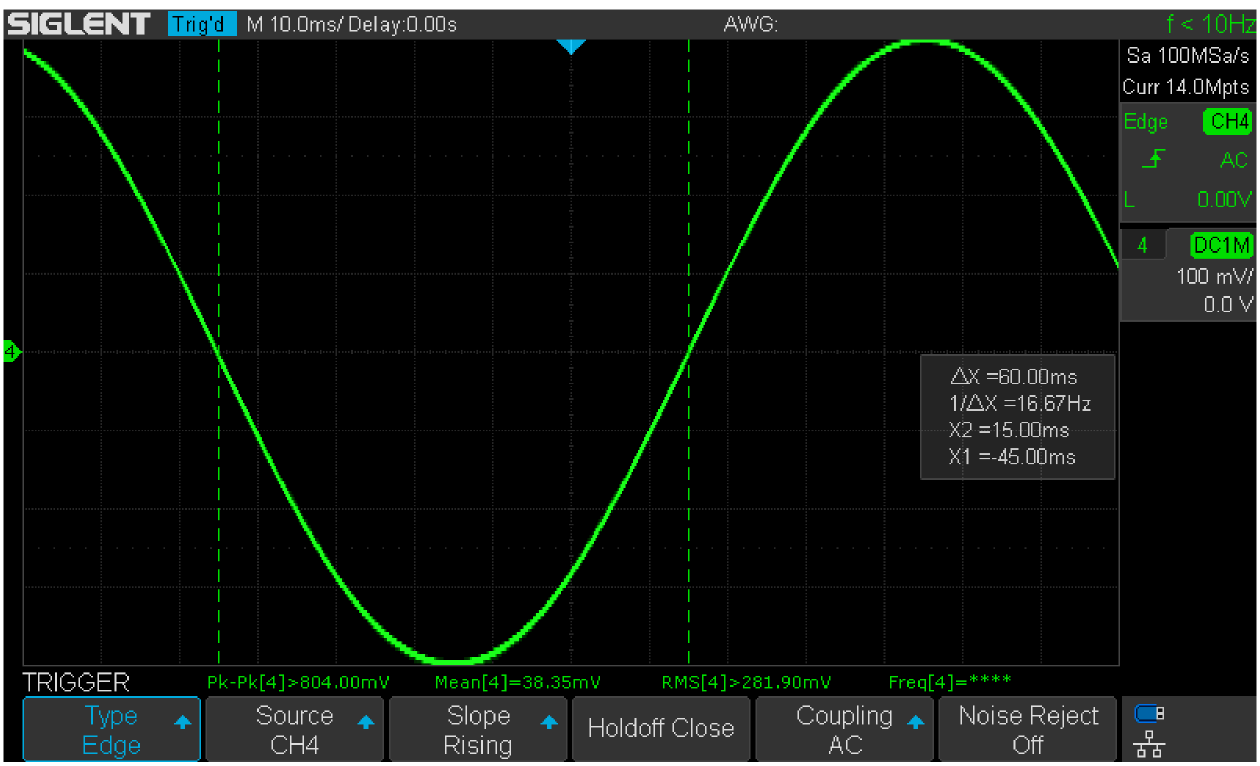

SDS1104X-E_Trig_AC_Corner

Cursor measurement is used to determine the phase shift. The difference between both cursors X2-X1 is half the period of the input signal, hence its frequency is 1/(2 x 0.06) = 1/0.12 = 8.33Hz. The easier way is to just divide the frequency calculation of the cursor measurement by two: 16.67Hz/2 = 8.33Hz.

We have 45deg phase shift when the trigger reference point sits at precisely 75% of the distance between X1 and X2. Since X1 – Ref. = 45.00ms and X2 – Ref. = 15.00ms and 3 x 15.00 = 45.00, this condition is perfectly met in the screenshot above.

For LF/HF-Reject trigger coupling, the corner frequencies cannot be determined in the same way, as apparently a higher order filter is implemented. Instead, a reliable triggering on a 1div peak-peak signal has been used as criterion in order to get the numbers listed above.

AC Trigger Coupling

Most of the time we use DC coupling for the trigger, and there is not much to tell about it, other than it works just as it should. We rather want to examine AC trigger coupling now.

Why and when would we need AC coupling for the trigger at all? Usually, we make that choice for the channel input and if we select AC coupling there, the trigger will inevitably be AC coupled as well. So there we already have the answer – we have the opportunity to force the trigger into AC coupling, even when the corresponding input channel is DC coupled. This can be useful for AC signals that have a DC offset that we want to watch on the screen. The offset might change with time and we still don't want to lose triggering.
AC trigger coupling does not display a trigger level indicator, which I'm not happy with, but Siglent don't seem to be willing to change that.

The following test uses a 200mVpp 1MHz sine wave that is superimposed on a 600mVpp symmetrical ramp signal at 100mHz, which acts as a variable DC offset here. As if this weren't enough, this signal has a fixed offset of +1V on top of that, which needs to be removed by means of the vertical position control.

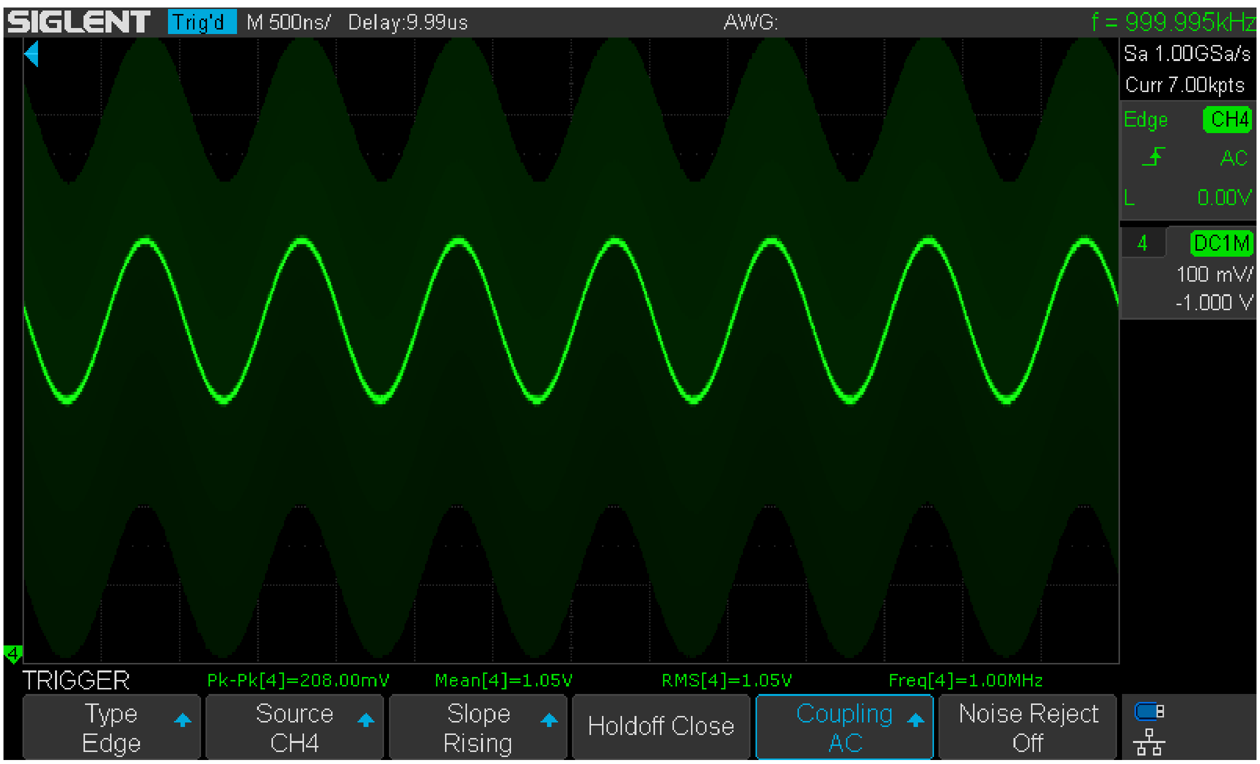

SDS1104X-E_Trig_AC

With DC trigger coupling, triggering would only occur about 1/3 of the total time in this scenario and even then the horizontal position would not be stable because of the ever changing trigger level (relative to the AC portion of the input signal).

It's totally different if we use AC trigger coupling. The trigger level is set to zero which is actually always equivalent to the mean level of a symmetrical input signal. With this, triggering occurs always at the same point on the X-axis, no matter what the DC offset or low frequency instantaneous signal level is. The waveform constantly changes its vertical position on the screen, but remains stable on the time axis.

Persistence mode has been used for the screenshot above, to indicate the slow vertical movement of the signal.

Trigger Sensitivity

Trigger sensitivity in Edge mode depends on the Noise Reject setting (hysteresis), but also the trigger level, whether it is automatically set (by pushing the trigger level knob) or manually tweaked. Here are the sensitivities expressed in vertical divisions for all possible combinations:

- Noise Reject off, manual trigger level: **0.32** div.
- Noise Reject off, automatic trigger level: 0.60 div.
- Noise Reject on, manual trigger level: 0.88 div.
- Noise Reject on, automatic trigger level: 1.60 div.

The following screenshots demonstrate all four settings listed above.

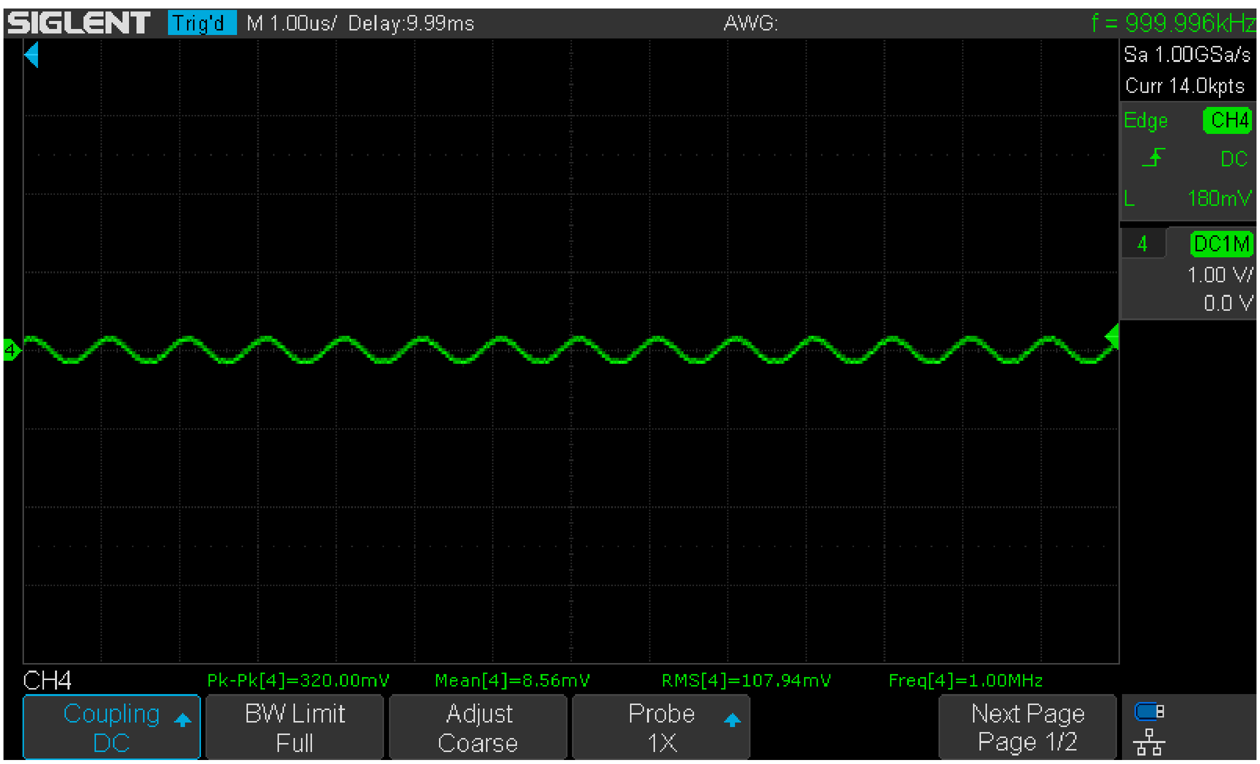

SDS1104X-E_Trig_32%_ML

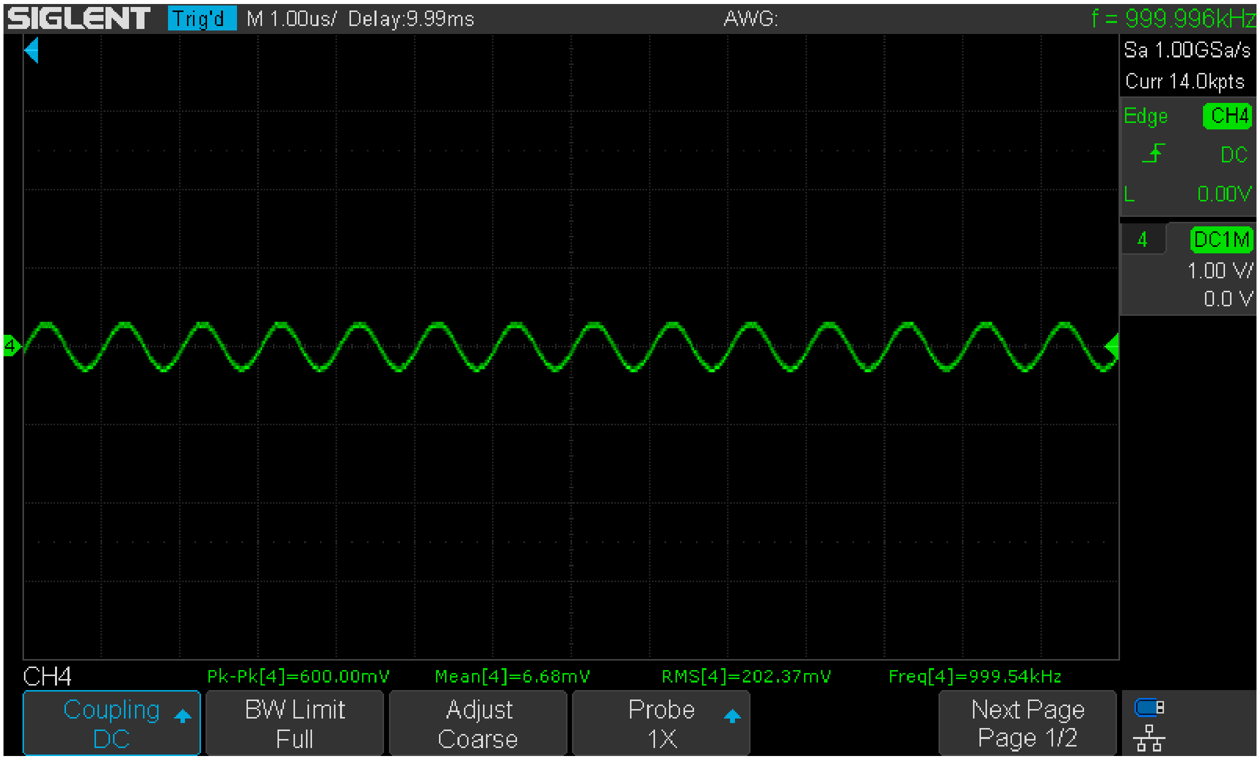

SDS1104X-E_Trig_60%_AL

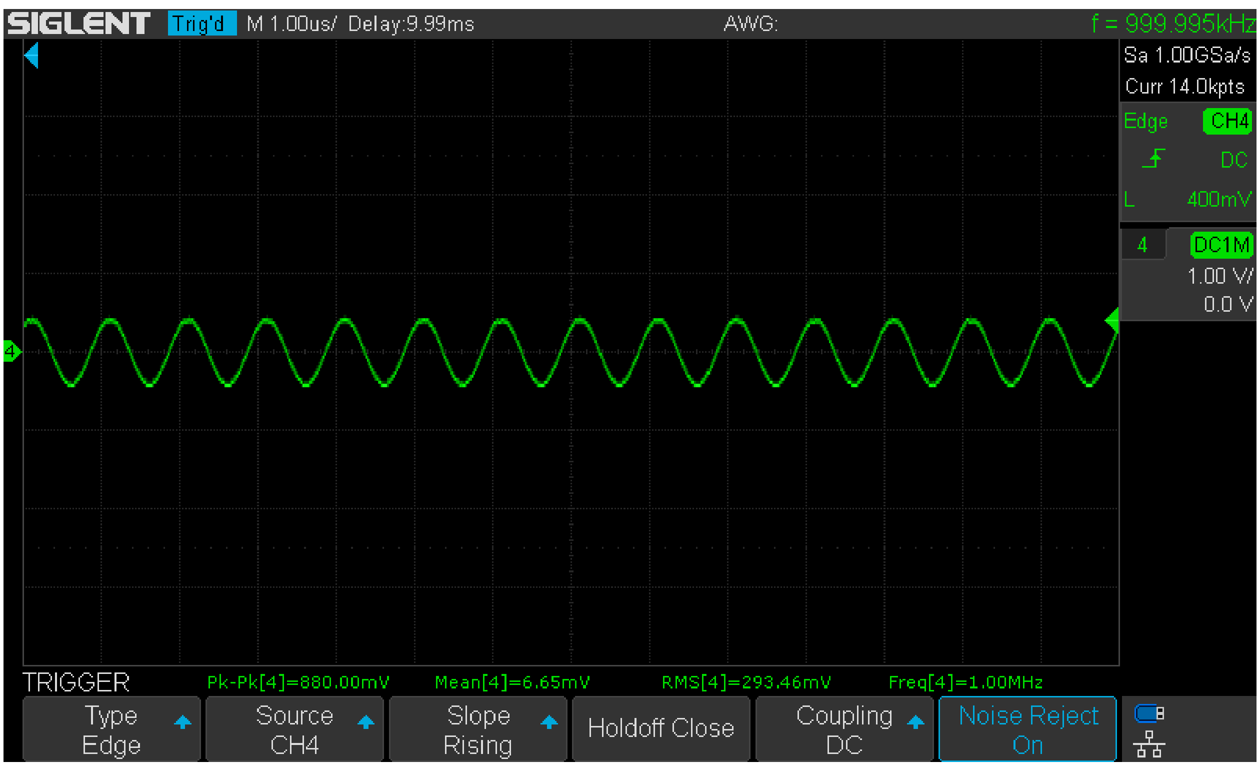

SDS1104X-E_Trig_NR_88%_ML

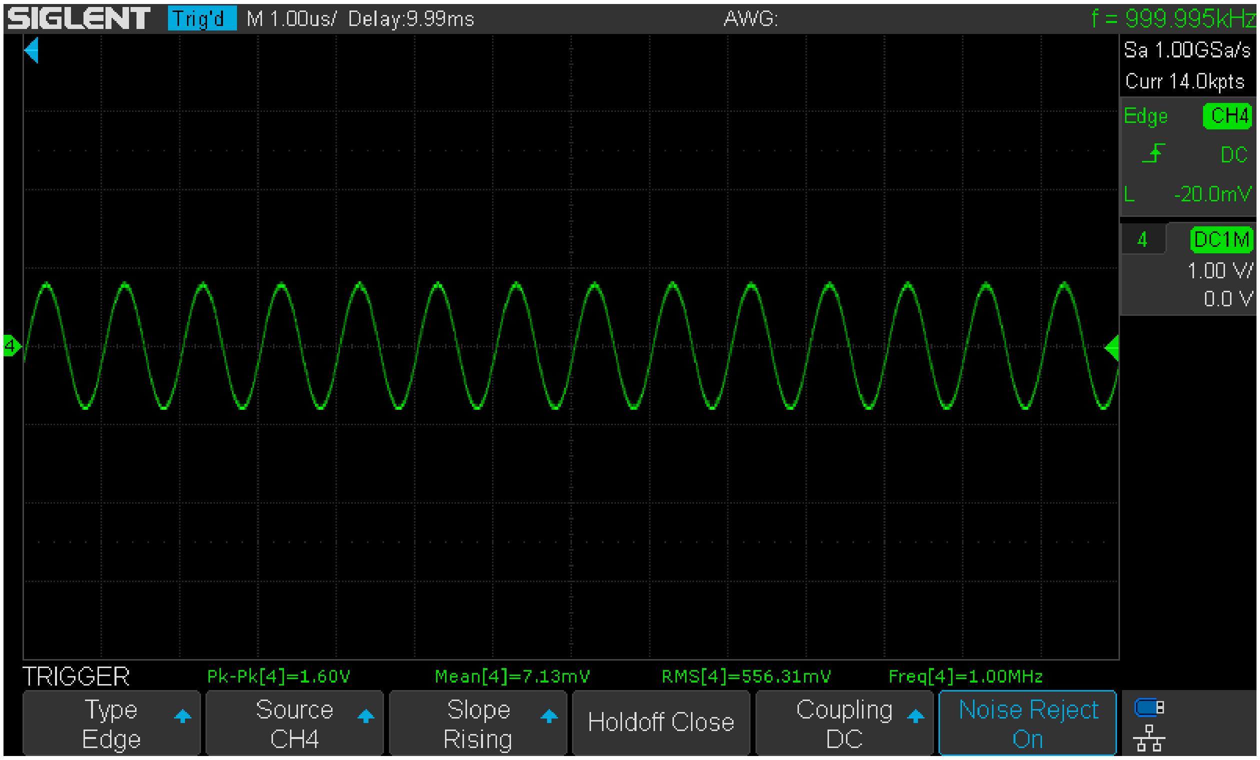

SDS1104X-E_Trig_NR_160%_AL

Trigger Stability

Another interesting aspect of a trigger system is stability, i.e. the absence of excessive jitter, as well as synchronous triggering on all channels so that accurate time measurements are possible.

Let's start with the stability of the trigger point which actually means stability of the trigger level and in case of a digital trigger – as it is implemented here – also the proper interpolation of samples to determine the exact trigger point.

A very accurate and stable 100MHz signal derived from an OCXO is fed into channel 4 of the SDS1104X-E and the fastest timebase of 1ns/div is used together with its associated maximum time delay of 9.99µs. Generally, the trigger delay is limited to 9999 times the time base setting. Persistence mode with 10s decay time has been enabled in order to make any signal variation clearly visible on the screenshot.

Please note that the actual trigger point is far outside the visible screen area by some 714 screen widths and the delay has been fine adjusted so that the falling edge is exactly at the centre of the screen, while the original trigger is on the rising edge.

As can be seen, there is not the slightest sign of jitter in either direction. The signal trace is just rock-stable and automatic measurement has even accurately determined the signal frequency.

SDS1104X-E_Trig_1ns_D10us

The second test looks at the skew between channels, which of course should ideally be nonexistent, i.e. zero. A stable 100MHz signal was used to feed all four channels in parallel through a 4-way power splitter and 4 coaxial RG58 cables with the exact same length of 1 meter. No external input termination was used for the scope channels this time, as I didn't have four identical ones at hand. I certainly do not recommend feeding signals into the scope this way, let alone at high frequencies like this, but for this particular test it is acceptable and even so the result leaves nothing to be desired.

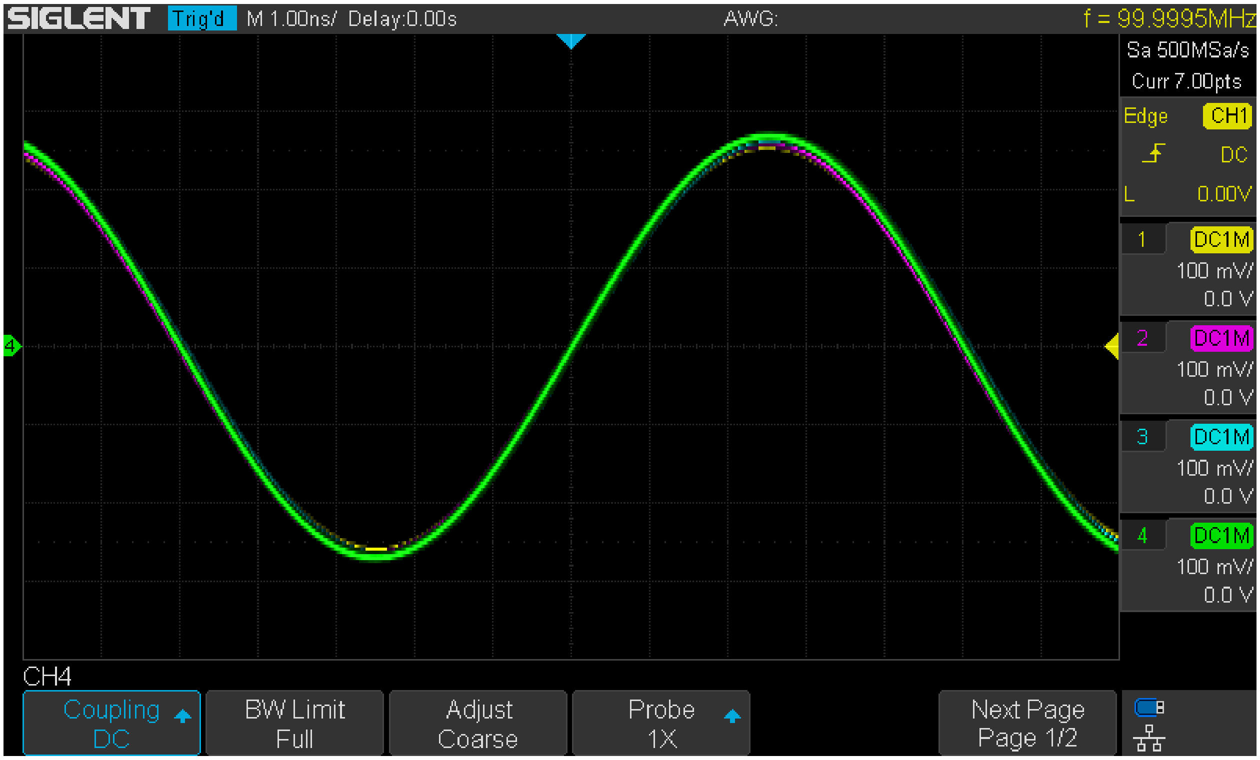

SDS1104X-E_Trig_1ns_Skew

At the fastest timebase of 1ns/div, we can see the small differences in the frequency response of the individual channels, as the amplitudes are slightly different, but there is barely any skew visible between channels. In any case, possible deviations are well below ±100ps and I don't have the means to guarantee that all four signals actually arrive at the scope inputs at the exact same picosecond anyway. However, this is a very pleasing result.

Triggering Noisy Signals

Triggering isn't always straightforward; complex signals, crosstalk and noise often require additional measures to get a stable signal trace that can be properly analyzed.

First decision we have to make is the proper trigger coupling. As has been stated earlier, this means limiting the trigger bandwidth in a way that suits the application. For this purpose, LF and HF reject trigger coupling are most important.

- LFRJ (LF-Reject): High-pass filter with a corner frequency of about 2MHz
- HFRJ (HF-Reject): Low-pass filter with a corner frequency of about 2.2MHz

Another means for aiding stable triggering on complex signals with a known period is the Hold-off time. This simply inhibits triggering for a certain customizable time window.

Finally, the trigger level hysteresis is a critical parameter as well. A higher hysteresis can help to deal with fuzzy signals, but this also reduces trigger sensitivity significantly. Consequently, the Siglent SDS1104X-E provides two levels of trigger hysteresis, which can be toggled by means of the [**Noise Reject**] menu button.

Now let's have a look at some practical examples for the use of trigger coupling, noise reject and hold-off to help with the triggering in certain situations.

First a 1kHz sine wave with random spikes superimposed. Trigger conditions are ambiguous and regular DC coupled edge trigger initially does not yield a stable waveform display.

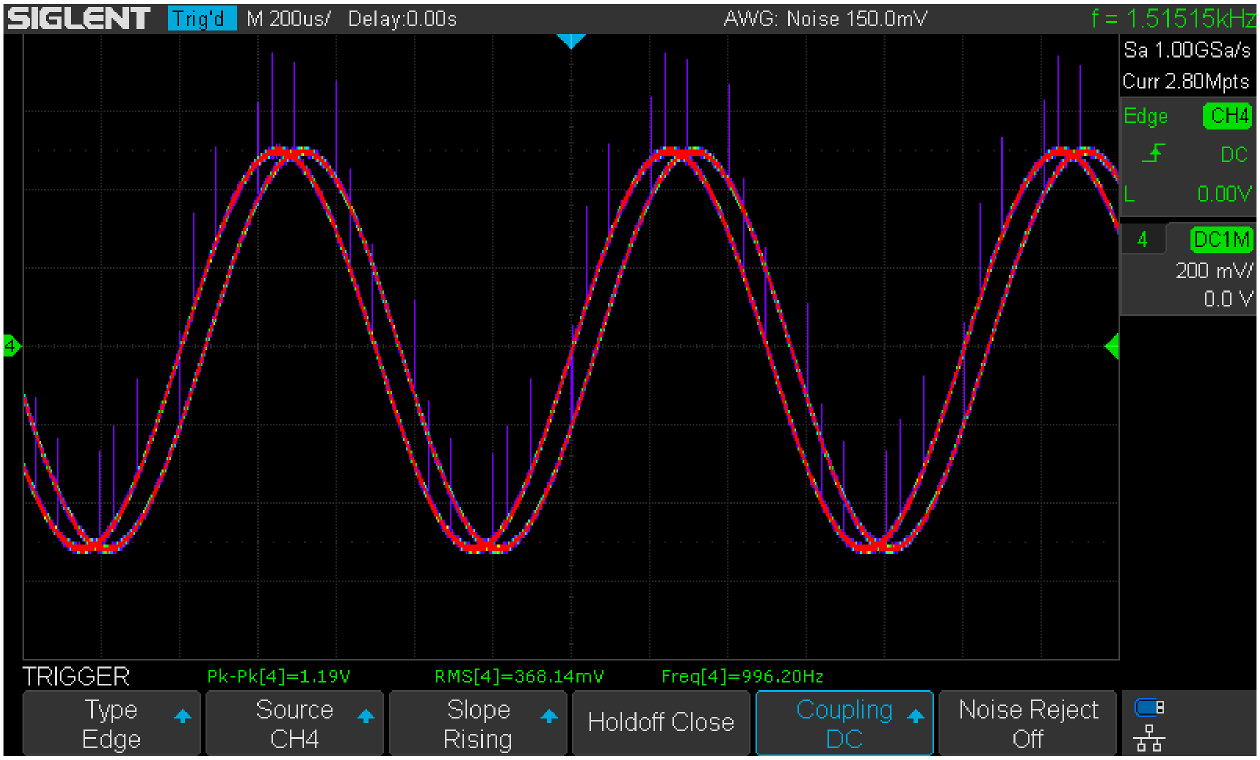

Trig_Spikes_DC

A hold-off time slightly less than the signal period works a treat.

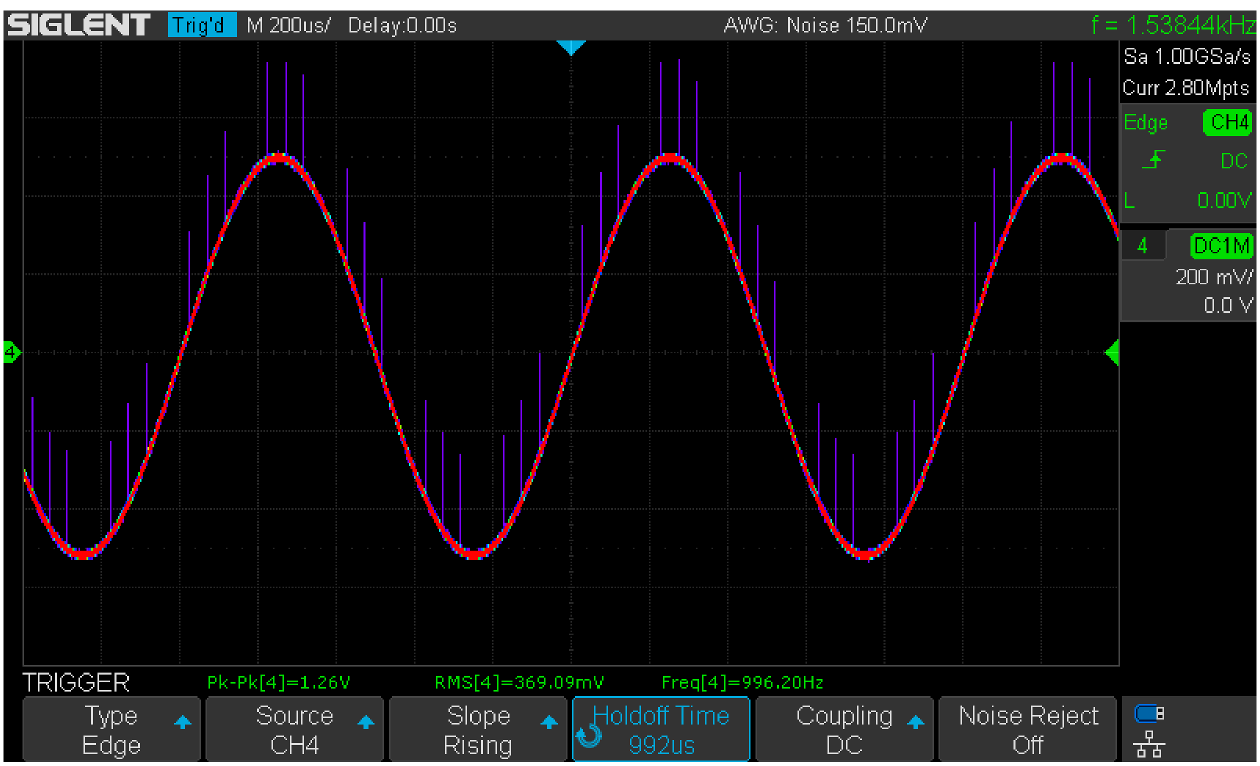

Trig_Spikes_DC_Holdoff

Alternatively, we can use HF-Reject trigger coupling (to ignore the spikes) to get a stable trigger.

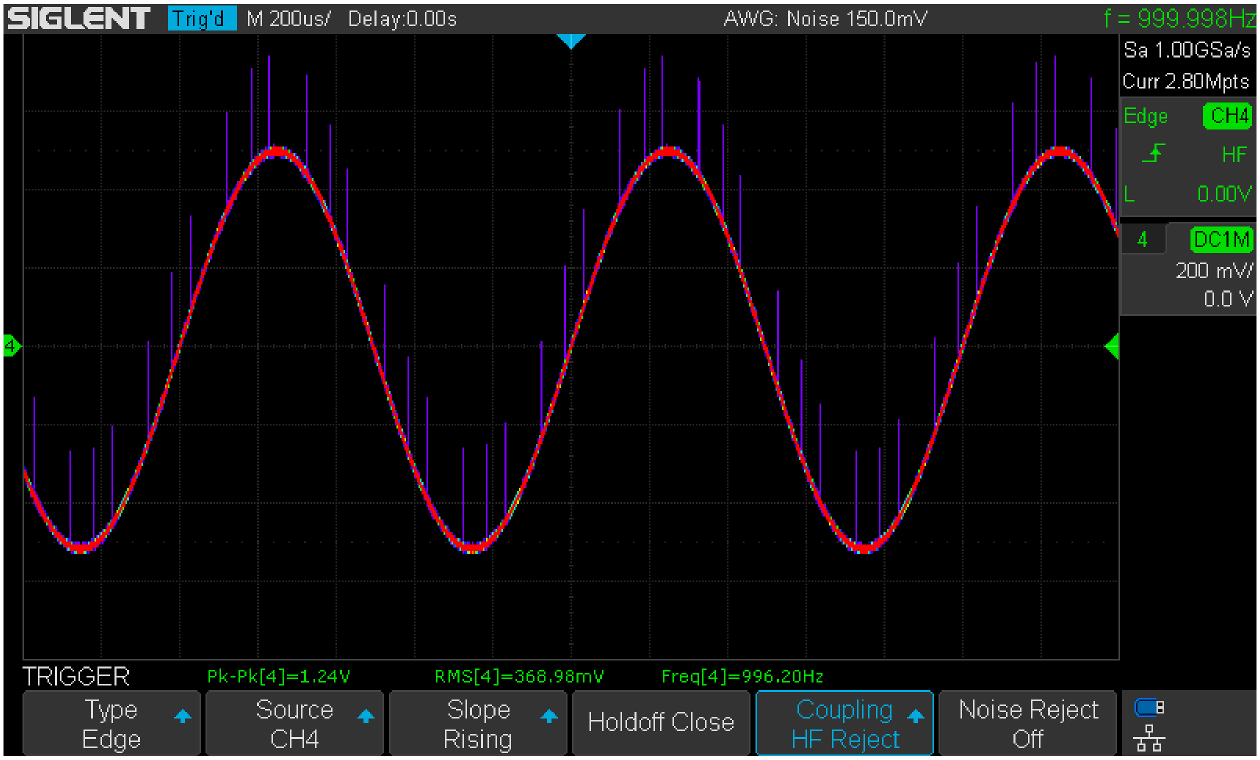

Trig_Spikes_HF_Reject

A similar example shows a very noisy 1MHz sine wave. Standard DC trigger does not work well at all.

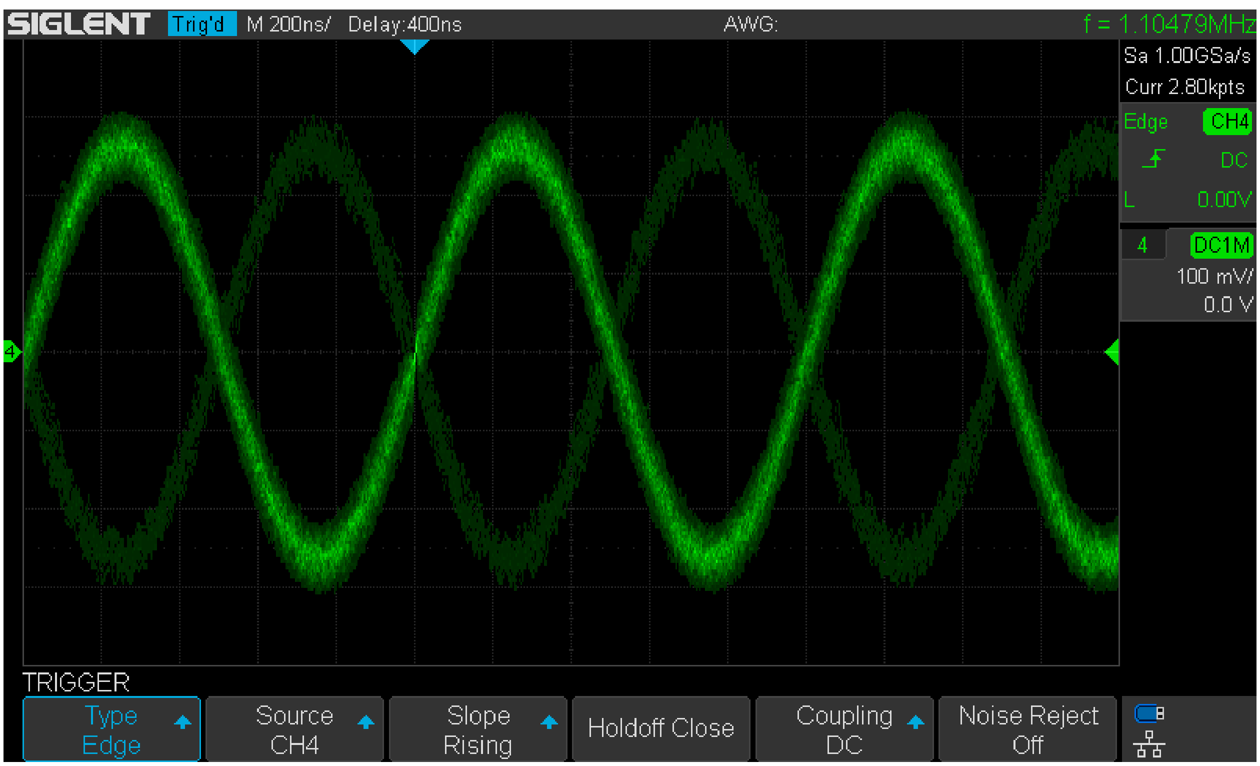

Trig_1MHz_Noise_-20dB_DC

Noise Reject does exactly what its name implies – problem solved.

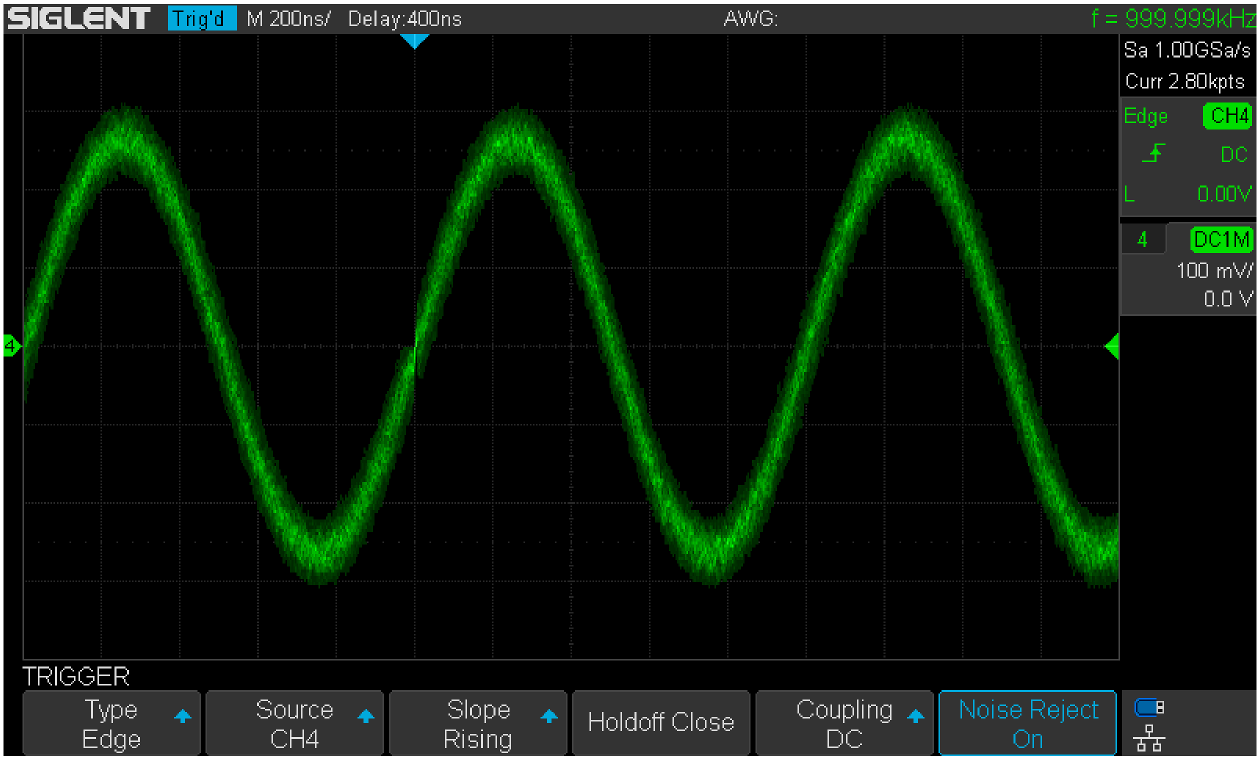

Trig_1MHz_Noise_-20dB_DC_NR

Once again, HF-Reject works just as well.

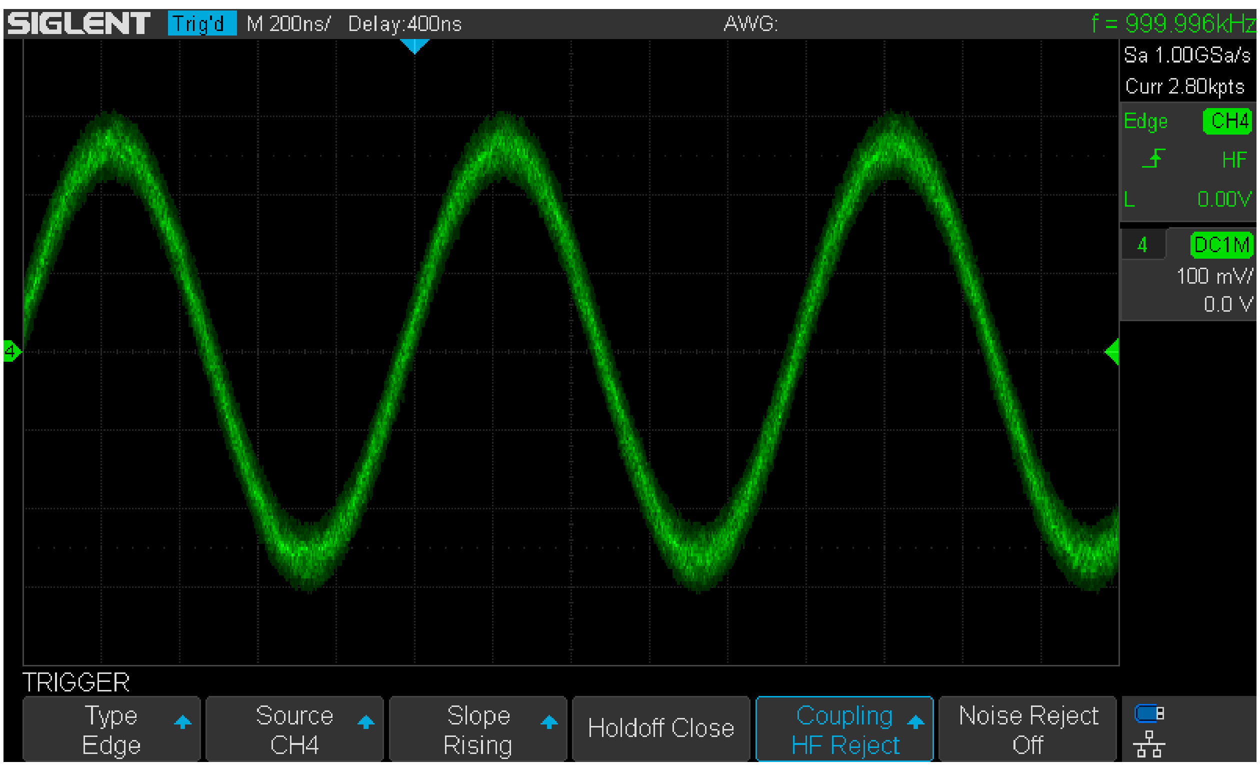

Trig_1MHz_Noise_-20dB_HFRej

Trigger Types

The SDS1104X-E comes with a fairly comprehensive set of advanced trigger types, each of them having a bunch of options, hence not all combinations could be tested. The following sections demonstrate the application of advanced triggers on complex signals, where it would be very difficult or even impossible to just use the default edge trigger.

Edge

Edge trigger is used throughout this review, even with its additional features like Hold-off and Noise Reject in order to deal with difficult signals, so no need to introduce it here again.

Slope

We have to define two threshold levels which are used to measure the slope (transition time) of an edge. Options are rising or falling slope, transition time greater or less than a customizable reference value, transition time in- or outside a customizable reference window and Noise Reject.

The example shows the zoomed view of a swept sine 500mVpp, 1kHz to 10kHz, where we want to trigger exactly on the transition from 10kHz to 1kHz, or more precisely, on the rising edge of the very first sine period at 1kHz. This is impossible with edge trigger, but a rising slope trigger set for transition times >=130µs within a ±100mV window does the trick.

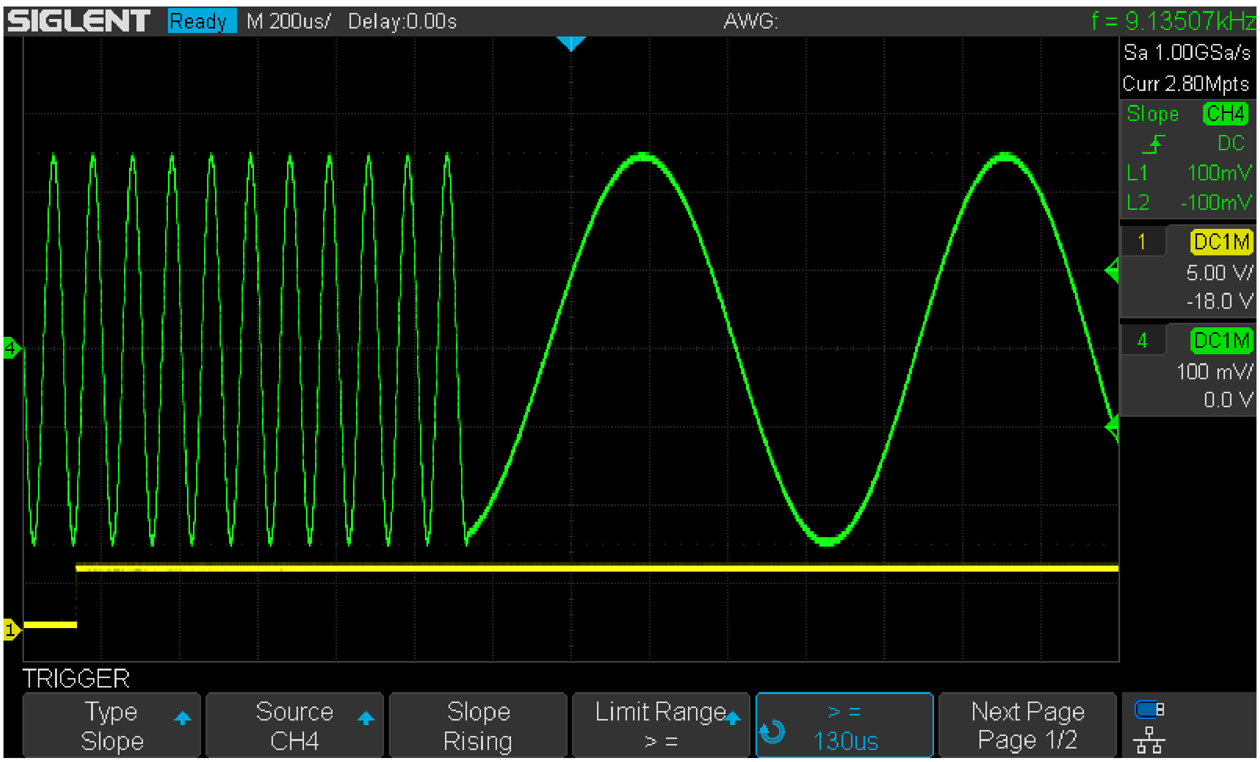

Trig_Slope

Pulse

Triggers on a certain pulse width range. Options are positive or negative pulse, pulse width greater or less than a customizable reference time, pulse width in- or outside a customizable reference time window and Noise Reject.

The example shows the zoomed view of a swept sine 500mVpp, 1kHz to 10kHz, where we want to trigger on the falling edge of the very first sine period at 1kHz. This is impossible with edge trigger, but a positive pulse trigger set for pulse widths >=300µs works quite nicely.

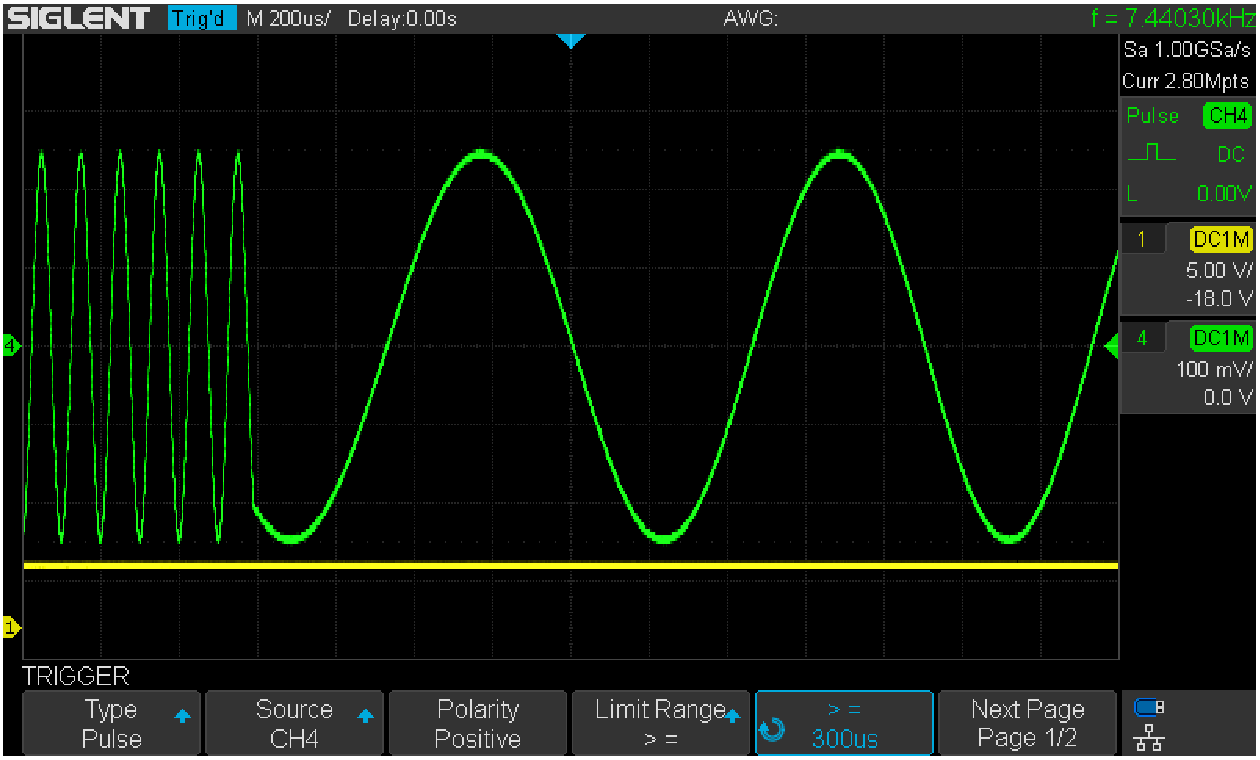

Trig_Pulse

Video

Video trigger is not often needed nowadays and I don't have any trigger source for that. It looks pretty comprehensive though. Please refer to the user manual for more details.

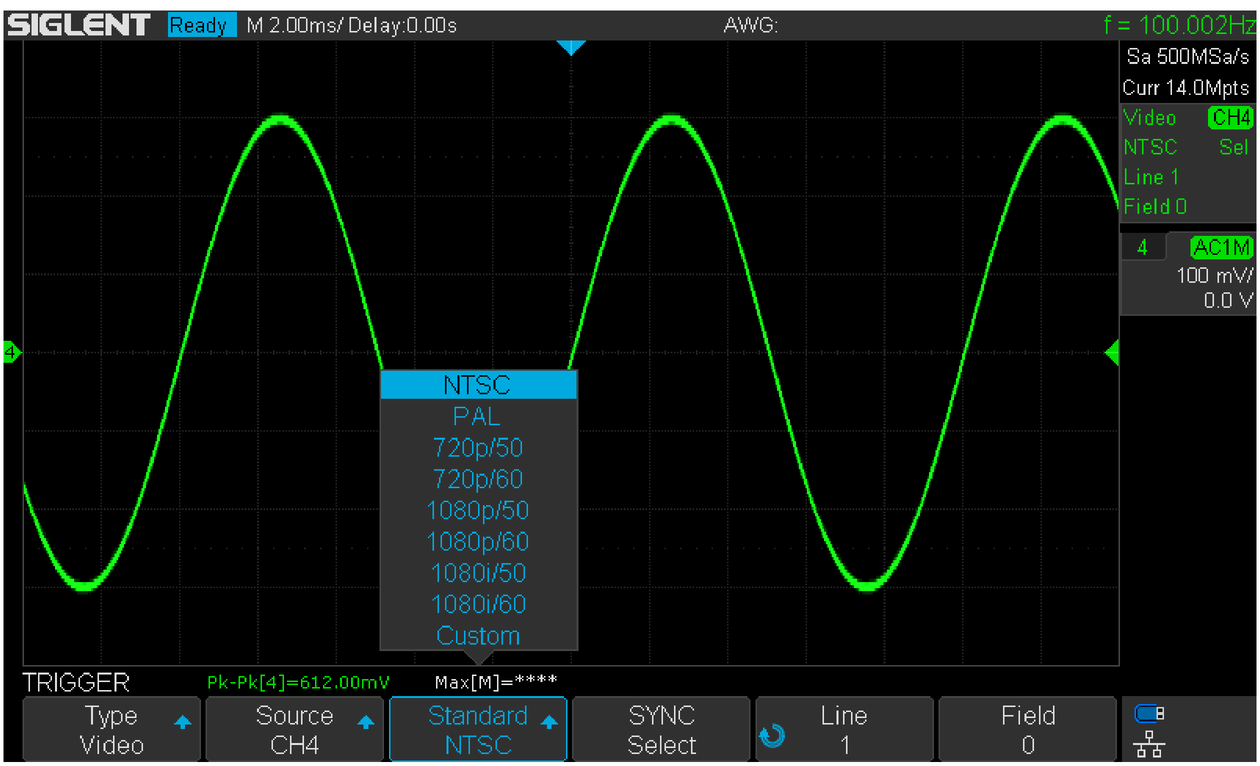

Trig_Video

Window

We have to specify two threshold levels to form a window that defines the valid signal amplitude range. Options are absolute or relative window definition and Noise Reject.

The example shows a random signal 500mVpp at 10kHz, where we want to trigger on the prominent peaks. This is perfectly possible with edge trigger for one polarity, but a window trigger set for a valid range of -220mV/+228mV works for positive and negative spikes at the same time.

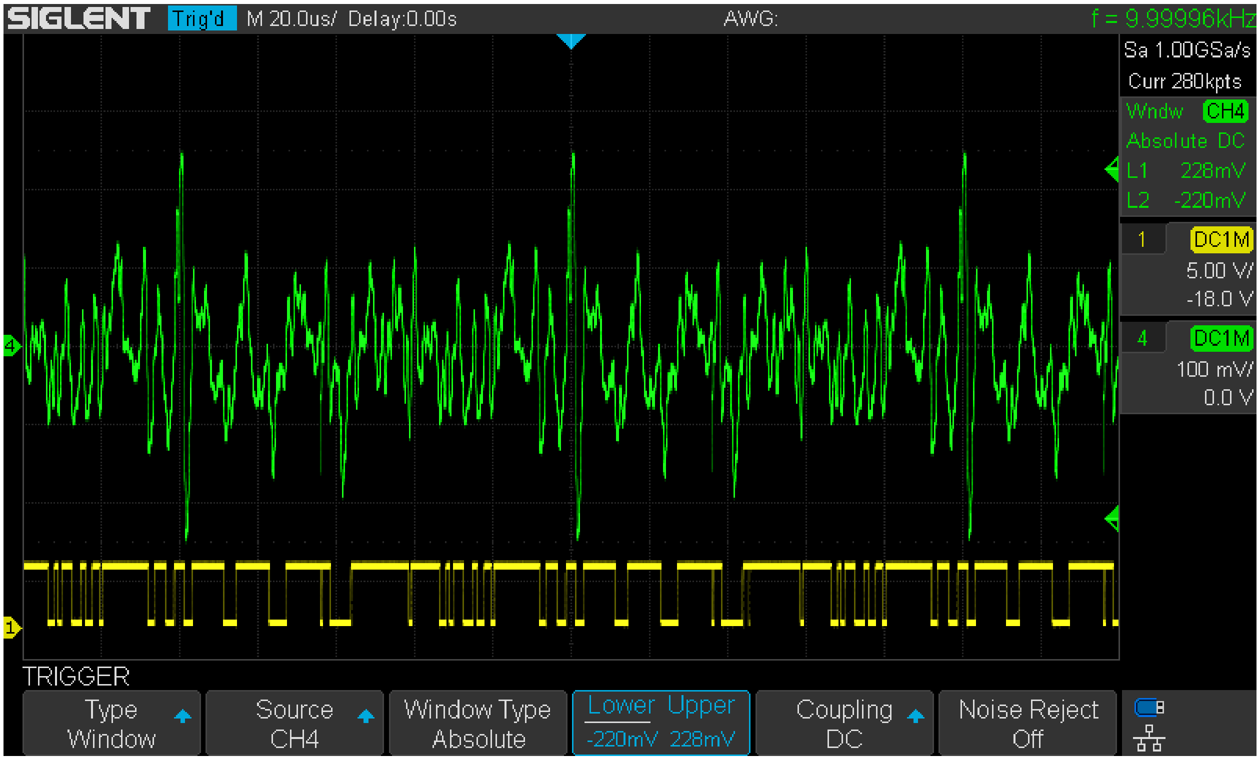

Trig_Window

Interval

This triggers on a certain signal period range. Options are rising or falling slope, signal period greater or less than a customizable reference time, signal period in- or outside a customizable reference time window and Noise Reject.

The example shows a swept sine 500mVpp, 1kHz to 10kHz, where we want to trigger exactly on the transition from 10kHz to 1kHz. This is impossible with edge trigger, but a rising slope interval trigger set for periods <=101µs does an almost perfect job.

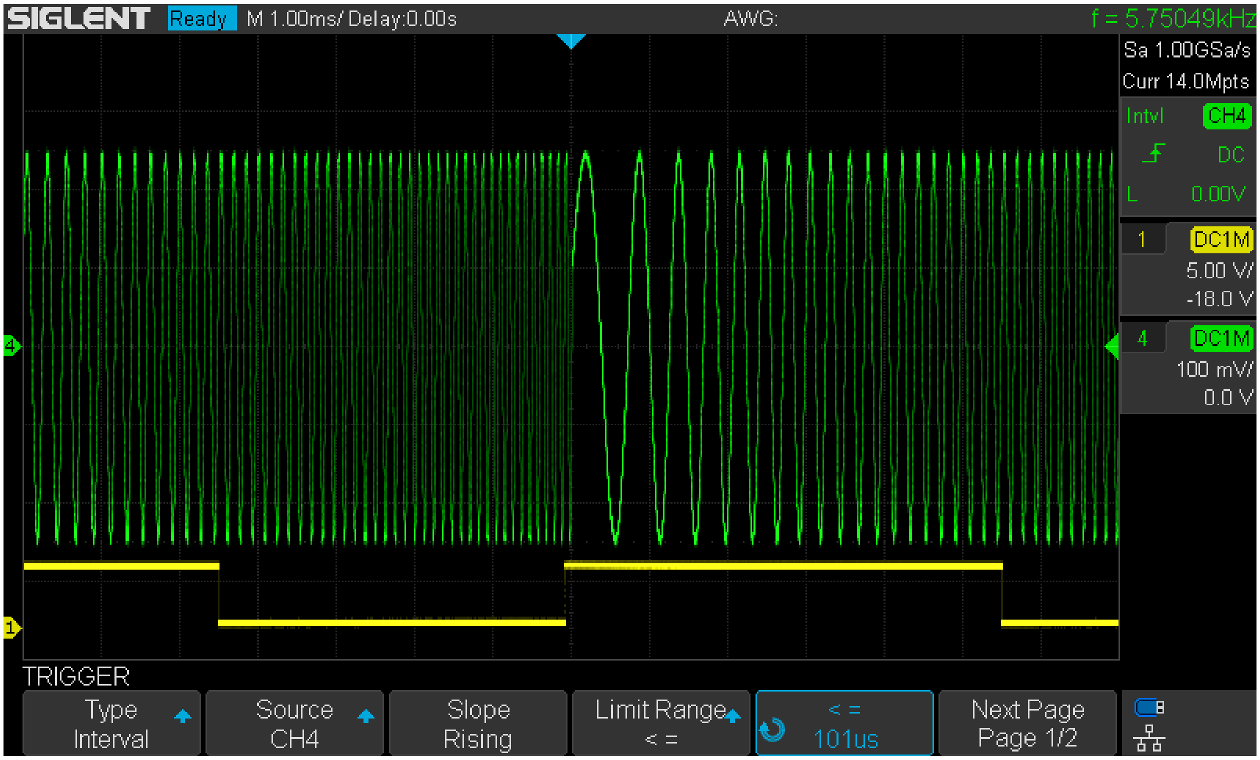

Trig_Interval

DropOut

Also known as Timeout trigger, and that's exactly what it does; it triggers when there is no signal transition within a specified time span. Options are rising or falling slope, timeout condition defined by edge or signal state and Noise Reject.

The example shows a 500mVpp 10kHz burst signal, where we want to trigger on the end of the burst packet. This is impossible with edge trigger, but a rising slope DropOut trigger set for a timeout >=101µs does just that.

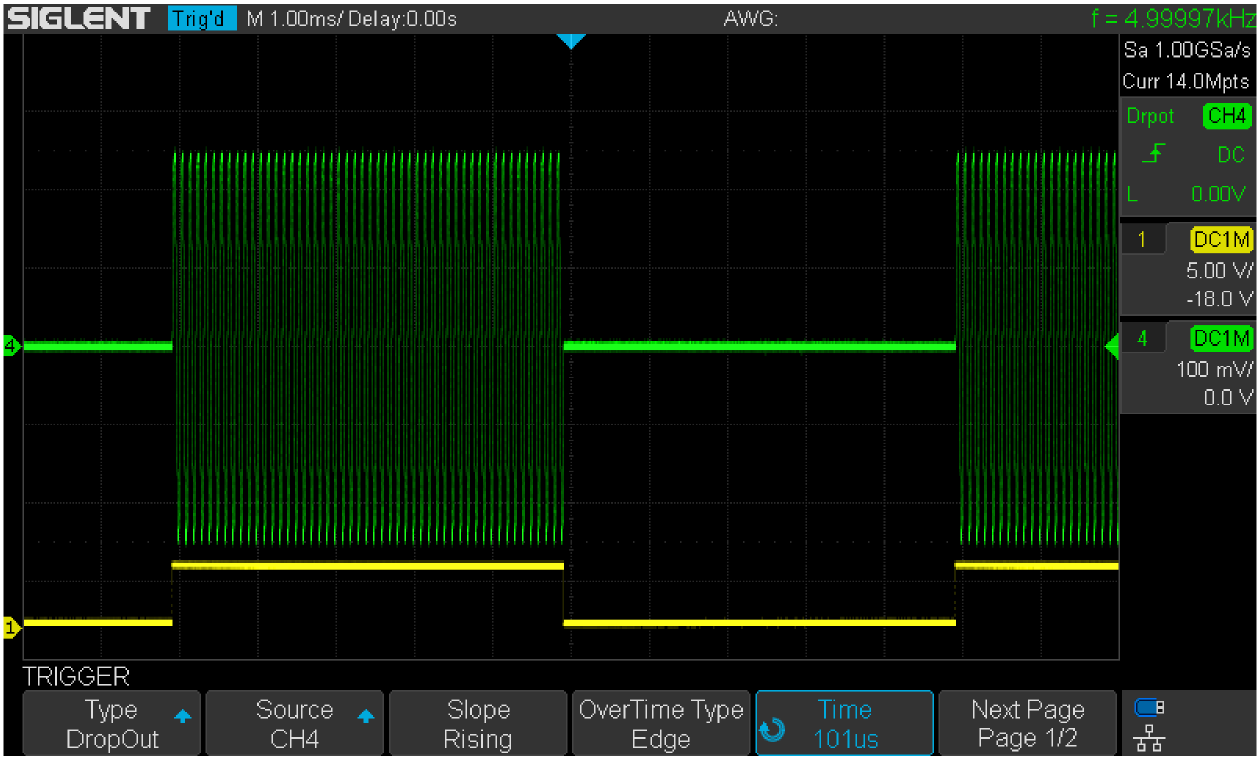

Trig_DropOut

Runt

We have to define two threshold levels. The trigger occurs whenever the signal crosses one of the thresholds twice without crossing the other one within a specified time window. In other words, runt pulse triggering can be limited to a certain pulse width range. Options are positive or negative polarity, pulse width greater or less than a customizable reference time, pulse width in- or outside a customizable reference time window and Noise Reject.

The example shows a pulse train 1Vpp at 1MHz that contains some runt pulses we want to trigger on. This is not possible with edge trigger, but a Runt trigger with the thresholds set at -100mV/+500mV works just fine.

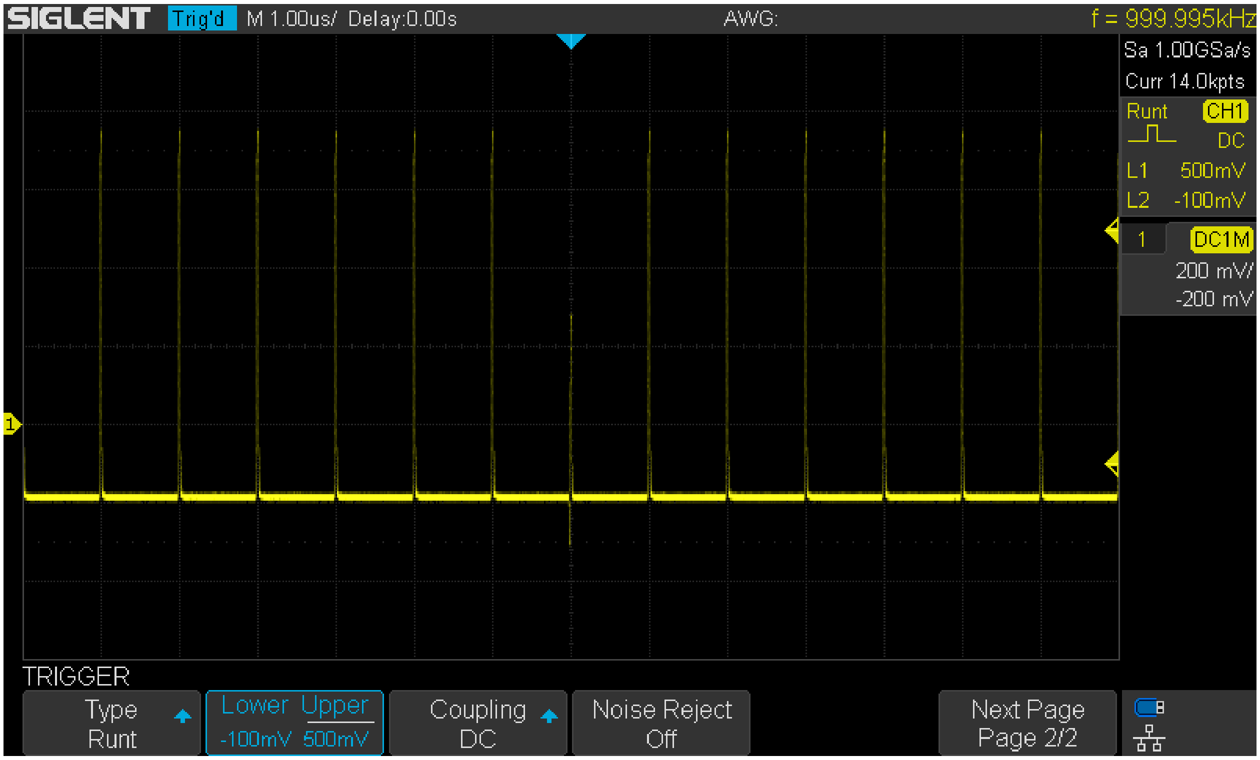

Trig_Runt

Pattern

Pattern trigger gives us the opportunity to define a simple logic AND, OR, NAND or NOR operation for triggering on up to four channels at the same time. Options are a separate threshold value for each channel, a minimum duration for the condition to become valid, and Hold-off.

The first example shows four pulse trains 2Vpp at 100kHz with some increasing delay from channel 1 to channel 4. We want to trigger on the last rising edge of all four channels when they remain above their respective thresholds for at least 10ns. Pattern trigger logic AND operation is used.

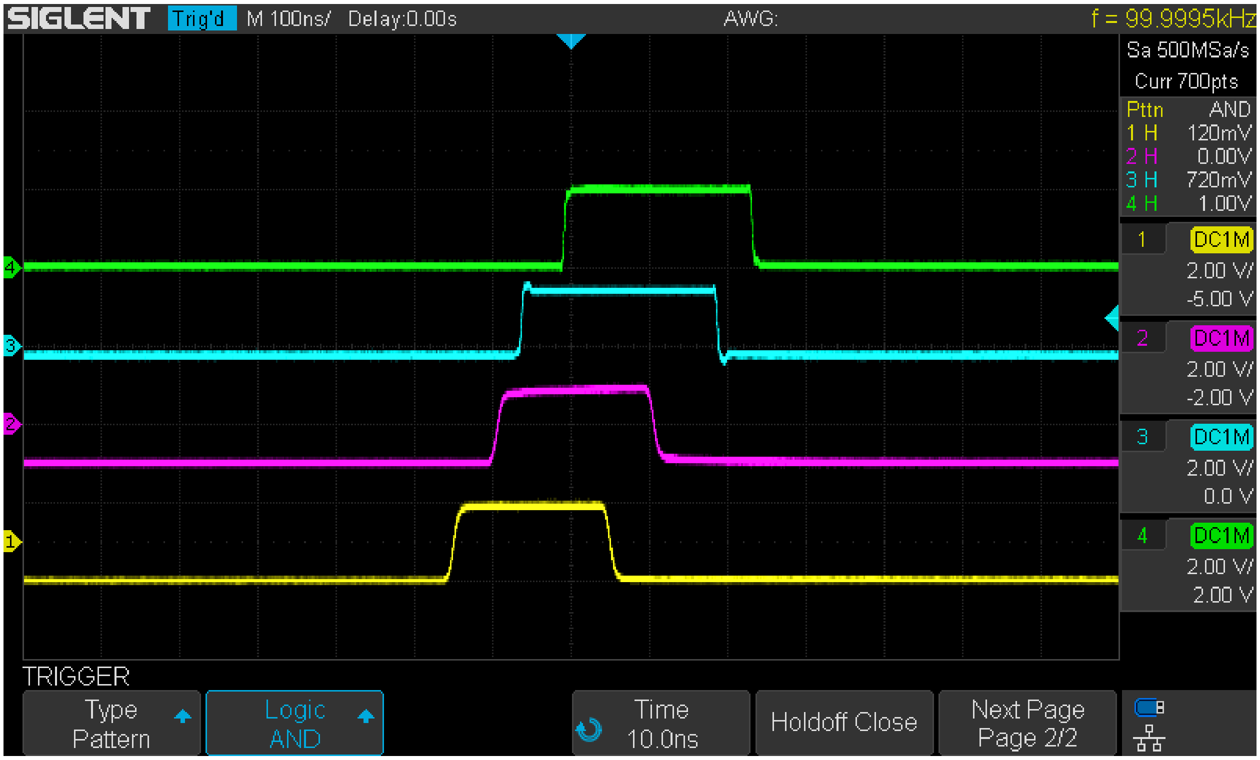

Trig_Pattern_AND

The second example shows four pulse trains 2Vpp at 100kHz with some increasing delay from channel1 to channel 4. We want to trigger on the first rising edge of all four channels when any channel remains above its respective threshold for at least 10ns. Pattern trigger logic OR operation is used.

Trig_Pattern_OR

The third example shows four pulse trains 2Vpp at 100kHz with some increasing delay from channel1 to channel 4. We want to trigger on the first falling edge of all four channels when any channel remains below its respective threshold for at least 10ns. Pattern trigger logic NAND operation is used.

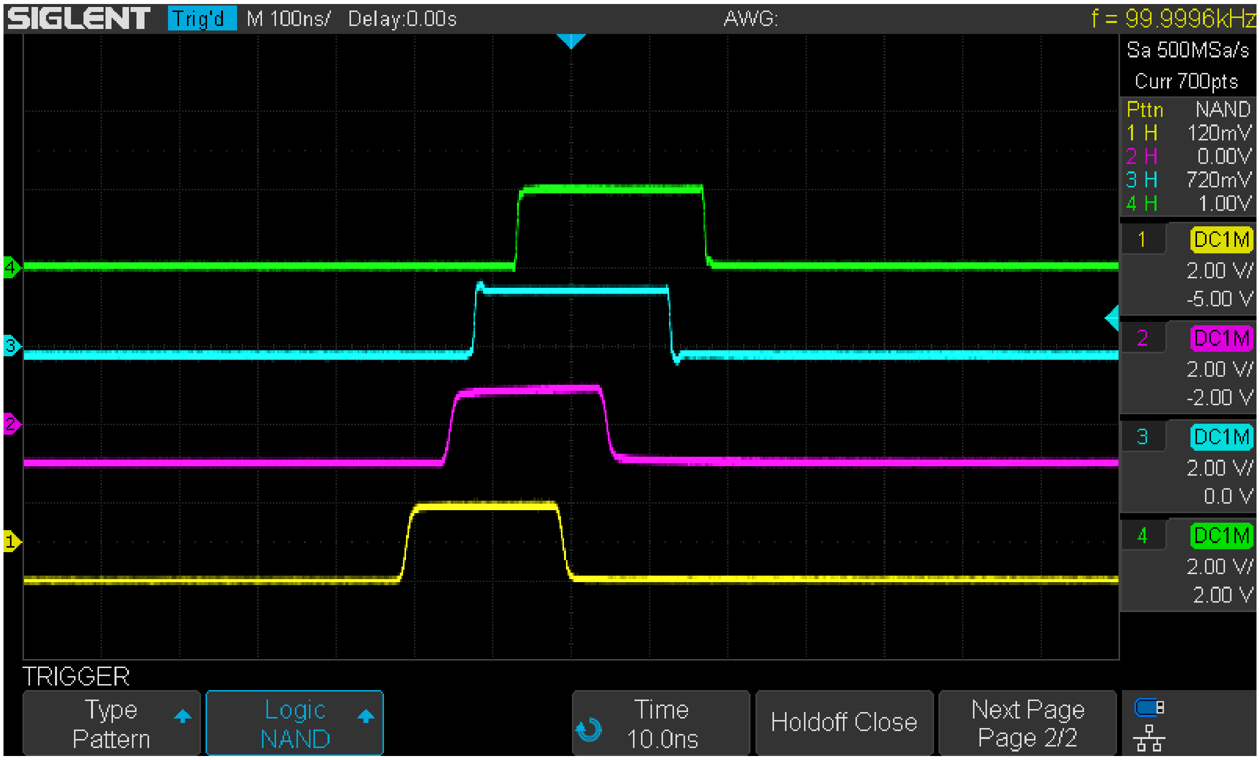

Trig_Pattern_NAND

The fourth example shows four pulse trains 2Vpp at 100kHz with some increasing delay from channel1 to channel 4. We want to trigger on the last falling edge of all four channels when they remain below their respective thresholds for at least 10ns. Pattern trigger logic NOR operation is used.

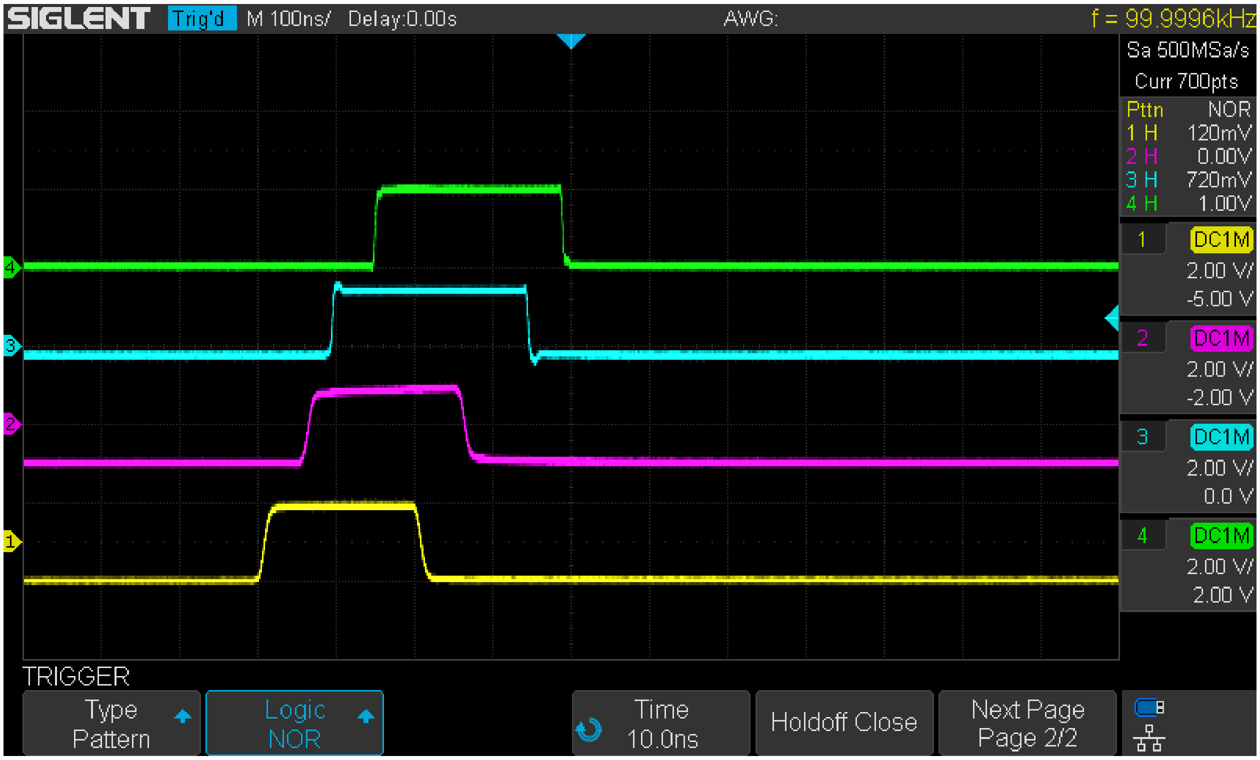

Trig_Pattern_NOR

Serial

Serial trigger will be dealt with later in this review, together with the serial decoders.

Trigger Frequency Counter

There are tasks where precise frequency measurements to at least 6 digits are essential. It is utterly convenient if a scope can do this, as it's usually just the signals we are monitoring with the oscilloscope where we also want to know their exact frequency. Back in the days of analog scopes, some offered a Y-output for the amplified signal of the current trigger channel and a frequency counter could be connected there – this was the ideal solution. Using the frequency counter alone was far less convenient, because it doesn't have a versatile input amplifier like the scope does and it is most desirable to be able to watch the waveform during the frequency measurement and vice versa without having to connect two instruments at the same test point.

Nowadays, most scopes don't have a Y-output and there are only very few exceptions. Unfortunately, Siglent X-series aren't among them. So we have to stick with the second best thing, i.e. the integrated trigger frequency counter. For the Siglent X-series, it has 6 digits resolution, but it still cannot replace a real frequency counter, yet might be good enough for most of the less demanding tasks.

To illustrate this, let's start with an accurate 10MHz input signal, derived from the OCXO. The automatic measurement with its limited 4-digit resolution does a nice job by displaying 10.00MHz. The trigger frequency counter has 6 digits resolution, but that reveals the cheap TCXO in the low cost SDS1104X-E, because the frequency reading is off by 4ppm. That's still not too bad and can certainly compete with several cheap frequency counters that don't have a high stability option (quartz oven) fitted.

The next screenshot shows the same test with a 200MHz signal, which vastly exceeds the bandwidth of the SDS1104X-E. Yet it is nicely displayed (at a significantly reduced amplitude) and the automatic measurement is spot-on with now 5 digits, whereas the trigger frequency counter reads low again by some 5ppm. Nevertheless a very respectable result!

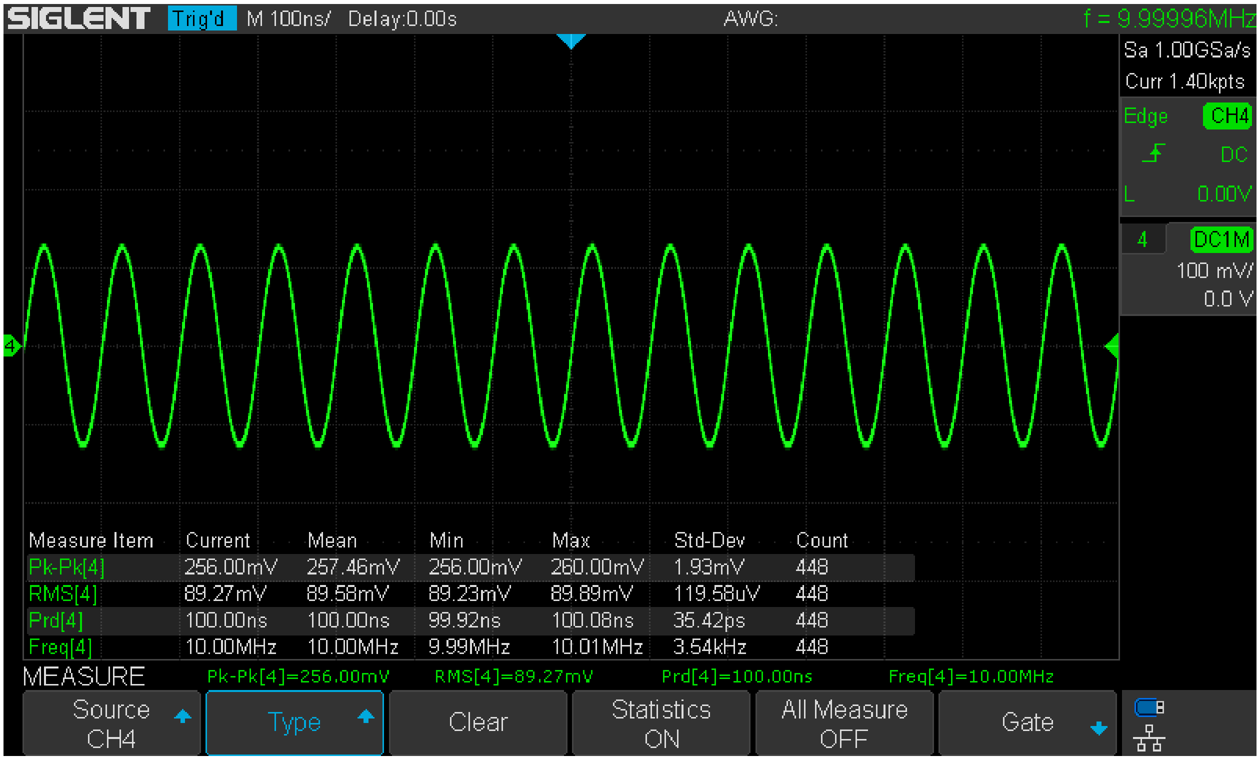

SDS1104X-E_CNT_10MHz

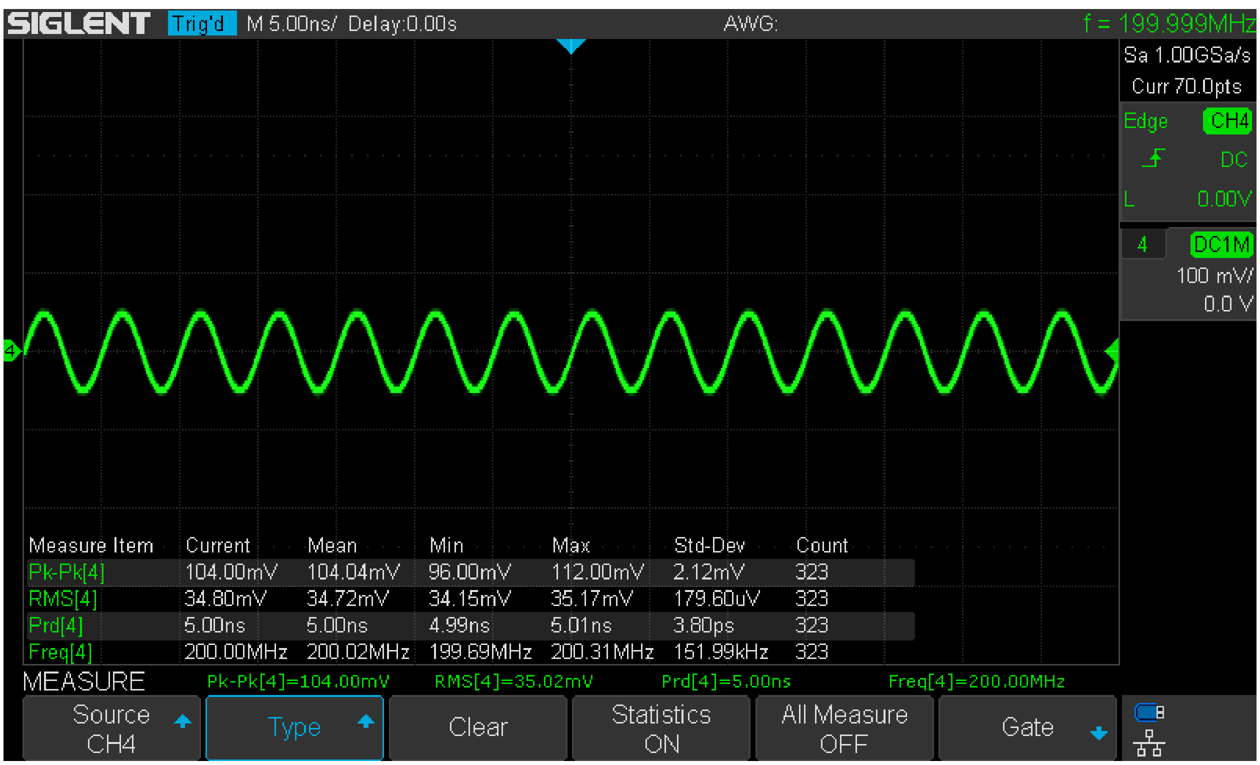

SDS1104X-E_CNT_200MHz

The last screenshot demonstrates a low frequency measurement. Since the trigger frequency counter does not work below 10Hz, the input signal was set to 15Hz to ensure smooth operation.

The readout of the trigger frequency counter doesn't look bad at all and an error of just 20ppm certainly appears acceptable. But the screenshot doesn't tell the whole story, as the frequency display isn't very stable and hardly more than 4 digits are actually usable at frequencies that low.

Automatic measurements on the other hand are spot-on again – at only 4 digits resolution, that is.

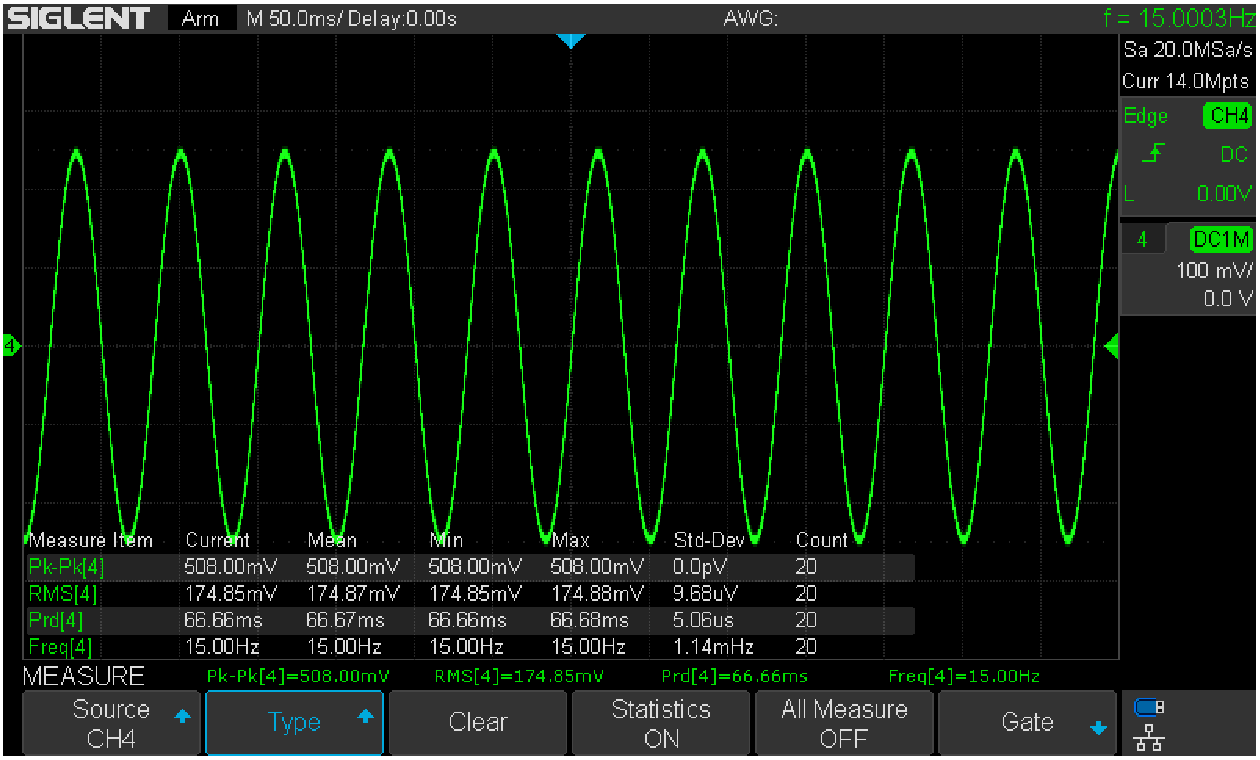

SDS1104X-E_CNT_15Hz

The situation improves quickly as the frequency goes up and at 1kHz nearly full accuracy and stability can be obtained. The graph below shows the frequency measurement error of both the counter and the automatic measurement for a number of input frequencies from 10Hz up to 50MHz. As can be seen from the results, the error reaches a first minimum at 10Khz and does not exceed 5ppm for any frequencies above. The 2nd screenshot above has already confirmed that for a 200MHz input signal.

It should be noticed, that this graph shows the maximum error, i.e. the maximum deviation from the actual frequency when the display fluctuates because of poor stability at low frequencies. In practical use we can watch the trigger frequency counter for a short while and do some averaging in our heads in order to obtain much more accurate results than what the graph below indicates.

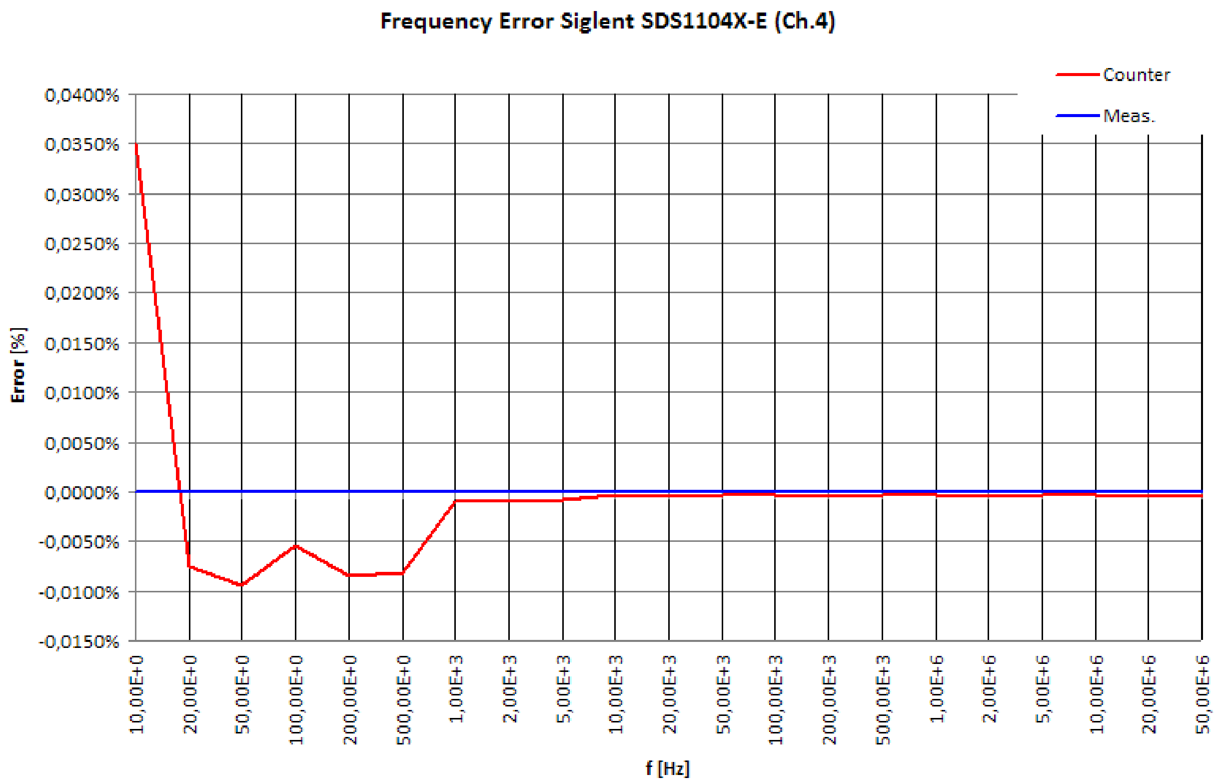

Counter Accuracy Graph

Trigger Rate

By now, the hype of high waveform update rates seems to have calmed down a bit; yet the simple truth remains: a DSO with reasonably fast trigger rate is a joy to use and provides useful additional information.

The actual waveform update rate is affected by so many things like number and combination of active analog and digital channels, display options and background tasks like automatic measurements. Because of the sheer number of combinations on a 4 channel MSO, tests have been limited to one and two channels in one group and all combinations of interpolation x or sin(x)/x and display mode Dots or Vectors. The MSO option was not yet available at the time of testing anyway.

The test results are summarized in the first table below. All measurements have been made with maximum memory length of 14Mpts per channel group enabled.

As it turns out, reconstruction has no effect on the trigger rate, but display mode does. This is not actually a problem, since Dots mode is fine for most applications and consequently Vector mode is rarely needed.

When the 2 channels are activated in different channel groups, the single channel waveform update rate is halved. Speed wise, this is often a better option than using two channels of the same group.

Automatic measurements may decrease the waveform update rate up to -30%.

The second table shows the result for all four acquisition modes with and without Fast Acquisition enabled exemplarily at the two timebase settings with the fastest waveform update rates, i.e. 5ns/div and 50ns/div. It turns out, that Average and Eres are hardware accelerated and only moderately slower than Normal and Peak Detect. Without hardware acceleration, they are even slightly faster.

Trigger: Ch.1		Measurements: OFF							
		Waveforms / second (Dots)				Waveforms / second (Vectors)			
		Ch. 1 only		Ch. 1 + 2		Ch. 1 only		Ch. 1 + 2	
Timebase	In Freq.	x	sin(x)/x	x	sin(x)/x	x	sin(x)/x	x	sin(x)/x
1 ns	50 MHz	6090	6090		3760		6085		3347
2 ns	50 MHz	9844	9844		3049		9400		3000
5 ns	20 MHz	34215	34215		12966		15486		7339
10 ns	10 MHz	12891	12891		17254		12143		8032
20 ns	5 MHz		13430		6470		12544		6247
50 ns	2 MHz		107694		20577		20948		8528
100 ns	1 MHz		19115		36542		18221		10462
200 ns	500 kHz		13362		9049		13089		8751
500 ns	200 kHz		8875		5900		8850		5875
1 µs	100 kHz		7288		3842		7288		3842
2 µs	50 kHz		5082		2851				
5 µs	20 kHz		2280		1512				
10 µs	10 kHz		1289		818				
20 µs	10 kHz		694		471				
50 µs	10 kHz		297		198				
100 µs	10 kHz		148		99				
200 µs	10 kHz		74		49				
500 µs	10 kHz		12,4		16,5				
1 ms	10 kHz		12,4		9,9				

Dots	sin(x)/x	Normal		Peak Detect		Average 16		Eres 2Bit	
Timebase	In Freq.	Acqu. Fast	Acqu. Slow	Acqu. Fast	Acqu. Slow	Acqu. Fast	Acqu. Slow	Acqu. Fast	Acqu. Slow
5 ns	20 MHz	34250	24,8	34250	24,8	27300	29,7	27300	29,7
50 ns	2 MHz	107694	24,8	107694	24,8	85982	29,7	85982	29,7

Trigger Rate

Another interesting – and even more important – aspect is the blind time of the DSO acquisition. This is simply expressed as the percentage of the total acquisition time where no sample data is collected.

For example, at a timebase of 100µs/div in single channel mode we get 148 waveforms per second. Since the screen is 14 divisions wide, a single waveform equals 1.4ms, hence 148 waveforms are equivalent to 148 x 1.4ms = 207.2ms. This in turn means that during one second, only 207.2ms worth of data is actually captured. This is just 20.72% of the total time, thus resulting in 79.28% blind time.

So even for a fast DSO with fairly high waveform update rates, the blind time is still substantial. The table below shows the blind time for all the previous measurements. Please note that even the fast waveform update rate of 107k at 50ns/div still leaves a blind time of 92.461%. While this might appear frustrating, there are still slower scopes around and if we look at the waveform update rate, it should immediately become obvious that for a DSO with twice as many waveforms per second at a certain time base also the “non-blind time” is doubled. It should also be noticed, that even if waveform update rates were identical, a scope with more horizontal divisions will have less blind time.

Trigger: Ch.1		Measurements: OFF							
		Waveforms / second (Dots)				Waveforms / second (Vectors)			
		Ch. 1 only		Ch. 1 + 2		Ch. 1 only		Ch. 1 + 2	
Timebase [s/div]	In Freq. [Hz]	Wfm/s	BT [%]	Wfm/s	BT [%]	Wfm/s	BT [%]	Wfm/s	BT [%]
1E-9	50E+6	6090	99,991%	3760	99,995%	6085	99,991%	3347	99,995%
2E-9	50E+6	9844	99,972%	3049	99,991%	9400	99,974%	3000	99,992%
5E-9	20E+6	34215	99,760%	12966	99,909%	15486	99,892%	7339	99,949%
10E-9	10E+6	12891	99,820%	17254	99,758%	12143	99,830%	8032	99,888%
20E-9	5E+6	13430	99,624%	6470	99,819%	12544	99,649%	6247	99,825%
50E-9	2E+6	107694	92,461%	20577	98,560%	20948	98,534%	8528	99,403%
100E-9	1E+6	19115	97,324%	36542	94,884%	18221	97,449%	10462	98,535%
200E-9	500E+3	13362	96,259%	9049	97,466%	13089	96,335%	8751	97,550%
500E-9	200E+3	8875	93,788%	5900	95,870%	8850	93,805%	5875	95,888%
1E-6	100E+3	7288	89,797%	3842	94,621%	7288	89,797%	3842	94,621%
2E-6	50E+3	5082	85,770%	2851	92,017%				
5E-6	20E+3	2280	84,040%	1512	89,416%				
10E-6	10E+3	1289	81,954%	818	88,548%				
20E-6	10E+3	694	80,568%	471	86,812%				
50E-6	10E+3	297	79,210%	198	86,140%				
100E-6	10E+3	148	79,280%	99	86,140%				
200E-6	10E+3	74	79,280%	49	86,280%				
500E-6	10E+3	12,4	91,320%	16,5	88,450%				
1E-3	10E+3	12,4	82,640%	9,9	86,140%				

Blind Time

Measurements

Traditionally, oscilloscopes were not expected to be useful for highly accurate precision measurements. Back in the days when the screen graticule was the only means for measuring signal parameters, the accuracy of reading alone was worse than 1% full scale - this has slightly improved with digital cursor readout and changed completely with automatic measurements. Now the reading accuracy is not an issue anymore, and the time accuracy (X-axis in Y-t mode) has improved significantly by several orders of magnitude with digital scopes, where the timing is determined by a digital clock coming from a crystal rather than a free-running RC oscillator that used to produce the ramp for horizontal deflection of the CRT. Even the amplitude accuracy (Y-axis) can be quite good in modern high resolution oscilloscopes with analog to digital converters using more than the traditional 8 bits.

The Siglent SDS1104X-E uses a reasonable stable TCXO to generate the clock of which the horizontal timebase is derived. More important is the effective sample rate, which causes an uncertainty of at least ±1ns at 1GSa/s and quite obviously the uncertainty becomes correspondingly higher as the sample rate decreases by enabling both channels in a group and/or lowering memory depth and/or timebase.

Automatic Measurements

Back in the days of analog oscilloscopes, the screen grid was pretty much the only aid for measurements. Characterizing a signal used to be a time consuming and error prone task, yet some measurements like RMS required additional equipment, at least for non-textbook waveforms. Nowadays we can utilize a bunch of automatic measurements, which can make life so much easier – as long as they work reliably and provide reasonably accurate results. This is to be examined in the following sections.