Siglent SDS1104X-E Review

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.

Time Resolution

Accurate time measurements shouldn't be a problem for a modern DSO that has its sample clock derived from a TCXO (Temperature Compensated Quartz Oscillator). Yet there are many DSOs which use just the screen buffer for calculating the measurement results. This is fast and easy to implement, but requires zooming into the waveform in order to get meaningful measurement results, just like with an ancient analog oscilloscope. In contrast to this, the Siglent X-E series DSOs use the full acquisition memory to calculate the measurements – and this is still fast, thanks to the high processing power.

Consider a pulse train with fast transitions, maybe also narrow pulse widths, but relatively slow repetition rate, like a high resolution PWM signal. If we want to see the repetition rate, duty cycle and transition times at a glance, a scope that only uses the screen buffer just won't be up to the task.

Here we have a 100Hz PWM signal with a duty cycle of 90%, displayed at 1ms/div in order to see one full period:

SDS1104X-E_PWM_100Hz_90%_M1ms

In the screenshot above, we can clearly see the 100Hz repetition rate – the trigger frequency counter is off by 0.001Hz, but automatic period measurement is spot-on, showing 10.00ms. We can also accurately measure the duty cycle of 90.00% and despite the slow timebase of 1ms/div, we still get a reasonably close measurement of the transition times in the realm of single-digit nanoseconds.

For a high resolution PWM, we might use more than 8 bits and therefore be able to adjust the duty cycle in steps smaller than 0.4%. The following example shows such a situation with a duty cycle of 0.12%. The duty cycle measurement is still spot-on, but is somewhat compromised by its limited resolution. As it is now, it can handle just about 13 bits – with only one digit more, it could be used for 16 bits PWM as well. This is one of several instances, where the display resolution of the automatic time measurements should be enhanced and I hope Siglent will improve this eventually.

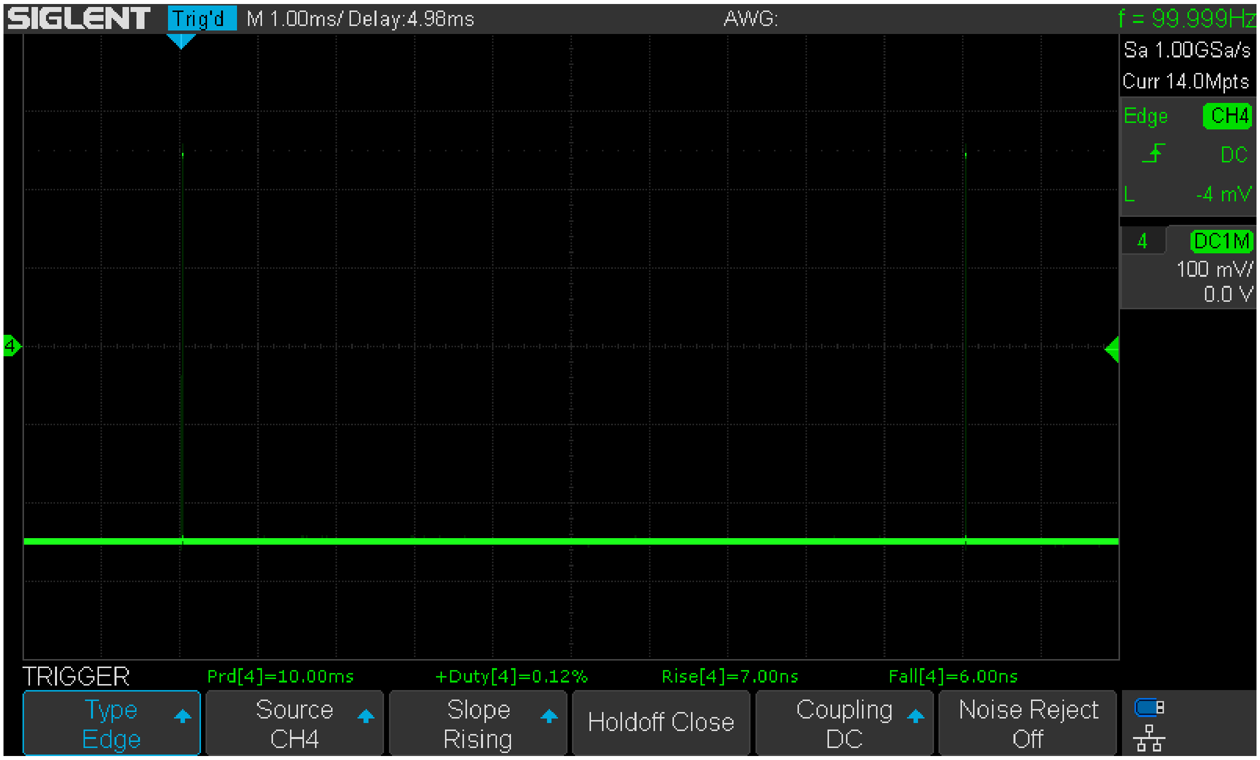

SDS1104X-E_PWM_100Hz_0.12%_M1ms

Now a duty cycle of only 0.02%:

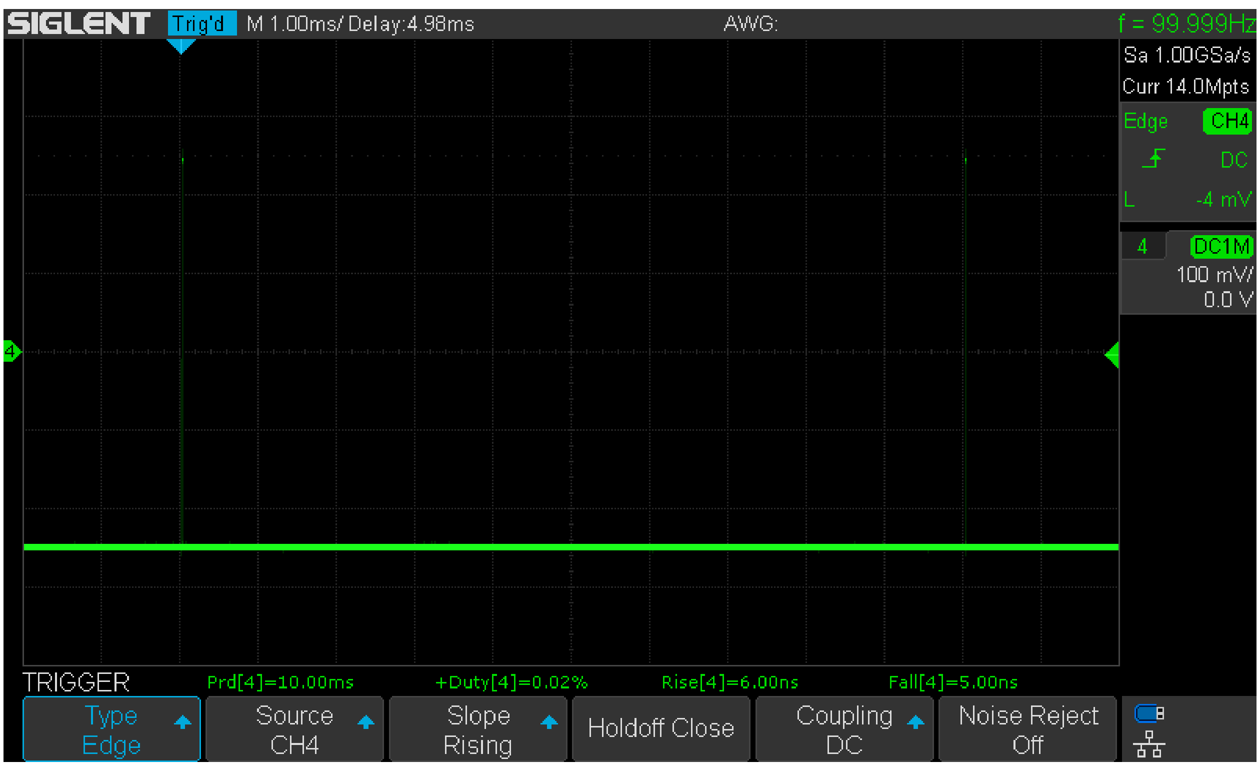

SDS1104X-E_PWM_100Hz_0.02%_M1ms

Finally, let's zoom into the waveform 5000 times (200ns/div) to see just the pulse:

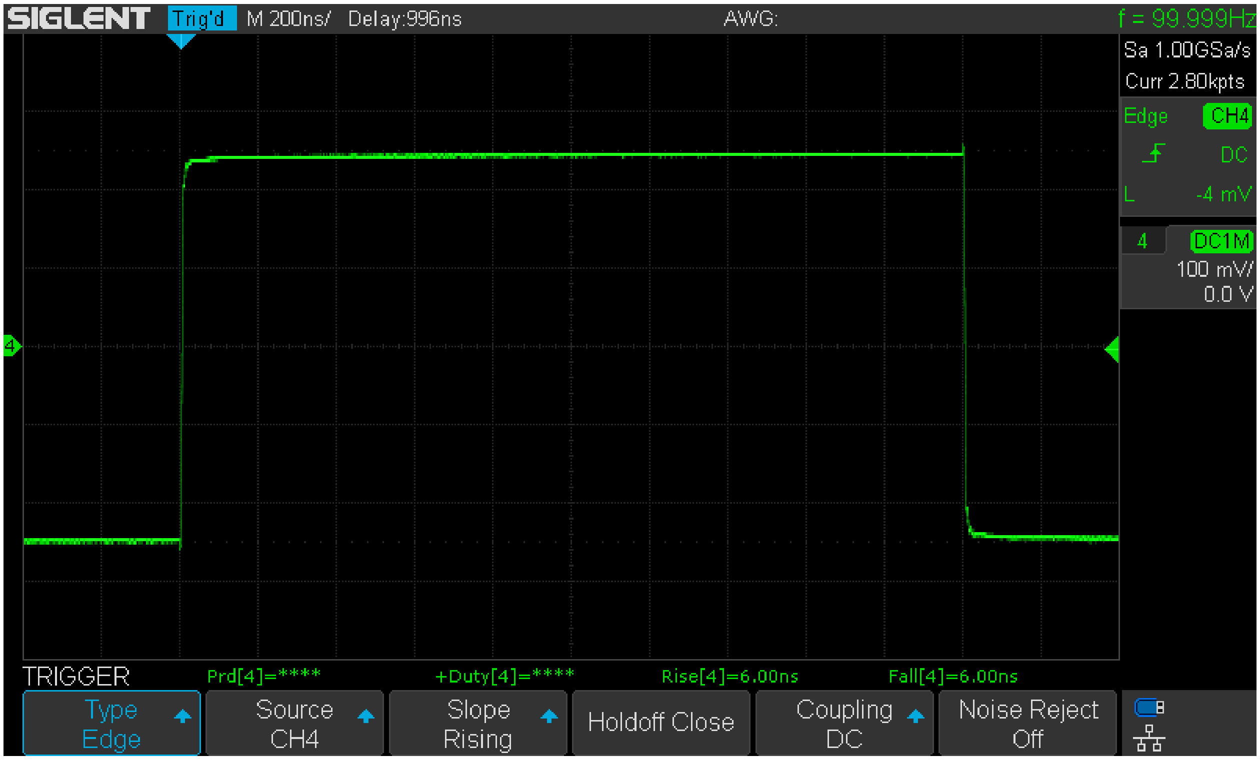

SDS1104X-E_PWM_100Hz_0.02%_M200ns

The trigger frequency counter still displays the (almost) correct frequency, but automatic measurements for period and duty cycle cannot provide any results, as they need at least one full signal period to work. Transition times still haven't changed, so there was no need to zoom anyway.

Sensitivity

Not a real test case, just a rather curious observation that should be documented here. A 250MHz sine is far above the capabilities of the 100MHz SDS1104X-E and even with the amplitude of 150mVpp it is just adding a little noise to the un-triggered trace line. Consequently, the trigger frequency counter shows nothing, but the automatic measurements still work and the frequency measurement is not far off.

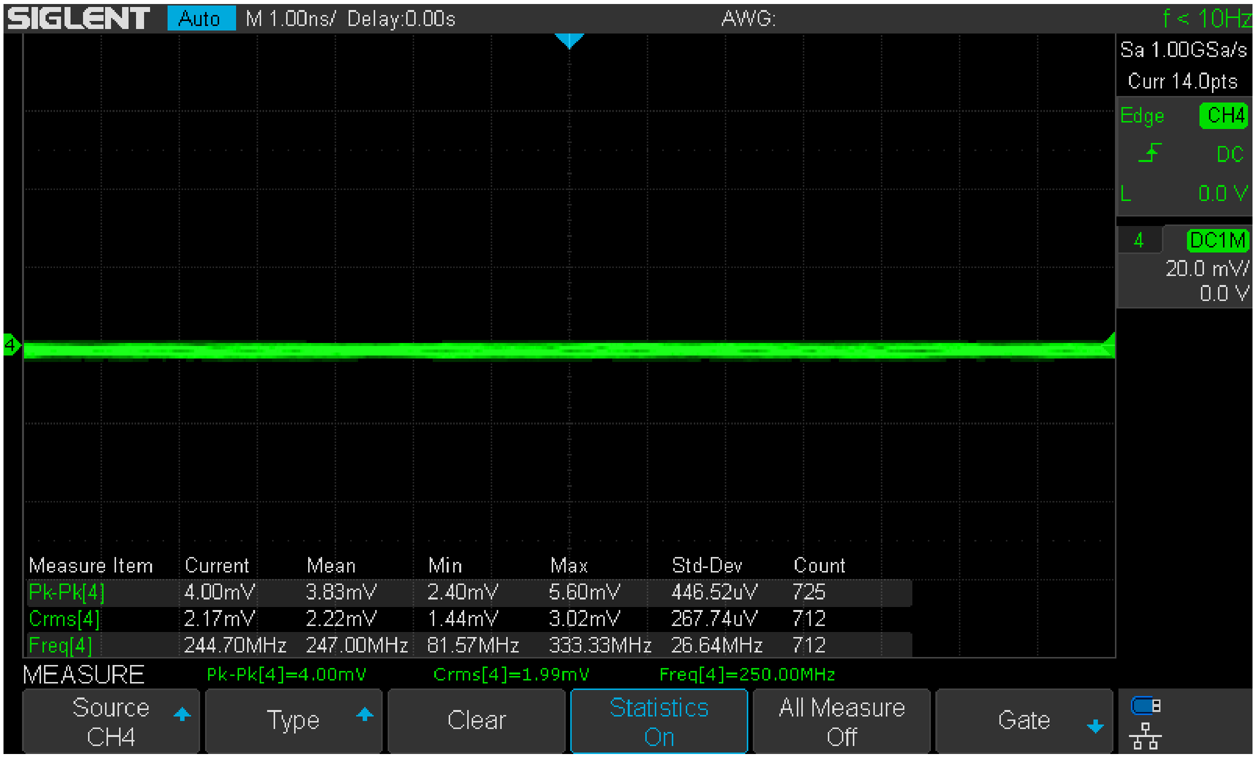

Sine_250MHz_150mVpp_Stat

Speed

We don't only want measurements to have high time resolution and good accuracy, they should also be fast. Speed matters, especially for dynamic signals. The SDS1104X-E certainly delivers in this regard, see the following examples.

The screenshots below show two modulated waveforms.

- First a 10MHz carrier, 50% amplitude modulated with 1kHz.
- Second a 10MHz carrier, with 1kHz frequency modulation and 1MHz peak deviation.

When looking at the min/max values in the measurement statistics, the modulation parameters can be easily confirmed. So measurements are indeed fast enough to characterize even modulated signals with good accuracy within acceptable time. Especially for FM, accuracy is impressive even after only 40 measurements (or less than 3 seconds).

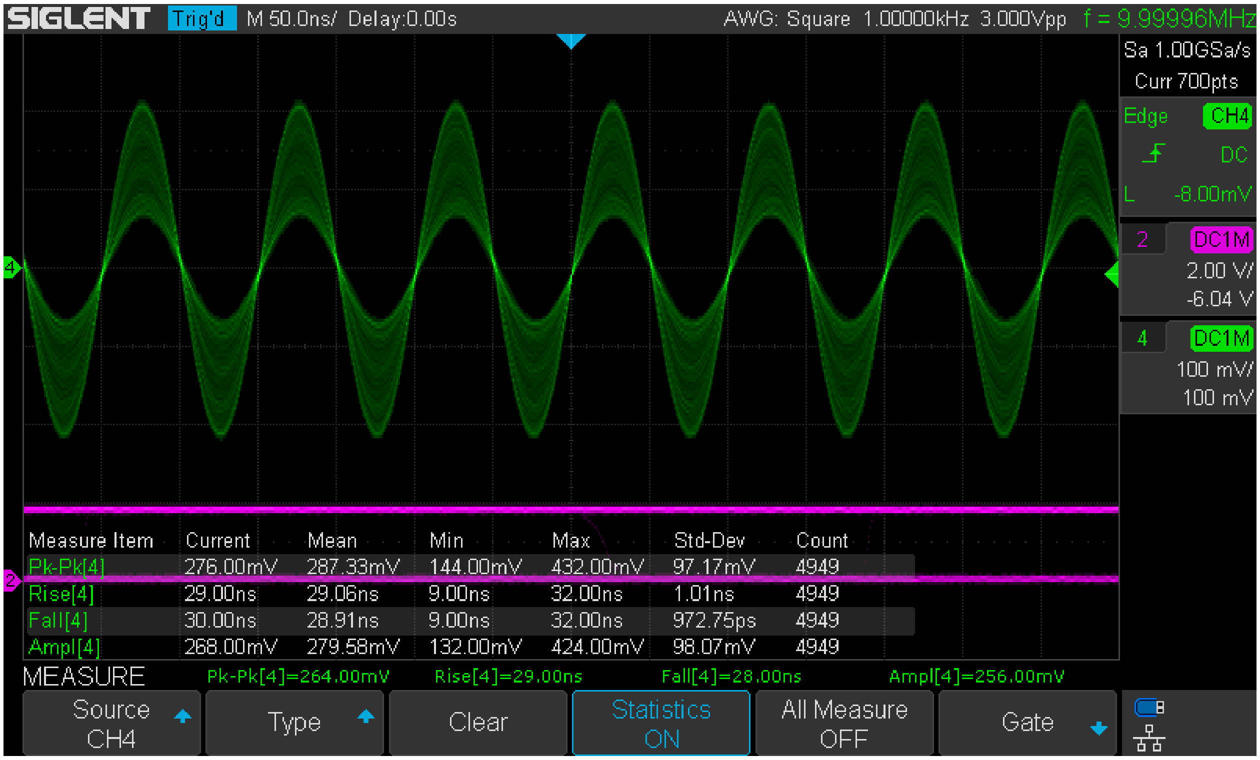

Sine_10MHz_AM1kHz_50%_M50ns_Dots

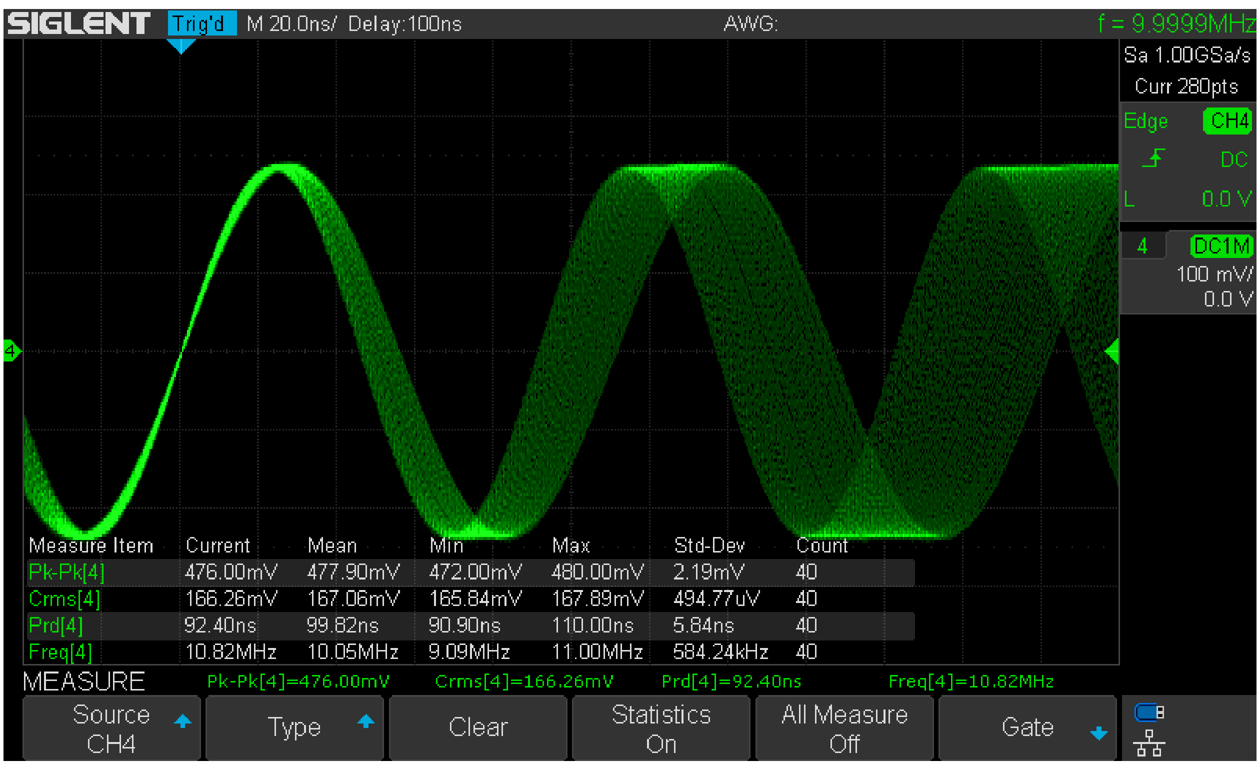

Sine_10MHz_FM1kHz_D1MHz_M20ns_Dots

Standard Measurements

Generally, we can have a maximum of 4 simultaneous measurements displayed at the bottom of the screen and one more in the statistics. Sometimes this is not enough, so we are given the opportunity to display all standard measurements together at the same time by selecting “All Measure ON”.

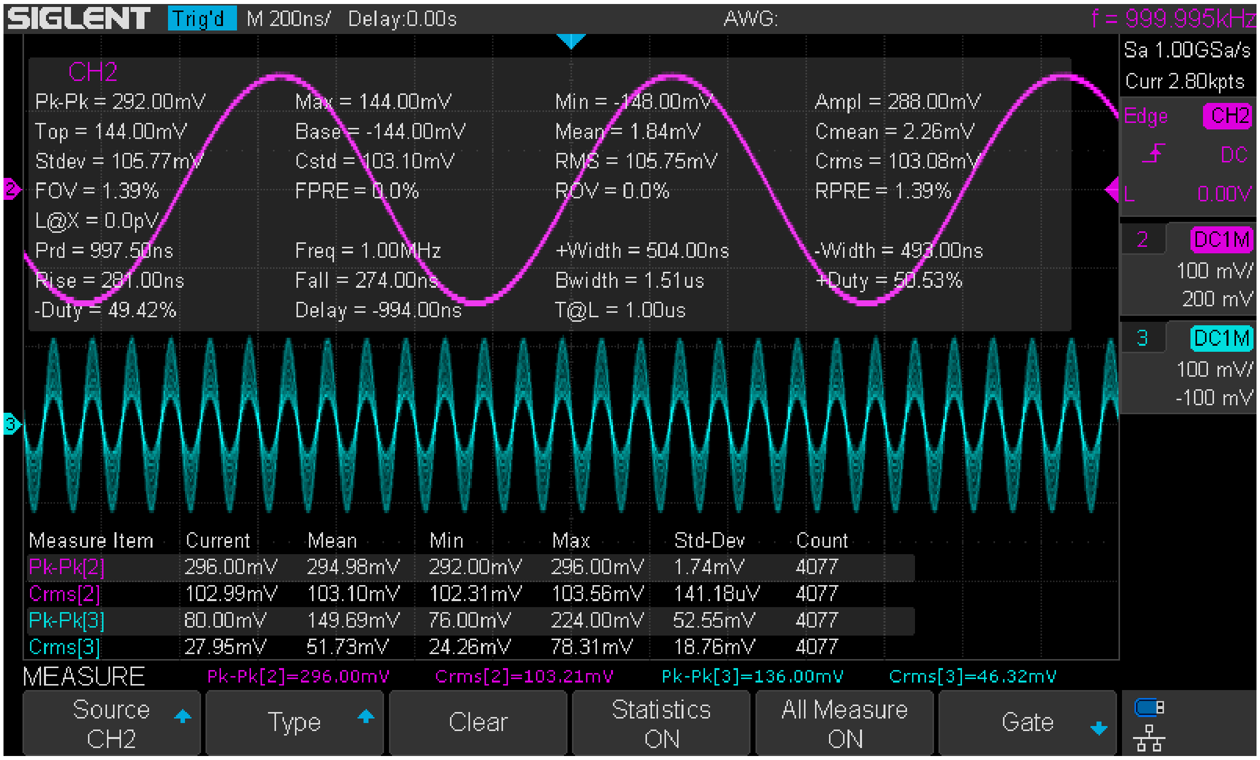

SDS1104X-E_all_measure

This is of course only possible for a single channel at a time, yet we can have all measurements for one channel plus 4 more for other channels at the bottom of the screen. The example above shows all measurements for channel 2 and the bottom line holds two Ch.2 measurements (which is of course redundant in this case, but we can have statistics for these), plus two Ch.3 measurements.

The big question remains, how accurate all these measurements are. There are no explicit specifications for that. In principle, we can use the specifications for amplitude accuracy as well as the sample clock as a basis for estimating the expected accuracy. I have done this exactly and run a number of tests in order to verify the accuracy of all measurements, one sample of such a test protocol is shown in the table below.

The resulting Error field is color coded:

Error is less than ¼ of what was to be expected.
Error is less than what was to be expected.
Error is more than what was to be expected.
Error is more than two times what was to be expected.

As can be seen, all measurements are pretty accurate except for L@X, which is a bit of a mystery anyway. According to the manual, this should be “The voltage value of the trigger point”, which in my understanding has to be pretty close to the trigger level. On the other hand, this wouldn’t make much sense since we certainly know the trigger level anyway and need no measurement for that. So it’s most likely a translation issue and/or lack of a clear description what this measurement really does and my expectations were just wrong.

Base Freq [Hz]	10,00E+3		Period [s]	100,00E-6
Base Ampl [V]	1,00E+0		Duty [%]	10,00%
Base Width [s]	10,00E-6		Base [-]	990,00E-3
Pulsewidth [s]	1,00E-6		Pulse [-]	10,00E-3
Pulseheight [V]	-100,00E-3		Duty [%]	1,00%
Transition [s]	12,00E-9		Duty+ [%]	10,00%
Pulsepos. [0/1]	1		Duty- [%]	89,00%
Samplerate (Hz)	1,00E+9			
Trigger [V]	128,00E-3			
Timebase [s]	50,00E-6		Periods [-]	6
Gain [V/div]	200,00E-3			
Gain Error	3,00%			
Offset Error [V]	2,00E-3			

Trigger on falling Edge!

Measurement	Expected	Measured	Error [%]	Error Limit [%]
Pk-Pk [V]	1,10E+0	1,11E+0	0,91%	3,00%
Max [V]	500,00E-3	496,00E-3	-0,80%	3,40%
Min [V]	-600,00E-3	-616,00E-3	2,67%	3,33%
Ampl [V]	1,00E+0	1,01E+0	1,00%	3,00%
Top [V]	500,00E-3	496,00E-3	-0,80%	3,40%
Base [V]	-500,00E-3	-512,00E-3	2,40%	3,40%
Mean [V]	-401,00E-3	-409,00E-3	2,00%	3,50%
Cmean [V]	-401,00E-3	-409,00E-3	2,00%	3,50%
Stdev [V]	305,78E-3	301,90E-3	-1,27%	3,65%
Cstd [V]	305,78E-3	301,90E-3	-1,27%	3,65%
RMS [V]	501,10E-3	511,50E-3	2,08%	3,40%
Crms [V]	501,10E-3	511,50E-3	2,08%	3,40%
FOV [%]	0	0,79	0,79%	3,00%
FPRE [%]	0	0	0,00%	3,00%
ROV [%]	0	0	0,00%	3,00%
RPRE [%]	10	10,32	0,32%	3,00%
L_At_X [V]	128,00E-3	64,00E-3	-50,00%	4,56%
Prd [s]	100,00E-6	100,00E-6	0,00%	0,01%
Frequ [Hz]	10,00E+3	10,00E+3	0,00%	0,01%
Pos_Width [s]	10,00E-6	10,00E-6	0,00%	0,01%
Neg_Width [s]	90,00E-6	90,00E-6	0,00%	0,01%
Rise [s]	12,00E-9	13,00E-9	8,33%	8,34%
Fall [s]	12,00E-9	13,00E-9	8,33%	8,34%
Bwidth [s]	610,00E-6	610,00E-6	0,00%	0,01%
Pos_Duty [%]	10,00%	10,00%	0,00%	0,01%
Neg_Duty [%]	90,00%	90,00%	0,00%	0,01%
Delay [s]	-310,00E-6	-310,00E-6	0,00%	0,01%
T_At_L [s]	290,00E-6	290,00E-6	0,00%	0,01%

SDS1104X-E_RUS-RPRE

Channel Delay Measurements

There is a second group of measurements, summarizing all the possible time measurements between different channels. Unfortunately, there is no option to display them all at once, so the screenshot below gives just an example of 4 selected measurements.

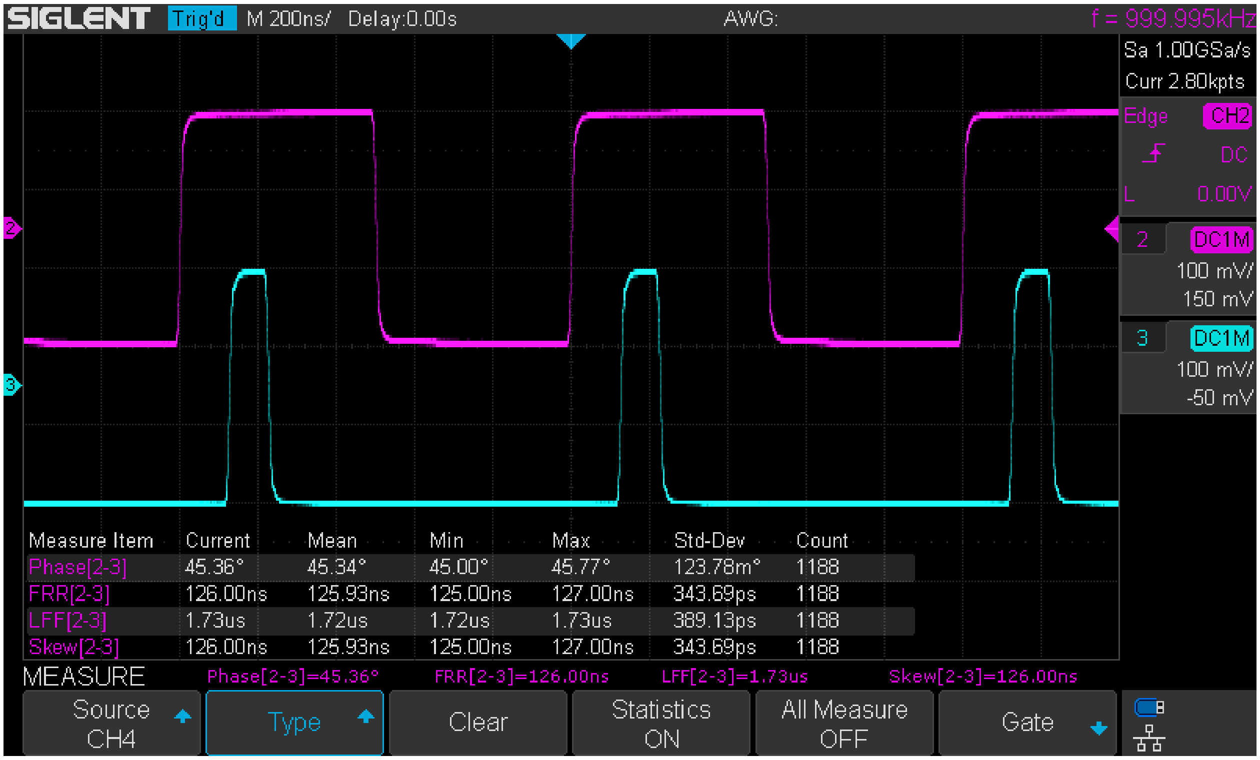

SDS1104X-E_Ch_Delay

A Freq [Hz]	1,00E+6	**A Period [s]**	1,00E-6
A Duty [%]	50,00%	**A Pulse [s]**	500,00E-9
B Duty [%]	10,00%	**B Pulse [s]**	100,00E-9
B Phase [°]	-45,00	**B Delay [s]**	125,00E-9
Samplerate (Hz)	1,00E+9		
Timebase [s]	200,00E-9	**Periods [-]**	2

Trigger on falling Edge!

Measurement	**Expected**	**Measured**	**Error [%]**	**Error Limit [%]**
Phase [°]	45,00	45,33	0,73%	1,61%
FRR [s]	125,000E-9	125,920E-9	0,74%	1,61%
FRF [s]	225,000E-9	225,940E-9	0,42%	0,90%
FFR [s]	625,000E-9	625,550E-9	0,09%	0,33%
FFF [s]	725,000E-9	726,000E-9	0,14%	0,29%
LRR [s]	2,125E-6	2,130E-6	0,24%	0,10%
LRF [s]	2,225E-6	2,230E-6	0,22%	0,10%
LFR [s]	1,625E-6	1,630E-6	0,31%	0,13%
LFF [s]	1,725E-6	1,730E-6	0,29%	0,13%
Skew [s]	125,000E-9	126,370E-9	1,10%	1,61%

SDS1104X-E_Ch_Delay_Accuracy

Again, the expected error limits and accuracies for all possible measurements have been determined, checked and summarized in the table above. Since this is all about time measurements, we could expect a very high accuracy, so it might be all the more surprising to find some error results highlighted in red. But it's not a faulty measurement, just insufficient resolution together with some rounding, which makes a basically correct result look bad. For some weird reason, measurements are limited to 3 digits in this area, thus turning e.g. 2.125µs into 2.13µs and making the error much bigger than necessary. One more time I do hope Siglent will take notice and improve that in a future firmware release.

Gated Measurements

Sometimes we don't want to measure over the whole record length and need a means to limit measurements to just a part of it. This is called a gated measurements and the Siglent SDS1104X-E provides dedicated X-Cursors for this.

The most popular example for gating is the integrating amplitude measurements, such as RMS and Mean. They can only work correctly on integer multiples of the signal period. Since the screen width (= record length) often does not meet this criterion, the automatic measurements already include a special set for these cases: it's the Cycle measurements Cycle Mean, Cycle Stdev and Cycle RMS, so we don't need to set up a gate for these common tasks.

The next example uses the PWM signal with 10% duty cycle again and at 1ms/div the record length is 140% of the signal period, which certainly isn't an integer multiple. So there's no surprise that the Cycle measurements yield different results compared to the normal ones:

SDS1104X-E_PWM_100Hz_10%_NoGate

As was to be expected, the error would be huge if we forget to use the Cycle measurements. 186.18mV RMS vs. 155.78mV Cycle RMS and 72.43mV Mean vs. 51.34mV Cycle Mean.

Since this chapter is about gated measurements, let's see if we can get the correct results with normal measurements when using a properly defined gate. The RMS and Mean measurements are now pretty close to Crms and Cmean from the previous screenshot:

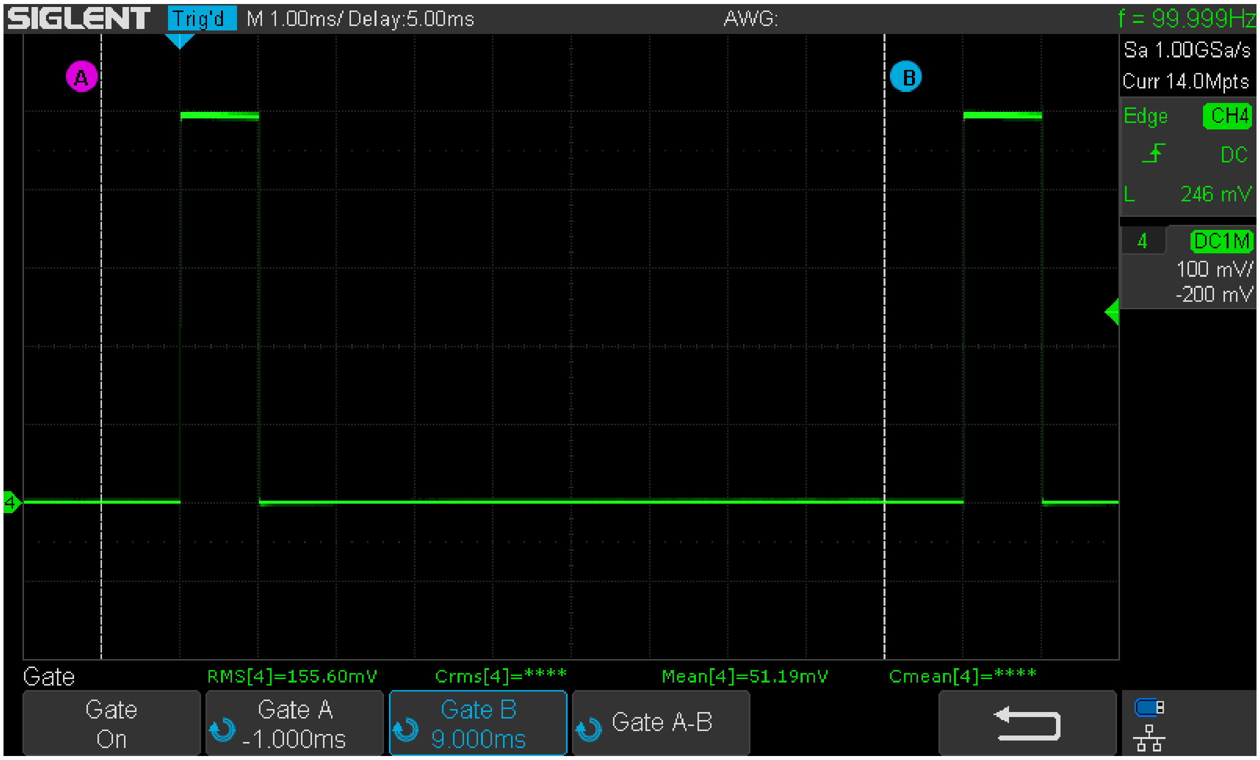

SDS1104X-E_PWM_100Hz_10%_Gate10ms

Why don't we define the gate from rising edge to rising edge?

Well, a proper cycle should include exactly one rising and one falling edge. We should avoid setting the gate border at a fast signal transition, at least when using time bases so slow that we cannot actually see the transition, because it's nothing but a vertical line. For accurate measurements, the signal level needs to be the same at both gate borders and we cannot guarantee this for transitions in the realm of single-digit nanoseconds by setting a gate with 20µs resolution like in the example above.

Currently, we don't get any results from the dedicated Cycle measurements, because the gate limits the view and makes it impossible to detect two rising (or falling) edges, which would be essential for recognizing a full signal cycle. That's not a problem here, as Cycle measurements don't make sense anyway – after all we have set the gate for exactly one cycle just to make the Cycle measurements superfluous.

Just for fun, we can still have both types of measurements at the same time by making the gate a tad wider and shift it a bit, so that it includes two rising edges. This will introduce a small error for the normal measurements (because this now sees a little bit more than just one cycle), but lets us compare both approaches within one screenshot. As can bee seen, RMS and Crms as well as Mean and Cmean measurements are now very similar as was to be expected.

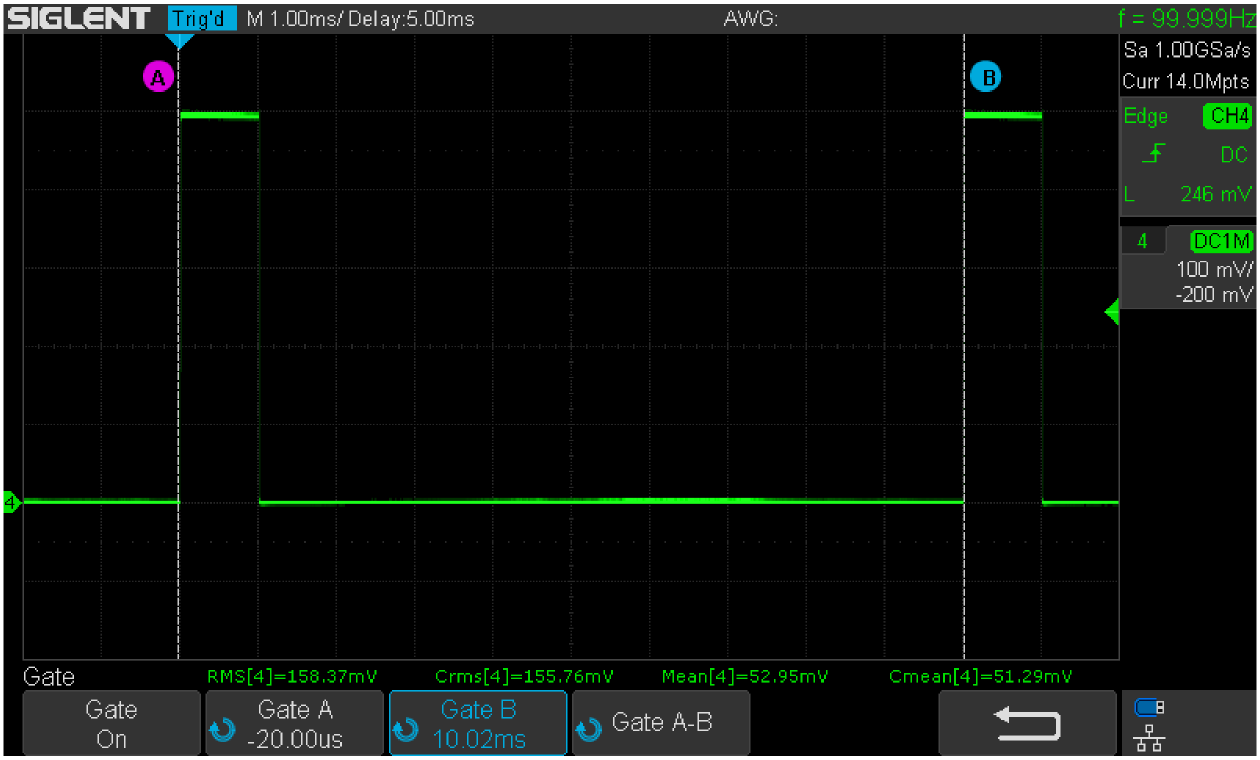

SDS1104X-E_PWM_100Hz_10%_Gate10.04ms

Cursors

The main applications for Cursor measurements are:

- No automatic measurement available for a particular measurement task.
- Visual highlighting of a particular measurement.
- Measuring and comparing time and level at an arbitrary point of the displayed waveform(s).

Manual

Manual cursor measurements are pretty straightforward, so they are dealt with here just for the sake of completeness.

The following screenshots demonstrate the use of X, Y and X/Y cursors for measuring some signal properties where no corresponding automatic measurements exist. Furthermore, the use of cursors can make a screenshot more descriptive, as it becomes immediately clear what exactly has been measured.

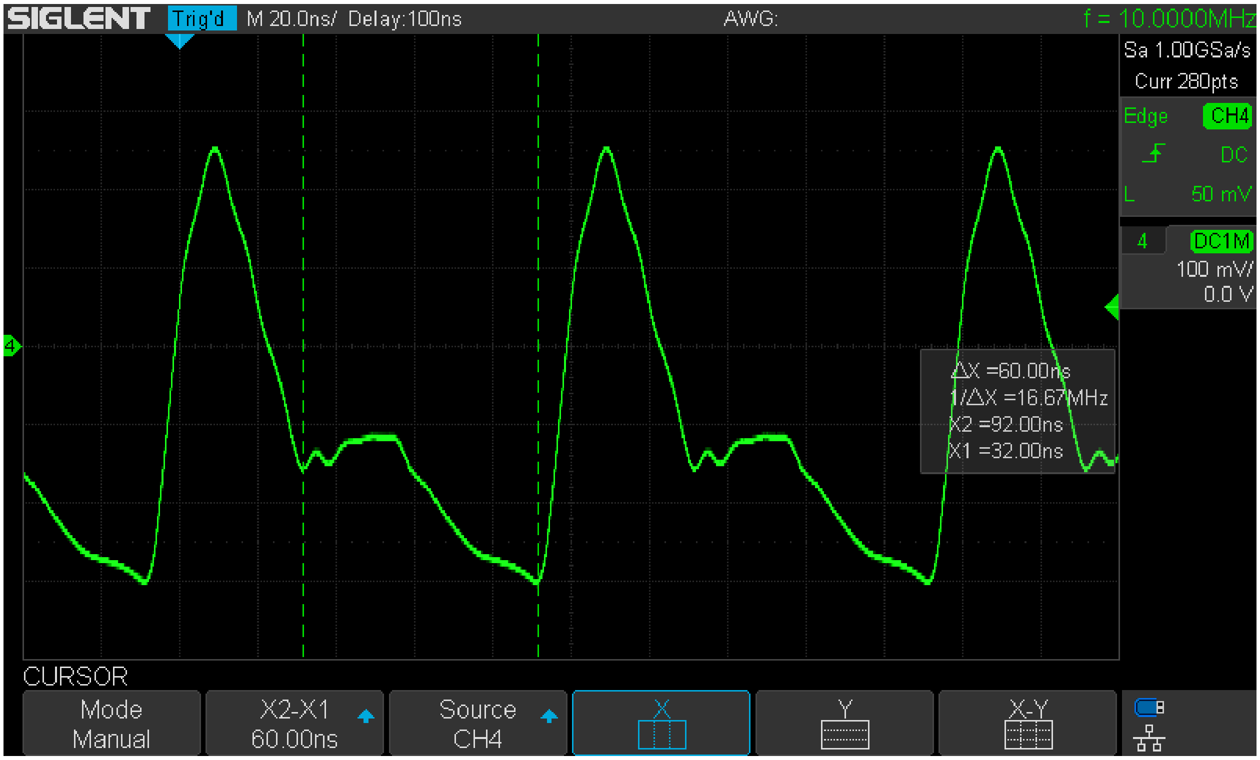

Cursors_Man_X

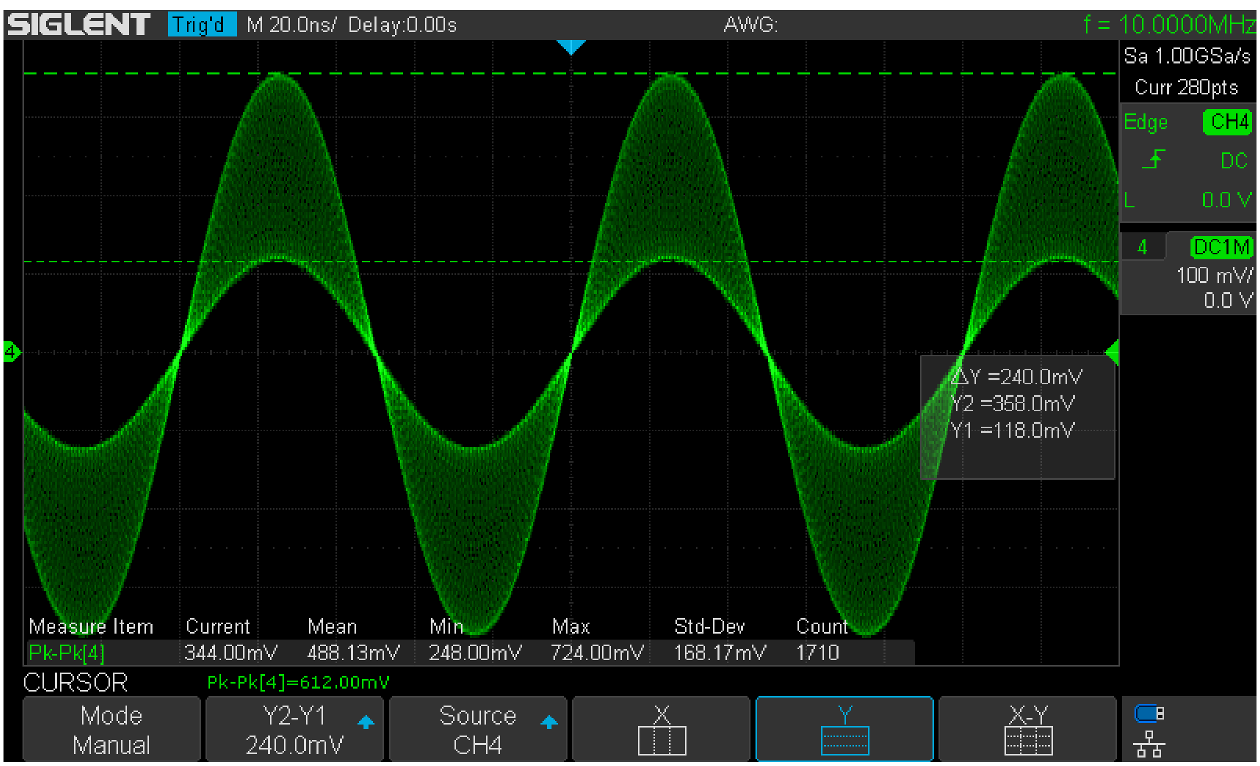

Cursors_Man_Y

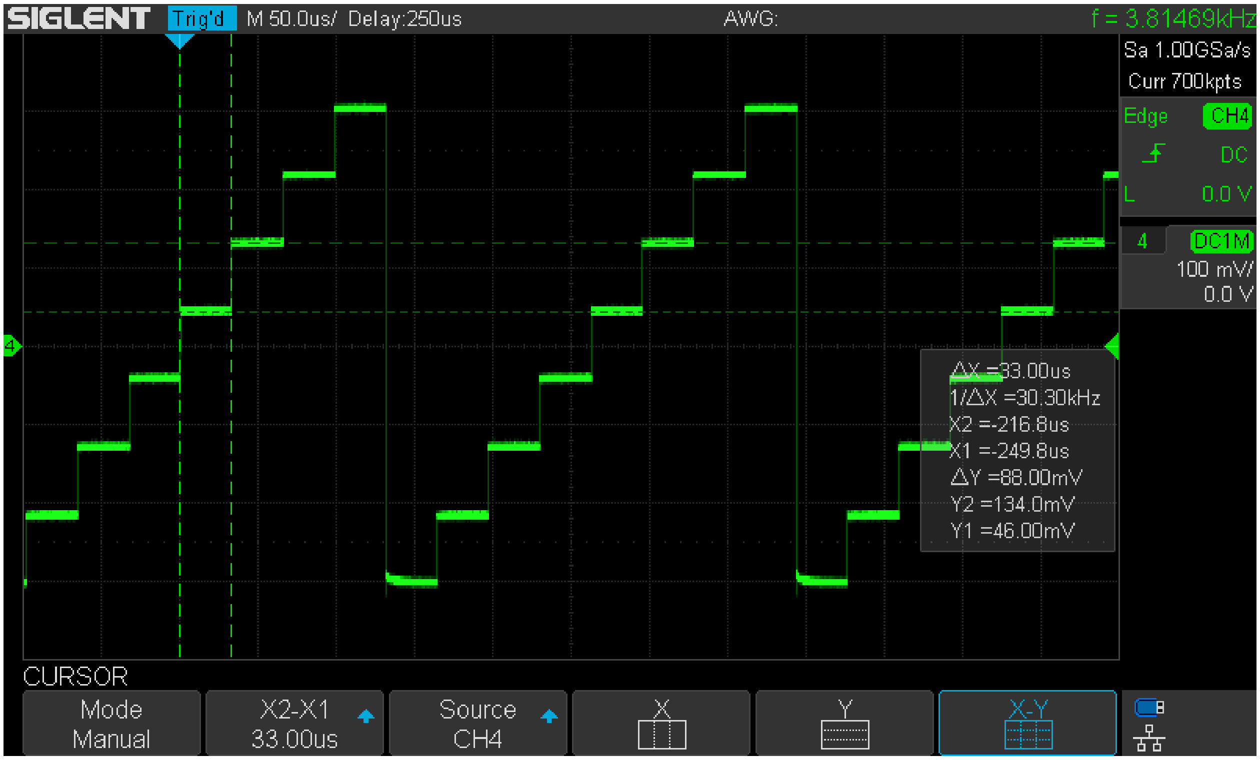

Cursors_Man_XY

Tracking

Tracking cursors let us view the time and level at any point of the displayed waveform. At the same time, this can be compared to any other point, even on a different waveform on a different channel.

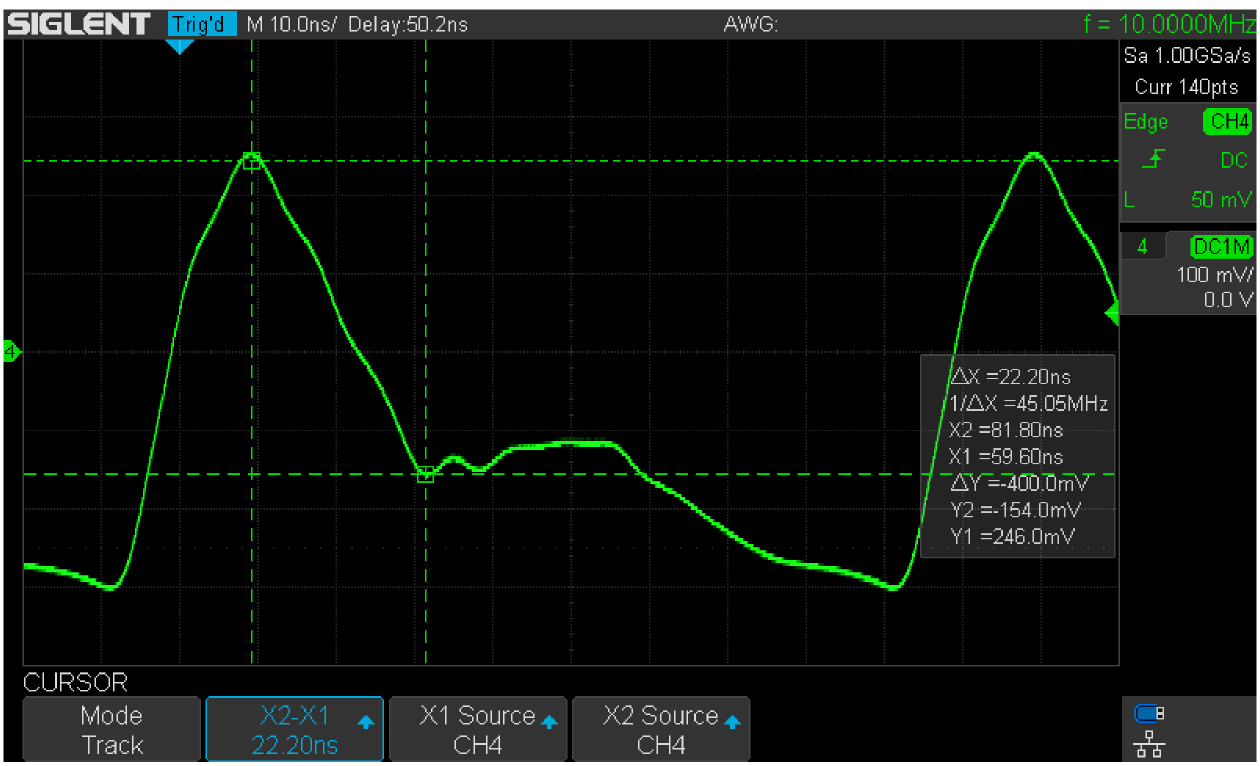

Cursors_Aut_Ch4

Measurement across channels is demonstrated by the example below, were we are just measuring the time delay and by adjusting the cursors so that the level difference is zero we ensure an accurate measurement.

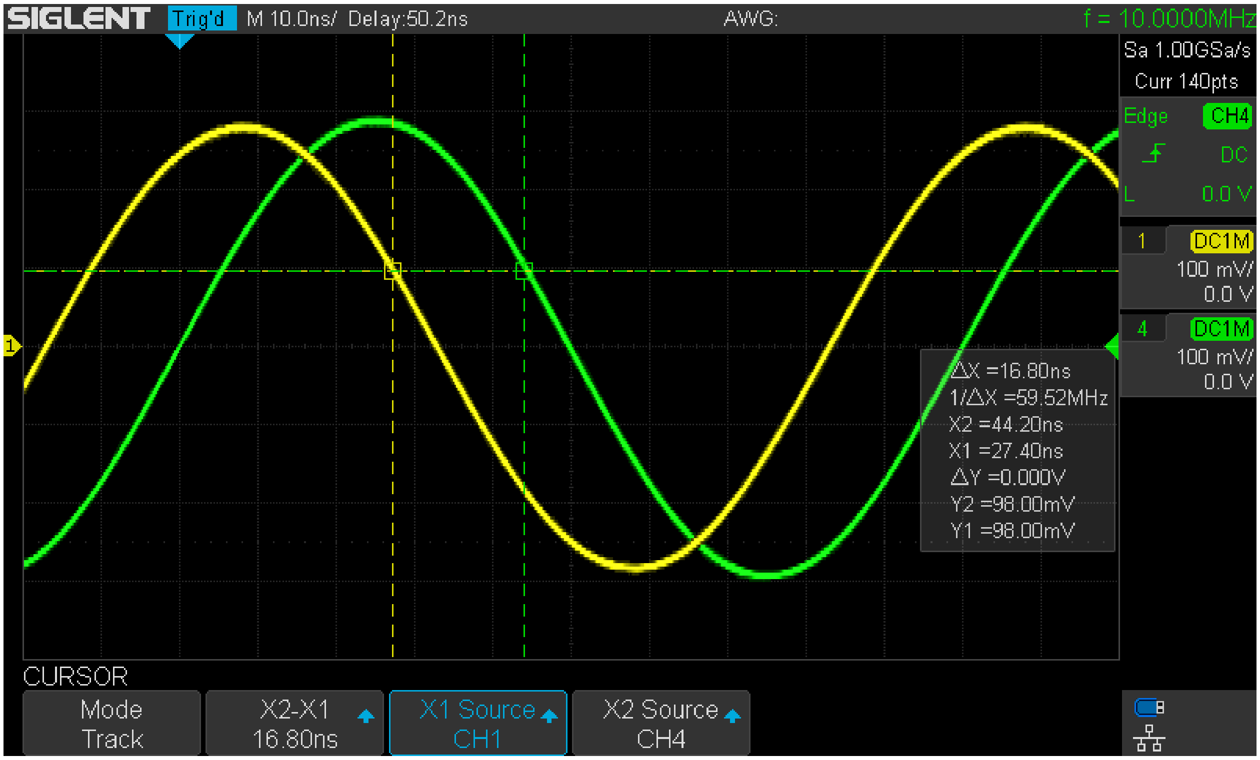

Cursors_Aut_Ch14

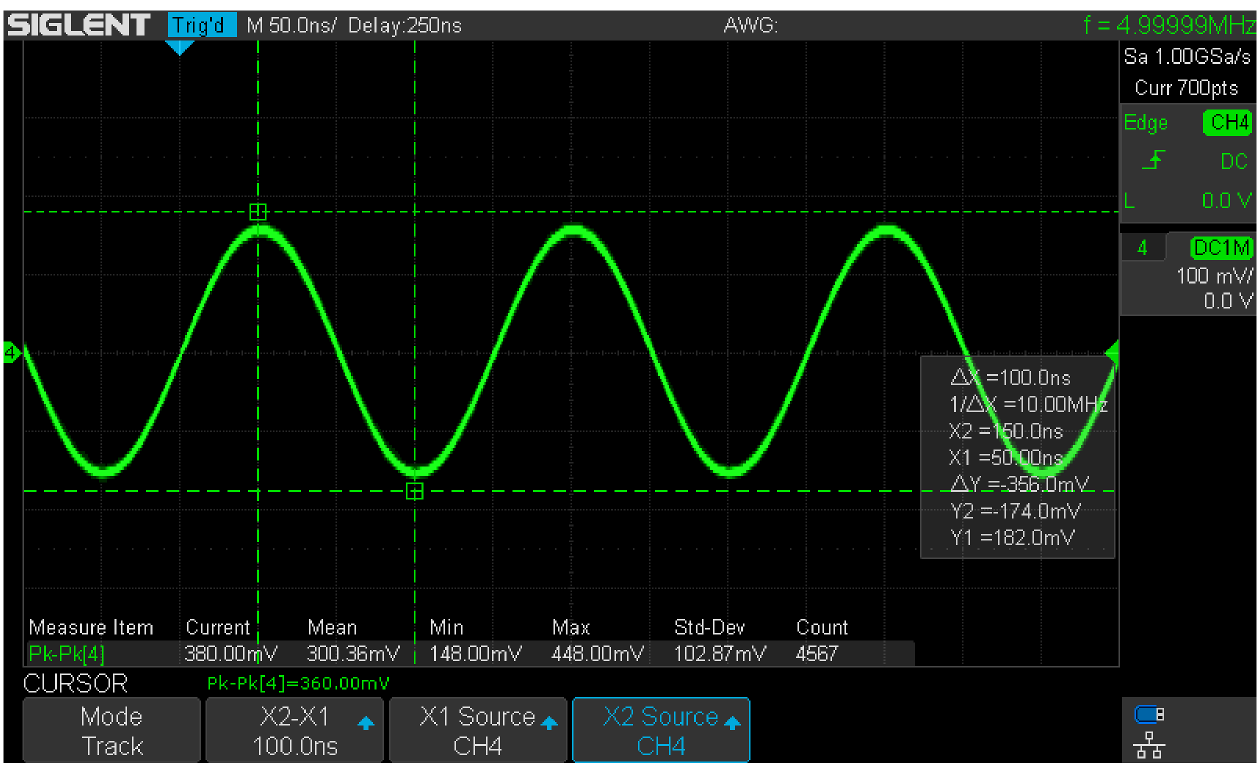

Cursors_Aut_Ch4_Mod_0.5Hz

The last example demonstrates tracking cursors on a dynamic signal – a 50% amplitude modulated waveform in this case. Modulation frequency is quite low at 0.5Hz, still the cursors lag behind a bit. Plotting the cursors on the screen is not at all optimized for speed, even though the actual measurements in the cursor box update much faster than that; so it's really just the cursor display being rather slow.